PHENIX CDR Update

An Experiment to be Performed at the
Brookhaven National Laboratory
Relativistic Heavy Ion Collider

November, 1993
(Printed August, 1994)

Spokesperson:

Shoji Nagamiya,
Columbia University

Project Director:

Samuel H. Aronson,
Brookhaven National Laboratory

Deputy Project Director:

Glenn R. Young,
Oak Ridge National Laboratory

Project Engineer:

Leo Paffrath,
Brookhaven National Laboratory

TAC Review Copy

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Editor's Note

Since the production of the PHENIX Conceptual Design Report Update for the Technical, Cost and Schedule Review in November 1993, the project and the collaboration have evolved to meet the funding and management challenges inherent in such a large scientific enterprise. PHENIX earned technical approval at the November 1993 review; however, the collaboration did not receive cost and schedule approval at that time. Further difficult deferments in the baseline detector were required of the collaboration before PHENIX gained final approval as a construction project at a March 1994 Cost and Schedule Review.

This is the *same* document that was prepared for the first review, and, as such, it represents the status of the PHENIX detector and collaboration in November 1993. After all the reviews, I had hoped to continue editing the CDR Update, but, as many of you know, I ran out of time. My thanks to everyone who made this document possible.

Walter L. Kehoe
12 August 1994

Contents

1	Overview	1-1
1.1	Introduction	1-1
1.2	Physics	1-1
1.3	Detector	1-3
1.4	Strategy	1-5
2	Baseline Detector	2-1
2.1	Introduction	2-1
2.2	Coverage and Geometry	2-1
2.3	Magnets	2-3
2.4	Central Arms	2-3
2.4.1	Charged Particle Tracking and Momentum Measurement	2-4
2.4.2	Electron and Photon Energy Measurements	2-4
2.4.3	Particle Identification	2-4
2.5	Muon Arm	2-7
2.6	Inner Detectors	2-8
2.7	Trigger and Data Acquisition	2-9
2.7.1	Front-end Electronics	2-9
2.7.2	Level-1 Trigger	2-10
2.7.3	Level-2 and Level-2.5 Triggers	2-10
2.7.4	Event Builder, Data Archiving and On-Line Computing System	2-11
2.7.5	Off-line Computing	2-11
2.8	Facility Engineering and Detector Integration	2-12
3	Physics Capabilities	3-1
3.1	Overview	3-1
3.2	Simulations: Strategy and New Results	3-5
3.2.1	Secondary Electrons on RICH	3-6
3.2.2	MVD Induced Conversion Electrons	3-10
3.2.3	Mass Resolution in the Central Arm	3-15
3.2.4	Dimuon Measurements	3-16
3.3	Revision for Photon Simulations	3-18
3.4	QED Background Studies in the BB	3-18

4	Magnets	4-1
4.1	Overview	4-1
4.2	Changes	4-1
4.3	Central Magnet	4-7
4.3.1	Inner and Outer Coils	4-8
4.3.2	Pole Pieces and Support Structure	4-10
4.3.3	Assembly and Installation	4-11
4.3.4	Testing	4-12
4.3.5	Cost and Schedule	4-12
4.4	Muon Magnet	4-14
4.4.1	Coil	4-14
4.4.2	Pole Pieces and Support Structure	4-15
4.4.3	Assembly and Installation	4-16
4.4.4	Testing	4-16
4.4.5	Cost and Schedule	4-17
4.5	Muon Identifier Steel	4-19
4.5.1	Description	4-19
4.5.2	Fabrication	4-19
4.5.3	Assembly and Installation	4-19
4.5.4	Cost and Schedule	4-20
5	Inner Detectors	5-1
5.1	Silicon System Overview	5-1
5.2	Silicon System Changes	5-2
5.3	Mechanical	5-3
5.3.1	Silicon detectors	5-3
5.3.2	Support Structures	5-4
5.3.3	Cabling	5-8
5.3.4	Cooling	5-8
5.3.5	Assembly	5-10
5.4	Electronics	5-11
5.4.1	Front end	5-11
5.4.2	Hybrid circuit	5-12
5.5	Silicon System Performance Issues	5-13
5.5.1	Electronics Location	5-13
5.5.2	Silicon Coverage	5-14
5.6	Cost and Schedule	5-17
5.7	Beam-beam Counter Overview	5-17
5.8	Beam-beam Counter Changes	5-18
5.9	Mechanical	5-19
5.9.1	Detector Elements	5-19
5.9.2	BB Location	5-20
5.9.3	Electronics	5-21
5.10	Beam-beam Counter Performance Issues	5-23

5.11	Cost and Schedule	5-26
6	Tracking System	6-1
6.1	Overview	6-1
6.2	Changes	6-1
6.3	Drift Chambers	6-2
6.3.1	Mechanical Design	6-2
6.3.2	Electronics	6-4
6.3.3	Prototype Operation and Performance	6-5
6.3.4	Cost and Schedule	6-6
6.4	Pad Chambers	6-7
6.4.1	Mechanical Design	6-7
6.4.2	Electronics	6-9
6.4.3	Calculations	6-11
6.4.4	Cost and Schedule	6-12
6.4.5	Pixel Pad Cathode Option	6-15
6.5	Time Expansion Chamber	6-19
6.5.1	Mechanical Design	6-20
6.5.2	Electronics	6-21
6.5.3	Operation and Performance	6-23
6.5.4	Cost and Schedule	6-24
6.6	System Performance	6-25
6.6.1	Momentum and Mass Resolution	6-25
6.6.2	Pattern Recognition	6-27
7	Ring Imaging Cherenkov Detector	7-1
7.1	Overview	7-1
7.2	Changes	7-2
7.3	Mechanical	7-2
7.3.1	Global Geometry and Optics	7-2
7.3.2	Radiator Gas Vessel	7-3
7.3.3	Mirror System	7-4
7.3.4	Gas Handling System	7-6
7.3.5	Focal Plane Detector	7-7
7.4	Electronics	7-9
7.4.1	Front-end Electronics	7-9
7.4.2	High Voltage Power Supply	7-10
7.5	Performance	7-10
7.6	Cost and Schedule	7-10
8	Time-of-Flight System	8-1
8.1	Overview	8-1
8.2	Changes	8-1
8.3	Mechanical Design	8-2

8.4	Electronics	8-4
8.5	Performance	8-4
8.6	Cost and Schedule	8-8
9	Electromagnetic Calorimetry	9-1
9.1	Overview	9-1
9.2	Changes	9-2
9.3	Lead-scintillator Calorimeter	9-2
9.3.1	Supermodule and Wall Design	9-2
9.3.2	Module Design	9-4
9.3.3	Calibration System	9-6
9.3.4	Photomultipliers and High Voltage Supply	9-8
9.3.5	Cost and Schedule	9-11
9.4	Lead-glass Calorimeter	9-11
9.4.1	Monitoring System Design Criteria	9-11
9.4.2	LED for Monitoring System	9-12
9.4.3	Pulse Generators	9-12
9.4.4	Light Distribution System	9-13
9.4.5	Reference Detector	9-13
9.4.6	Cost and Schedule	9-13
10	Muon Arm	10-1
10.1	Overview	10-1
10.2	Changes	10-1
10.3	Acceptance and Coverage	10-3
10.4	Layout and Mechanical	10-3
10.4.1	Interface to Magnet	10-3
10.4.2	Specialized Absorber Sections	10-4
10.4.3	Muon Arm Tracking Chambers	10-5
10.4.4	Muon Identifier	10-7
10.5	Electronics	10-9
10.5.1	Front-end	10-9
10.5.2	Trigger	10-9
10.6	Performance and Backgrounds	10-11
10.6.1	Absorber	10-12
10.6.2	Tracking and Spectrometer	10-12
10.7	Cost and Schedule	10-12
11	On-line System	11-1
11.1	Overview	11-1
11.1.1	General On-line Scheme	11-3
11.1.2	Event Rates at RHIC	11-6
11.2	Changes	11-7
11.2.1	Front-end Electronics Changes	11-7

11.2.2	Trigger Changes	11-9
11.2.3	DACQ Changes	11-9
11.2.4	On-Line Computing System Changes	11-10
11.3	Front-end Electronics	11-10
11.3.1	Subsystem FEE and Data Volumes	11-10
11.3.2	Development of Functional Blocks	11-13
11.3.3	Control and Bus Lines	11-14
11.4	Trigger System	11-16
11.4.1	Level-1 Trigger	11-16
11.5	Data Acquisition System	11-35
11.5.1	Overview of Data Acquisition System	11-35
11.5.2	Data Formats	11-40
11.5.3	DCM Cards	11-44
11.5.4	Event Builder	11-49
11.6	On-line Computing System	11-55
11.6.1	Overview	11-55
11.6.2	Changes	11-57
11.6.3	On-line Framework	11-58
11.6.4	Generic Services	11-62
11.6.5	Development plan	11-64
11.7	Cost and Schedule	11-66
12	Off-line System	12-1
12.1	Overview	12-1
12.2	Requirements	12-2
12.3	Computing Models	12-3
12.3.1	The RHIC Computing Facility	12-3
12.3.2	Main Code Modules of the PHENIX Off-line System	12-6
12.4	Data Storage and Management	12-9
12.4.1	Event Storage and Management	12-9
12.4.2	Detector Database	12-11
12.4.3	Information Database	12-12
12.4.4	Data Structures	12-12
12.5	Software	12-13
12.5.1	Code Development Environment	12-13
12.5.2	High Level Program Packages	12-16
12.5.3	Core Modules	12-16
12.6	Cost and Schedule	12-17
13	Systems Engineering and Integration	13-1
13.1	Overview	13-1
13.2	Changes	13-1
13.3	Detector Systems and Central Arm Carriages	13-2
13.4	Facilities	13-4

13.4.1 Counting House	13-4
13.4.2 Utilities	13-4
13.5 Integration	13-6
13.6 Installation	13-6
13.7 Test and Commissioning	13-9
13.8 Cost and Schedule	13-9
14 Project Management	14-1
14.1 Introduction	14-1
14.2 Personnel	14-1
14.2.1 Scientists and Engineers	14-1
14.2.2 Technical and Administrative Personnel	14-1
14.3 Cost Estimate and Schedule	14-2
15 Schedules and Profiles	15-1
15.1 Summary Cost Estimate	15-1
16 Upgrade Options	16-1
16.1 Introduction	16-1
16.2 Tracking TRD	16-1
16.3 High Resolution Photon Detector	16-2
16.3.1 Overview	16-2
16.3.2 Changes	16-2
16.4 Hadron-blind Detector	16-5
16.4.1 Introduction	16-5
17 Collaboration and Organization	17-1
17.1 PHENIX Collaboration	17-1
17.2 Management Structure	17-1

PHENIX Experiment at RHIC

Participants

November, 1993

University of Alabama, Huntsville, Alabama

J. C. Gregory, T. Shiina, Y. Takahashi*, M. Tocci

Bhabha Atomic Research Centre, Bombay, India

R. K. Choudhury, S. K. Gupta, S. S. Kapoor*

Brookhaven National Laboratory, Upton, NY 11973

S. H. Aronson, G. David, M. Fatyga, S. Gavin, W. Guryn*, J. Harder,
A. Hershcovitch, B. Johnson, S. Kahn, W. L. Kehoe, E. Kistenev, C. Knapp,
P. Kroon, Y. Makdisi, M. Montag, J. Negrin, P. O'Connor, E. O'Brien, L. Paffrath,
S. Rankowitz, T. K. Shea, M. J. Tannenbaum, S. White, D. Whitehouse, C. L. Woody

University of California - Riverside, Riverside, CA 92521

J. Chang, Sun-Yiu Fung*, Ju-Hwan Kang, R. Seto, Quian Zhu

China Institute of Atomic Energy, Beijing, P. R. China

Xixiang Bai, Kexing Jing, Xiaoping Liu, Zuhua Liu, Yajun Mao, Benhao Sa,
Zuxun Sun*, Yude Wan, Zhongqi Wang, Jincheng Xu, Xiaoze Zhang, Yuming Zheng,
Shuhua Zhou

Chung-ang University, Seoul, 156-756, Korea

C. S. Lee*

Columbia University, Nevis Laboratories, Irvington, NY 10533

C. Y. Chi, B. Cole, J. Dodd, S. Nagamiya, M. Moulson, T. Nayak, G. Pettili,
W. Sippach, O. Vossnack, Fuqiang Wang, Yufeng Wang, Yuedong Wu, Xihong Yang,
W. A. Zajc*

Florida State University, Tallahassee, FL 32306

R. Chappell, A. D. Frawley*, S. Wang

Georgia State University, Atlanta, GA 30303

X. T. Liu, G. Petitt*, Z. Zhan

Hiroshima University, Saijo, Hiroshima, Japan

S. Esumi, H. Iwata, K. Kumagai, N. Maeda, O. Miyamura, A. Sakaguchi,
T. Sugitate, Y. Sumi*

Idaho National Engineering Laboratory, Idaho Falls, ID 83415

R. Aryaeinejad, R. DeVries, J. Cole*, M. Drigert

Institute of High Energy Physics, Academia Sinica, Beijing, P. R. China

Qi-ming Li, Jun-guang Lu, Man Wang, Ming-han Ye, Bin-yun Zhang, Wei-qin Zhao, Zhi-peng Zheng, Yu-can Zhu*

Institute of High Energy Physics, Protovino, Russia

V. V. Ammosov, A. Denisov, A. Durum, V. Gapienko, Yu. Galitsky, Yu. Gutnikov, A. Holodenko, A. Ivanilov, V. I. Kochetkov*, B. Korablev, A. Kozelov, V. V. Makeev, E. Melnikov, Yu. Mikhailov, V. Onuchin, Yu. Pischalnikov, Yu. Protopopov, V. I. Rykalin, K. Shestermanov, A. Starkov, D. Strustumov, A. Surkov, V. Zaetz, A. Zaichenkov,

Institute of Modern Physics, Academia Sinica, Lanzhou, P. R. China

Ruirong He, Genming Jin, Yixiao Luo, Baowan Wei, Hongfei Xi, Feng Ye, Wenlong Zhan*, Jingyu Zhang

Institute for Nuclear Research, Russian Academy of Sciences, Moscow, Russia

S. N. Filippov, Y. K. Gavrilov, F. F. Guber*, O. V. Karavichev, A. B. Kurepin, V. I. Rasin, A. I. Reshetin

Institute for Nuclear Study, University of Tokyo, Tanashi, Tokyo, Japan

Y. Akiba, H. Hamagaki*, S. Homma, H. Sako, M. Sekimoto

Institute of Theoretical and Experimental Physics - Moscow, Moscow, Russia

A. Bolozdynya, L. Gusev, A. Koutchenkov, A. Kovalenko, V. Lebedenko, V. Sheinknan, V. Tchernyshev, V. Vinogradov

Iowa State University and Ames Laboratory, Ames, IA 50011

L. A. Ewell, J. C. Hill, B. Libby, A. Petridis, D. Quingley, H. Skank, G. Sleege, W. Thomas, F. K. Wohn*

Joint Institute for Nuclear Research - Dubna, Dubna, Russia

S. V. Afanasiev, Yu. S. Anisimov, V. V. Arkhipov, V. V. Balashov, O. V. Egorov, A. F. Elishev, A. Yu. Isupov, A. N. Khrenov, V. A. Kuznetsov, A. G. Litvinenko, P. K. Maniakov, A. I. Malakhov, G. L. Melkumov, I. I. Miglina, A. S. Nikiforov, V. G. Perevozchikov, S. G. Reznikov, A. Yu. Semenov, A. B. Shabunov, V. A. Smirnov, V. A. Trofimov, V. V. Trofimov, P. I. Zarubin*,

KEK, Institute for High Energy Physics, Tsukuba, Japan

J. Chiba*, Y. Mori, M. Nomachi, O. Sasaki, H. Sakamoto, T. Shintomi, K. H. Tanaka

Korea University, Seoul, 136-701, Korea

S. J. Lee, S. K. Lee, G. Y. Lim, K. S. Sim*

Kurchatov Institute, Moscow, Russia

S. Belyaev, S. Fokin, M. Ippolitov, K. Karadjev, A. Lebedev, V. Manko*,
A. Nyanine, A. Tsvetkov, A. Vinogradov

Kyoto University, Kyoto, Japan

H. Enyo*, H. Kaneko

Lawrence Livermore National Laboratory, Livermore, CA 94550

J. Bowers, L. Hansen, D. Krofcheck, J. Luke, N. Namboordiri, C. Sangster, R. Soltz,
J. Thomas*, S. Tonse, R. Yamamoto

Los Alamos National Laboratory, Los Alamos, NM 87545

J. G. Boissevain, T. A. Carey, D. Clark, D. Fields, H. van Hecke, B. Jacak*,
J. Kapustinsky, W. W. Kinnison, M. Leitch, J. McClelland, P. L. McGaughey,
C. Morris, J. Moss, J. Simon-Gillo, W. E. Sondheim, J. P. Sullivan

Louisiana State University, Baton Rouge, LA 70803

P. N. Kirk*, Y. C. Wang, Z. F. Wang, Jiangchen Yu

Lund University, Lund, Sweden

L. Carlen, S. Garpman, H.-A. Gustafsson, A. Oskarsson*, I. Otterlund,
K. Soderstrom, E. Stenlund, T. Svensson

McGill University, Montreal, Quebec, Canada

J. Barrette, R. Fernholz, K. Filimonov, S. K. Mark*, L. Nikkinen, L. Normand,
A. Singh, F. Qi, M. Rosati, N. Starinsky

Massachusetts Institute of Technology, Cambridge, MA 02139

D. Akimov, M. Akopyan, A. Bolozdynya, M. Chen*, D. Churakov, A. Kovalenko,
V. Lebedenko, D. Simonichev, M. Smolin

University of Münster, Münster, Germany

C. Barlag, E. M. Bohne, A. Claussen, N. Heine, K.-H. Kampert, Th. Peitzmann,
R. Santo*

Nagasaki Institute of Applied Science, Nagasaki-shi, Nagasaki, Japan

K. Kimura, Y. Tanaka

National Institute of Radiation Science, Chiba-ken, Japan

E. Takada*

State University of New York - Stony Brook, Stony Brook, NY 11794

P. Braun-Munzinger, T. Hemmick, B. Hong, Y. Kwon, W. J. Llope, S. Panitkin,
M. Rao, J. Stachel*, S. Sedykh, N. Xu, Y. Zhang, C. Zou

Oak Ridge National Laboratory, Oak Ridge, TN 37831

G. Alley, T. Awes, C. Britton, M. S. Emery, Hee Kim, R. Maples, F. Obenshain*,
F. Plasil, K. Read, C. Q. Rennich, B. W. Reimer, S. Saini, M. Simpson, P. Stankus,
J. Walker, A. Wintenberg, G. R. Young

Peking University, Beijing, P. R. China

Shang-lian Bao, Jia-ehr Chen*, Dong-xing Jiang, Lun-cun Wei

St. Petersburg Nuclear Physics Institute, Gatchina, Leningrad, Russia

N. K. Abrosimov, V. Baublis, V. Ivochkin, M. Kann, A. Khanzadeev, B. Komkov,
L. Kochenda, A. Kudin, N. Miftakhov, V. Mulnikov, O. Novozhilov, E. Orishin,
V. Pylov, E. Roshin, A. Shetkovsky, E. M. Spiridenkov, O. Tarasenkova, I. Tkach,
L. Uvarov, A. Vorobyov*, E. Vznuzdaev, V. Yatsura

University of São Paulo, CP20516-01000, São Paulo, Brasil

O. Dietzsch*, N. Cansian de Silva, E. M. Takagui

Seoul National University, Seoul, 151-742, Korea

S. J. Chae, H. T. Chung, J. H. Ha, J. C. Kim*, Y. K. Kim, M. Y. Youn

Soong-Sil University, Seoul, 156-743, Korea

S. C. Jeong*

University of Tennessee, Knoxville, TN 37996

S. Garbett, Xiao-chun He, J. Kreke, S. Sorensen*

University of Tokyo, Tokyo, Japan

R. Hayano*, T. Ishikawa, K. Itahashi, H. Sakurai, K. Shigaki, H. Tamura

University of Tsukuba, Tsukuba, Japan

I. Arai, T. Ikeda, S. Kato, A. Kumagai, M. Kurata, K. Kurita, Y. Miake*, S. Sato,
K. Tomizawa, S. Ueno, K. Yagi

Vanderbilt University, Nashville, TN 37235

R. L. Clark, E. W. Cornell, A. V. Ramayya, C. F. Maguire*

Yale University, New Haven, CT 06511

M. Bennett, A. Chikanian, S. Coe, G. Diebold, B. Shiva Kumar*, J. Nagle,
A. Oulette, K. Pope, N. Smirnov (on leave from PNPI)

Individual Participation:

V. Emel'yanov (MEPI, Moscow), R. Schlueter (LBL), R. Vogt (GSI)

* Contact person for each institution.

Spokesperson: S. Nagamiya

Project Director: S. H. Aronson

Deputy Project Director: G. R. Young

Project Engineer: L. Paffrath

Assistant for Project Management: W. L. Kehoe

Liaison Physicist: B. Johnson

Chapter 1

Overview

1.1 Introduction

The PHENIX Conceptual Design Report Update (CDR Update) is intended for use together with the Conceptual Design Report (CDR). The CDR Update is a companion document to the CDR, and it describes the collaboration's progress since the CDR was submitted in January 1993. Therefore, this document concentrates on changes, refinements, and decisions that have been made over the past year. These documents together define the baseline PHENIX detector that the collaboration intends to build for operation at RHIC startup.

In this chapter the current status of the detector and its motivation are briefly described. In Chapters 2 and 3 the detector and the physics performance are more fully developed. In Chapters 4 through 13 the details of the present design status, the technology choices, and the construction costs and schedules are presented.

1.2 Physics

The physics goals of PHENIX collaboration have remained exactly as they were described in the CDR. Primary among these is the detection of a new phase of matter, the quark-gluon plasma (QGP), and the measurement of its properties. The PHENIX experiment will measure many of the best potential QGP signatures to see if any or all of these physics variables show anomalies simultaneously due to the formation of the QGP.

Table 1.1 summarizes the physics quantities the PHENIX detector will measure and their relationship to the collaboration's physics objectives. If the QGP phase transition occurs, one of our main goals is to characterize its nature. Among the most important elements of our physics program are measurements of:

1. Debye screening of QCD interactions.
2. Chiral symmetry restoration.
3. Thermal radiation of a hot gas.

Table 1.1: Physics Issues Related to Quark-Gluon Plasma

QGP Physics Issues	Probes
<u>Debye Screening of QCD Interactions</u> <ul style="list-style-type: none"> • $r(\Upsilon) = 0.13 \text{ fm} < r(J/\psi) = 0.29 \text{ fm} < r(\psi') = 0.56 \text{ fm}$ $J/\psi \rightarrow e^+e^-$ at $y \simeq 0$. $J/\psi \rightarrow \mu^+\mu^-$ at $y \simeq 2$. $\psi', \Upsilon \rightarrow \mu^+\mu^-$ at $y \simeq 2$. 	Electrons Muons
<u>Chiral Symmetry Restoration</u> <ul style="list-style-type: none"> • Mass, Width, Branching Ratio: $\phi \rightarrow e^+e^-, K^+K^-$ with $\Delta m \leq 5 \text{ MeV}$. • Baryon Susceptibility: Production of antinuclei. • Narrow σ-meson? 	Electrons Hadrons
<u>Thermal Radiation of Hot Gas</u> <ul style="list-style-type: none"> • Prompt γ, Prompt $\gamma^* \rightarrow e^+e^-$. 	Photons, Electrons
<u>Deconfinement: Nature of the Phase Transition</u> <ul style="list-style-type: none"> • First-order: Entropy Jump $\rightarrow$ Second rise in the $\langle p_T \rangle$ spectra of π, K, p. • Second-order: Fluctuation $\rightarrow N(\pi^0)/N(\pi^+ + \pi^-), d^2N/d\eta d\phi$. 	Hadrons Hadrons, Photons
<u>Strangeness and Charm Production</u> <ul style="list-style-type: none"> • Production of K^+, K^-, K_L^0. $\phi \rightarrow e^+e^-, K^+K^-$ at $y \simeq 0$, $\phi \rightarrow \mu^+\mu^-$ at $y \simeq 2$. D-meson: $e\mu$ coincidence. 	Hadrons Electrons Muons
<u>Jet Quenching</u> <ul style="list-style-type: none"> • High p_T jets via leading particle spectra. 	Hadrons
<u>Space-Time Evolution</u> <ul style="list-style-type: none"> • HBT correlations for $\pi\pi$ and KK. 	Hadrons

In order to accomplish these measurements, lepton, photon, and hadron signatures are studied simultaneously in the baseline PHENIX detector. The leptons and photons are thought to originate primarily in the plasma phase and, therefore, probe it directly. The hadrons arise later in the process and provide complementary information about the hadronization phase transition.

As Table 1.1 makes clear, in order to accomplish our goals we must measure electron pairs, muon pairs, photons, and charged hadrons. Therefore, particle identification is an important element of the PHENIX detector. What is less clear from Table 1.1 is that in order to use such signals to find and deduce the properties of the QGP, a *programmatic study* of these potential signatures is required, as a function of energy density in both $A + A$ and $p + A$ collisions. Furthermore, since some of the potential signatures involve rare processes and small effects, this program requires a detector capable of taking data at the highest

luminosities expected at RHIC.

1.3 Detector

The PHENIX detector is essentially unchanged from that presented in the CDR. Minor changes have occurred in the conceptual design, which were primarily related to the maturing of the design, results from more sophisticated simulations, technology choices being made, and more attention to issues of integration and installation. The overall changes to the PHENIX detector are described in Chapter 2, and specific modifications associated with each of the detector subsystems are described in Chapters 4-13.

Figure 1.1 shows a three-dimensional cutaway view of the PHENIX detector, showing the location of the various detector subsystems. The acceptance and function of the detector subsystems are summarized in Table 1.2

Table 1.2: Summary of the PHENIX Detector Subsystems.

Element	$\Delta\eta$	$\Delta\phi$	Purpose and Special Features
Magnet: central (CM) muon (MM)	± 0.5 1.1-2.5	360° 360°	Up to 1.0 T·m. 0.72 T·m for $\eta = 2$, 0.36 T·m for $\eta = 1.3$.
Silicon (MVD)	± 2.7	360°	$d^2N/d\eta d\phi$, precise vertex.
Beam-beam (BB)	$\pm(3.1-4)$	360°	Start timing and fast vertex.
Drift chambers (DC)	± 0.35	90° + 90°	Tracking, and good momentum and mass resolution, $\Delta m/m = 0.4\%$ at $m = 1\text{ GeV}$.
Pad chambers (PC)	± 0.35	90° + 90°	Pattern recognition and tracking for nonbend direction.
TEC	± 0.35	90° + 90°	Pattern recognition, tracking after RICH, and dE/dx .
RICH	± 0.35	90° + 90°	Electron identification.
TOF	± 0.35	30°	Good hadron identification, $\sigma < 100$ ps.
PbSc	± 0.35	90° + 45°	Good e/π separation at $p > 1$ GeV/c by EM shower and $p < 0.35$ GeV/c by time-of-flight. K/e separation at 1 GeV/c by time-of-flight for dE/dx .
PbGl	± 0.35	45°	Photon detection.
Muon chambers (μT)	1.15-2.35	360°	Tracking for muons.
Muon identifier (μID)	1.15-2.35	360°	Steel absorbers and chambers for μ /hadron separation.

Figure 1.1: A three-dimensional cutaway view of the PHENIX detector, showing the location of the various detector subsystems. The detector subsystems are labeled in the figure.

Ongoing R&D and simulation studies on all aspects of the detector appear to uphold our previous expectations of performance. These are summarized in Table 1.3.

Table 1.3: Performance of the PHENIX Detector.

Physics	Day-1 Version	Upgrade
Electrons	$\pi/e < 10^{-4}$ at $p \leq 4$ GeV/c • RICH for < 4 GeV/c • TEC (dE/dx) for < 2 GeV/c • EMCal for > 1 GeV/c	$\pi/e < 10^{-4}$ for all p_T : • TRD for > 2 GeV/c Low mass background rejection: • Dalitz Rejector
Photons	$p_T \geq 1$ GeV/c for 0.5 sr with PbGl. $p_T \geq 1-1.5$ GeV/c for 1.5 sr with PbSc.	$p_T \geq 0.5$ GeV/c for 0.17 sr with CsI or BaF ₂ .
Hadrons	≤ 2.5 GeV/c π -K for 0.36 sr • TOF with $\sigma < 100$ ps.	
Muons	$\pi/\mu < 10^{-4}$ at $p \geq 2.3$ GeV/c with 6 layers of μ ID.	Extended η coverage: • Additional arm(s)
Global	$d^2N/d\eta d\phi$ for $ \eta < 2.7$	

1.4 Strategy

Since the submission of the CDR, more refined and detailed cost estimates, construction schedules, and contingency estimates have been made. These are summarized in Chapter 15, and are presented in detail in the PHENIX Management Plan, the revised Work Breakdown Structure book and the WBS Notes book. As was the case in the CDR, the cost of the complete detector described above exceeds somewhat the allocated DOE budget of \$35.5M (FY 1993 dollars).

The plan to address this problem, involving deferral of certain detector components, remains essentially as it was presented in the CDR. We have not found it necessary to defer any additional subsystems or physics capabilities in the present Baseline. The deferred elements are:

1. Muon tracking (μ T) and muon identifier (μ ID) instrumentation.
2. High data acquisition bandwidth.
3. High level triggers.

The collaboration remains committed to its plan to obtain the necessary additional resources to install these deferred components as soon as possible.

The list of possible upgrades to PHENIX also remains the same, and R&D efforts on these options are continuing. These are:

1. A "hadron blind" Dalitz rejector for the electron arms.
2. A high resolution scintillating crystal calorimeter for low p_T photon measurements.
3. A transition radiation detector for high energy electron identification.

Additional coverage for muons and additional devices for event characterization have also been discussed.

Chapter 2

Baseline Detector

2.1 Introduction

The PHENIX detector configuration is essentially unchanged from that presented in the Conceptual Design Report (CDR). The solid angle coverage of the inner detectors, the central arms and the forward muon arm is the same, except for slight modifications that were required as the subsystem designs became more realistic. Channel counts remained nearly unchanged as the designs for the detector subsystems matured.

In the CDR a "Baseline" version of the detector was defined. In the baseline version, certain subsystem components were deferred for budgetary reasons. These included the muon arm instrumentation, the inner coils of the central magnet, and high level triggers for the data acquisition system. Other subsystem components were regarded as future upgrades. These included the transition radiation detection (TRD) capability of the time expansion chamber (TEC), the high resolution scintillating crystals (HRX) in the electromagnetic calorimeter (EMCal), and the inner hadron-blind detector (HBD). The present version of the detector retains the same menu of deferrals and upgrades.

At the time the CDR was written, a number of technology choices remained to be decided. These included the layout of the multiplicity/vertex detector (MVD), the cathode geometry of the pad chambers (PC) and the focal plane detector of the ring imaging cherenkov counter (RICH). These choices have now all been made.

Chapters 4 through 12 present the details of design changes and technology choices that have been made during the 10 months since the CDR was submitted. Chapter 12, which describes the offline computing effort, is new. The rest of this chapter is devoted to a brief summary of these updates to the baseline detector. Table 2.1 is a basic parameter list for the baseline detector.

2.2 Coverage and Geometry

As indicated above, the basic layout and coverage of the three spectrometer arms and the inner detectors remains unchanged. Small adjustments have been made in the course of

Table 2.1: Parameter List for the Baseline PHENIX Detector

Detector	Solid Angle		Radial Allocation (mm)	Spatial resolution (mm)	2 Particle Resolution (mm)	Timing Resolution	Number of Channels	Preamplifier Gain (mv/fc)	Shaping TC (ns)	Storage Type	ADC Word Width*	Digital Conversion
	$\Delta \eta$	$\Delta \phi$										
BB	$\pm(3.1-4.0)$	360°	-	25	-	70 ps	132	(pmt)	(pmt)	DMU	10/10	Flash ADC
MVD	± 2.7	360°	50-200	0.2 (z)	-	-	34816	40	60	AMU	8/8	ADC
DC	± 0.35	$90^\circ+90^\circ$	2050-2450	0.15 (r- ϕ)	1.5	-	10240	15	10	DMU	9/16	pTDC
PC1	± 0.35	$90^\circ+90^\circ$	2475-2525	2 (r- ϕ) 2 (z)	8 (r- ϕ) 8 (z)	-	15616	3	400	AMU	11/11	ADC
RICH	± 0.35	$90^\circ+90^\circ$	2575-4100	1° (θ) 1° (ϕ)	2° (θ) 2° (ϕ)	-	6400	(pmt)	(pmt)	DMU	8/8 8/8	TAC/ Flash ADC
PC2	± 0.35	$90^\circ+90^\circ$	4150-4210	3 (r- ϕ) 3 (z)	12 (r- ϕ) 12 (z)	-	15232	2	400	AMU	11/11	ADC
TEC	± 0.35	$90^\circ+90^\circ$	4217-4910	0.25 (r- ϕ)	4 (r- ϕ) 5 (z)	14 ns	28032	20	70	DMU	5/9	Flash ADC
PC3	± 0.35	$90^\circ+90^\circ$	4910-4980	4 (r- ϕ) 4 (z)	16 (r- ϕ) 16 (z)	-	15552	2	400	AMU	11/11	ADC
ToF	± 0.35	30°	5030-5180	15 (r- ϕ) 15 (z)	15 (z)	80 ps	2048	(pmt)	(pmt)	AMU	10/10 12/12	TAC/ADC
EMCal - PbSc	± 0.35	$90^\circ+45^\circ$	5030-5930	8 @ 1 GeV	70	180 ps @ 0.5 GeV	15552	(pmt)	(pmt)	AMU	11/15	TAC/ADC
EMCal - PbGl	± 0.35	45°	5230-6130	6 @ 1 GeV	50	310 ps @ 0.5 GeV	9984	(pmt)	(pmt)	AMU	11/15	TAC/ADC
Muon Tracking	1.15-2.35	360°	-	0.2	1.5	-	19680	15	30	DMU	7/16	pTDC
Muon ID	1.15-2.35	360°	-	1° (ϕ) 5.6° (θ)	2° (ϕ) 11.2° (θ)	<100 ns	18360	1	70	AMU	8/8	ADC

* For the ADC word width the first value is the number of ADC bits and the second is the effective resolution with ranging.

integrating the mechanical design. Table 2.1 summarizes the coverage (θ and ϕ) of the detectors in the central and forward arms, and the inner detectors subsystem.

2.3 Magnets

Numerous refinements in the magnet system have been made in the course of finalizing the design. However, the basic specifications of both magnets are unchanged since the CDR was submitted. The final design is reviewed in Chapter 4; salient changes are listed here.

1. The shape of the CM polefaces has been slightly modified to reduce the fringe field in the region of the RICH focal plane, and to make it possible to move the frames of DC and PC1 further from the active area of the central arms. This was found to be important in reducing the photon-conversion background in the RICH.
2. The bore in the MM piston was increased from 12 to 22 cm diameter. This allows sufficient clearance for the beams in asymmetric ($p + A$ and $A + p$) running conditions, with less than a 1% reduction in the field strength of the MM.
3. The five upper "lampshade" plates of the MM flux return have been segmented, and the ends of these lampshade plates nearest CM form a fixed arch. See Chapter 4 for figures and details. This change has several technical virtues, the main ones being that (a) the downstream, removable parts of the lampshade are now all independently removable, and (b) the upstream and downstream edges of the fixed arch provide simple, stable mounting surfaces for μT stations #1 and #2.
4. Following discussions with potential fabricators, the CM steel has been segmented differently than in the CDR version. The result is that all parts are now small enough to be cast or forged by all fabricators.
5. A rail system for the CM and the central carriages has been designed so that it will be relatively easy and quick to move CM 1.5-m upstream of the collision point, allowing access to the BB array near MM and to the μT station #1. The rail system, which is described in detail in Chapter 13, also allows motion of CM and the central carriages in and out of the collision hall, although this is not designed to be a quick operation.

2.4 Central Arms

The central arms are composed of detector elements that perform the functions of charged particle tracking and momentum measurement, electron and photon energy measurement, and particle identification. Although the arrangement and function of the central arm components is unchanged, there have been slight adjustments in the radial boundaries between detectors. These are summarized in Table 2.1. In addition, the subsystems outside the RICH (i.e., PC2, TEC, PC3, TOF, and EMCal) have now been subdivided into four sectors per arm instead of two. This is a better approximation to a cylindrical design and improves

the performance by limiting the range of incident angles of charged particles on these detectors. Figure 2.1 is a cross-section of the PHENIX detector that shows the central arm detector layout at the nominal collision point in the major facility hall, along with the radial dimensions of the detectors. An isometric view of the PHENIX detector is shown in Fig. 2.2.

2.4.1 Charged Particle Tracking and Momentum Measurement

The magnetic field and tracking chamber arrangement is the same as that in the CDR. The implementation of the TEC has been changed to reflect further study of its eventual upgrade to a TRD tracker. In the present design, all six layers of the TEC chamber will be built and installed in the baseline detector, whereas only four of six were previously planned. However, in the baseline, only four of the six layers will be instrumented with electronics. The radiator packs, the xenon recirculating system, and the additional electronics will be built and installed in the upgrade.

2.4.2 Electron and Photon Energy Measurements

A minor change in the dimensions of PbSc calorimeter modules and supermodules has been made in light of the resegmentation of the EMCal walls. The final dimensions of individual elements of the PbSc are $5.5 \times 5.5 \text{ cm}^2$. This has reduced the channel count, as reflected in Table 2.1.

2.4.3 Particle Identification

The philosophy of particle ID for electrons and hadrons in the central arms is unchanged from the CDR. The current status of particle ID components is listed below:

1. **Time-of-Flight:** A compact design of the precision TOF elements, with the PMTs folded back along the scintillator slats using 180° light guides, has been shown to work well. Beam tests give $\sigma_t \approx 88 \text{ ps}$. The mechanical design of the precision TOF wall has been finalized, using this type of element.
2. **Cherenkov Counter:** At the time the CDR was written the focal plane detector technology for the RICH had not been chosen. The choices were an array of PMTs, and a readout based on a parallel-plate avalanche counter with cathode pads and a solid CsI photocathode. In spite of very encouraging test results with the CsI-based technology, a decision was taken to use the PMT version in the Baseline. Although the CsI system might eventually be less expensive than a PMT-based readout, it was determined that it is now too early in its development to base such a large-scale detector upon it. An essentially similar readout technology is still being studied for a hadron-blind detector (HBD) that might serve as an inner tracker for Dalitz rejection. The HBD is a possible upgrade for PHENIX and is presented in Chapter 16.

Figure 2.1: Cross-section of the detector through the nominal collision point, showing the layout of detector subsystems in the central arms and the radial boundaries between subsystems.

Figure 2.2: Isometric view of the PHENIX detector.

3. **Time Expansion Chamber:** As mentioned above, the current plan calls for construction of all six layers of the TEC in the Baseline, as compared to four in the CDR version. This plan simplifies the conversion of the TEC to a TRD tracker later on.

Figure 2.3: Cross-section showing the layout of detectors in the muon arm, and its longitudinal (z) location in relation to the position of the central magnet.

2.5 Muon Arm

Since the submission of the CDR, there has been considerable refinement of the design concepts for the muon arm instrumentation. Table 2.1 summarizes the operating parameters and coverage of the detector elements in this subsystem. Figure 2.3 shows in cross-section the detector layout in the muon arm.

Progress has been made in the design of the muon arm since the submission of the CDR in the following areas:

1. The muon arm implementation plan, involving early fabrication of the MM and deferral of the μT and μID instrumentation, was reviewed by a special committee convened in June 1993. The physics goals of the muon arm, the present design of the instrumentation, and the schedule impact of various installation scenarios were all reviewed. In brief, the committee had positive findings in all these areas. As a result, design work on the MM (and the μID absorber) has continued, and is now essentially finished, on the basis of an early, fixed installation.

2. The μ ID absorber, which was shown as two layers of concrete in the CDR, has reverted to an earlier design of five layers of steel. This was shown at the June muon arm review to be much superior in performance, although more costly. The present design consists using five layers behind the MM backplate, allowing for a total of six active μ ID detector layers. The details of the design, including simulations of trigger performance and hadron rejection, are given in Chapter 10.
3. Optimization of the design of the μ T instrumentation led to the relocation of station #2 inside MM, moving it "upstream" by about 0.5 m. This improved the resolution slightly and reduced the size of the station #2 chambers.
4. Although the present design of the muon arm instrumentation (multiplane drift chambers in μ T and limited streamer tubes in μ ID) was found to be adequate at the June review, R&D work is in progress on other detector technologies that could improve performance. Cathode strip chambers are under consideration for muon tracking. These are in an advanced state of development for the GEM detector at SSC and could provide more robust chambers with less dead area and better spatial resolution. Resistive plate counters are under study as an alternate to the limited streamer tubes. Development of this technology was also supported by GEM R&D. These could provide a muon trigger within the required 112 ns; limited streamer tubes of smaller than standard size would be required to meet this requirement. Because the muon instrumentation is presently deferred, it is felt that there is some time available to evaluate potentially better technologies. Decisions will be made on both during FY 1994.

2.6 Inner Detectors

The functions and the coverage of the inner detectors, listed in Table 2.1, remain as they were in the CDR. An issue of the photon conversion background produced by the MVD has been resolved. The result is a single configuration that retains full azimuthal coverage for multiplicity measurements, provides adequate vertex reconstruction efficiency for all RHIC running conditions, and does not dominate other sources of background electrons. Compared to the CDR version, the present configuration also has somewhat reduced silicon area and channel count. Chapter 5 gives a full account of the new baseline configuration and simulation results.

The beam-beam (BB) counter design and function has remained essentially unchanged, although subsequent R&D has led to an improved design of the quartz radiators and a final specification for the PMT. Since the submission of the CDR, the location of the BB arrays has undergone reevaluation. The issue is access to the array on the same side as the muon arm. As described above, the current design of the rail system on which the CM and carriages sit allows relatively easy access to the BB array.

2.7 Trigger and Data Acquisition

The electronics for PHENIX deals with signals coming from a variety of different detectors, applies selective level-1 (LVL-1), level-2 (LVL-2) and level-2.5 (LVL-2.5) triggers to the data stream, and, because of the large event size and high luminosity, has to acquire data and build events at a very high rate. The front-end electronics is described in the chapters dealing with the individual subsystems. The design philosophy of the Baseline electronics is basically the same as in the CDR, although much progress has been made in the detailed implementation of this philosophy. The hardware and software for triggering, data acquisition and on-line computing and archiving are described in detail in Chapter 11. The Baseline configuration is described briefly below.

2.7.1 Front-end Electronics

Although PHENIX contains a number of different detector technologies, a strong effort has been made in the design of the front-end electronics (FEE) to minimize the number of different types of electronics chains. Reasonable success was achieved in this effort. However, it is not possible, for reasons of cost and power consumption, to digitize everything in real time directly on the detectors for all subsystems. The PHENIX detector will digitize in real time for all signals from the BB, RICH, DC, TEC, and μ T, but not for the MVD (power/heat) or for the PC, TOF, EMCal, and μ ID (cost). Consequently, digitized output is stored in digital memory units (DMU) for the first group and in one of at most three types of switched-capacitor analog memory units (AMU) for the second group. Digitization of the data in the AMU will be done following an "accept" from the LVL-1 trigger in a local ADC servicing a channel or group of channels. The AMU/DMU pipelines are made 64 cells deep in order to accommodate LVL-1 time latency and the buffering of up to seven events pending digitization following a LVL-1 "accept."

Another design goal for the FEE is to include the capability of calibrating every channel, under program control. This feature is included explicitly in the cost estimate.

The readout electronics must perform zero-suppression at the earliest possible stage in order to shrink the data volume and to reduce the computing power needed to do subtractions and corrections to the data. This is done immediately after digitization and the initial ordering and labelling of the data for readout. Zero-suppression reduces the data volume by nearly a factor of 20 for central Au + Au events and nearly a factor of 1000 for p + p events. Gain corrections and calibrations are applied to the remaining valid data after this step. Both zero-suppression and the correction/calibration step are performed in data collection modules (DCM) sitting in Futurebus+ crates located near the detector or on the support carriage. Data from the DCMs are made available to the Level-2 trigger, to the event builder, and to the monitoring ports. The DCMs contain reprogrammable elements so that raw data may be transmitted and the parameters for the zero-suppression and calibration algorithms may be adjusted.

2.7.2 Level-1 Trigger

The LVL-1 trigger will be deadtimeless to enable PHENIX to examine all bunch-crossings at the highest available luminosities for physics events. It is synchronized with the machine bunch-crossing clock and must be “pipelined” in order to be deadtimeless. The LVL-1 delivers a trigger decision each beam crossing for a beam crossing occurring a fixed time earlier. The “local” LVL-1 algorithms deliver detector information to the global level-1 (GL1) trigger within 32 crossings after the event. The GL1 trigger examines information from the entire detector and makes the initial “accept” or “reject” decision for the event in another 10 crossings. This is accomplished by means of routing layers (cascaded parallel lookup tables and scaledown units), all of which are programmable. As described in Chapter 11, the design for LVL-1 uses input from BB, MVD, TOF, RICH, EMCal and μ ID.

The LVL-1 trigger also provides several thousand pointers to detector elements satisfying certain conditions. These are passed on as data to the DCMs and the LVL-2 trigger in order to steer computational loops at that trigger level.

2.7.3 Level-2 and Level-2.5 Triggers

The LVL-2 and LVL-2.5 triggers are also pipelined processes. They are event-driven, unlike LVL-1 which is clock-driven. The distinction between levels 2 and 2.5 is that LVL-2 operates on partial (local LVL-2) or full (global LVL-2) trigger data but only local digitized detector data, whereas LVL-2.5 operates upon the full digitized data set from the detector, as constructed by the event builder. The LVL-2 trigger is event pipelined and the LVL-2.5 trigger is event-parallel because completion of the LVL-2 and LVL-2.5 algorithms will take longer than the mean time (selected as 40 μ s) between events accepted by LVL-1. Typical processing by LVL-2 includes the following steps:

1. Additional computation in the DCMs and following processors at the subsystem level (e.g., application of corrections to get better primitives for time of flight, calorimeter energy clusters, track segment direction cosines, or vertex location).
2. A number of algorithms, running simultaneously on parallel processors, to generate higher level primitives (e.g., identification of $e^\pm$, $\mu^\pm$, $K^\pm$, and γ).
3. A global LVL-2 merger of the relevant primitives to produce event signatures (e.g., “ e^+e^- pair” or “high p_T photon” or “hadron pair with $\Delta Q < 10\text{MeV}/c$ ”).

The LVL-2.5 processors operate on fully assembled events and can apply mass versus p_T cuts, refine tracking fits and tag rare events, such as those containing J/ψ or η mesons. Events accepted by LVL-2.5 are recorded on archival storage and sent to the on-line computing system (ONCS) for on-line monitoring, display and analysis. For the Baseline configuration described here, the LVL-2 trigger is deferred in the same way as the inner CM coils and instrumentation for the muon arm. We believe that deferral of this second level of trigger selectivity is consistent with the expected demands on the data acquisition system during early operations at RHIC. The overall trigger and data acquisition architecture is

designed to accommodate the LVL-2 trigger and increased bandwidth at all levels, as required by the design luminosity and future enhancements of luminosity. Implementation of the LVL-2 trigger is expected as soon as funding permits. This is motivated by the need to gain experience with the LVL-2 hardware and algorithms well before allowing them direct the event selection in PHENIX.

2.7.4 Event Builder, Data Archiving and On-Line Computing System

By the parallel nature of processing in the trigger, accepted events must be assembled from subevents in an "Event Builder". The event builder must be able to assemble the correct subevents in a robust and reliable fashion. It must accommodate expansion on both the input and output sides to cope with higher event and data rates. The design thus employs parallelism on the trigger side and a large number of output streams on the data logging side. The design is able to build events at data rates in excess of 1 GB/s. Initial operation of PHENIX at RHIC design luminosities only requires operation at rates up to 50 MB/s, but planned RHIC luminosity upgrades and the goal of PHENIX to measure infrequent events require ability to increase the bandwidth by at least an order of magnitude.

The on-line computing system is responsible for archiving data, monitoring and controlling the data flow through the system, downloading the more than 2000 embedded processors in the data acquisition system, configuring the entire on-line system, handling calibrating and monitoring tasks, providing needed databases, presenting a uniform interface to users wishing to operate the system or analyse results, providing needed software tools and code management facilities to enable software development, and monitoring and controlling all "slow" aspects of PHENIX such as high voltage systems, gas systems, AC and DC power systems, and temperature monitors. An approach based on a standardized realtime workstation and a distributed database and monitoring processors is being developed. This system will enable users to access all parts of PHENIX susceptible to computer monitoring and control via a common set of protocols.

The baseline data acquisition system will be able to archive data at 20 MB/s. At RHIC design luminosities this can keep up with the 2% most central Au + Au collisions and with p + p collisions, although not with p + A and some lighter A + A minimum-bias collisions. Thus in early running, or until RHIC reaches the design luminosities, the data acquisition system will be able to collect data with this taping capability and without the LVL-2 trigger. The bandwidth of the entire system, including the implementation of the LVL-2 trigger, will increase to keep up with expected increases in rates.

2.7.5 Off-line Computing

The high data archival rate of 20 MB/s results in very large requirements for data storage and cpu power. A combined total of 0.3-0.5 PB of raw, reconstructed and simulated data will have to be stored each year. The main storage medium is likely to be 19-mm video tapes recording at 16 MB/s and each storing 100 GB. A selected sample of events will be stored

in 10-20 TB storage robots. The event reconstruction and simulation will require 50-100 GFlops of continuous CPU power, if all events are fully reconstructed. The current "straw-man" model for RHIC computing assumes that the main CPU servers used for PHENIX event reconstruction and data analysis, as well as data recording and storage facilities, will be provided by the RHIC computing facility.

An important task for the off-line system will be to organize the large amount of very varied data within the collaboration. This will be accomplished through a set of databases tailored to the different types of data, but with a common user interface. The reconstruction and data analysis codes will be highly modular and be designed with special emphasis on utilizing the parallel architecture of computer farms or massive parallel processors.

2.8 Facility Engineering and Detector Integration

This item, which deals with detector infrastructure, has been expanded to include all integration and installation efforts, as well as structures and facilities. Significant progress, described in Chapter 13, has been made in this area, which benefited from an increase in the amount of available engineering resources at BNL devoted to this effort. Among the important areas of progress are:

1. As mentioned above, a system of permanent and temporary rails for the major facility hall (MFH) has been designed, allowing several kinds of detector motion for installation, servicing, and major upgrading.
2. Additional space near the MFH for run control and on-line computing has been specified and discussions are underway with Plant Engineering on the detailed design. The space requirements have been derived from an understanding of where various online functions will be located.
3. An installation plan has been worked out consistent with the overall RHIC schedule constraints and with the detector construction timelines. Resources for all of the installation effort have been integrated within this area of the PHENIX work break down structure (WBS) to make possible a reliable estimate of the total.

Chapter 3

Physics Capabilities

3.1 Overview

The major physics goals of the PHENIX collaboration were already described in the CDR. These goals are reviewed here briefly, together with new results from additional simulation studies since the submittal of the CDR. The physics issues related to the quark-gluon plasma, in particular those related to the PHENIX experiment, are summarized in Table 3.1.

Two important aspects exist for the quark-gluon plasma. One is “deconfinement”, which is closely related to the Debye screening of the QCD potential. The basic mechanism of J/ψ suppression is that, if the J/ψ radius is longer than the screening length, then a $c\bar{c}$ pair would not be able to form a bound state due to a Debye suppression of the long-range term of the $c\bar{c}$ potential and, thus, J/ψ production is suppressed [1]. The degree of suppression depends strongly on the relative size difference between the meson radius and the screening length. Because $r(\psi') = 0.51 \text{ fm} > r(J/\psi) = 0.25 \text{ fm} > r(\Upsilon) = 0.13 \text{ fm}$, we expect that the ψ' must melt first, then the J/ψ , and finally the Υ . Recently, a stronger suppression of the ψ' than the J/ψ was reported by the NA38 group at CERN [2]. This result compels us to study these vector mesons systematically at RHIC. We plan to measure the J/ψ in the midrapidity region by dielectrons, and the J/ψ , ψ' and Υ at forward angles by dimuons.

The other important element is “chiral symmetry restoration”. Because the mass of the ϕ -meson is close to twice the kaon mass, and because both ϕ and K could be distorted in the quark-gluon plasma, it was predicted [3] that a change would occur in: (1) the branching ratio between leptonic and hadronic channels, (2) the mass of the ϕ , and (3) the width of the ϕ . Here, high resolution ϕ -spectroscopy is required to study these points. We plan to measure ϕ -mesons using both the electron and hadron channels.

Thermal radiation from a hot gas has been a subject of controversy for many years [4]. There was a confusing period in which it was debated whether or not the radiation was enhanced or suppressed when the phase transition occurred. Recently, a consensus among theorists seems to be that the gluon content is high at an early stage of the quark-gluon plasma and an enhancement must be expected in the region of m_T (or, p_T for photons) greater than 2–3 GeV [5]. We plan to investigate this topic using the photon measurement capabilities in PHENIX.

Table 3.1: Physics Issues Related to Quark-Gluon Plasma

QGP Physics Issues	Probes
<u>Debye Screening of QCD Interactions</u> <ul style="list-style-type: none"> • $r(\Upsilon) = 0.13 \text{ fm} < r(J/\psi) = 0.29 \text{ fm} < r(\psi') = 0.56 \text{ fm}$ • $J/\psi \rightarrow e^+e^-$ at $y \simeq 0$. • $J/\psi \rightarrow \mu^+\mu^-$ at $y \simeq 2$. • $\psi', \Upsilon \rightarrow \mu^+\mu^-$ at $y \simeq 2$. 	Electrons Muons
<u>Chiral Symmetry Restoration</u> <ul style="list-style-type: none"> • Mass, Width, Branching Ratio: $\phi \rightarrow e^+e^-, K^+K^-$ with $\Delta m \leq 5 \text{ MeV}$. • Baryon Susceptibility: Production of antinuclei. • Narrow σ-meson? 	Electrons Hadrons
<u>Thermal Radiation of Hot Gas</u> <ul style="list-style-type: none"> • Prompt γ, Prompt $\gamma^* \rightarrow e^+e^-$. 	Photons, Electrons
<u>Deconfinement: Nature of the Phase Transition</u> <ul style="list-style-type: none"> • First-order: Entropy Jump $\rightarrow$ Second rise in the $\langle p_T \rangle$ spectra of π, K, p. • Second-order: Fluctuation $\rightarrow N(\pi^0)/N(\pi^+ + \pi^-), d^2N/d\eta d\phi$. 	Hadrons Hadrons, Photons
<u>Strangeness and Charm Production</u> <ul style="list-style-type: none"> • Production of K^+, K^-, K_L^0. • $\phi \rightarrow e^+e^-, K^+K^-$ at $y \simeq 0$, • $\phi \rightarrow \mu^+\mu^-$ at $y \simeq 2$. • D-meson: $e\mu$ coincidence. 	Hadrons Electrons Muons
<u>Jet Quenching</u> <ul style="list-style-type: none"> • High p_T jets via leading particle spectra. 	Hadrons
<u>Space-Time Evolution</u> <ul style="list-style-type: none"> • HBT correlations for $\pi\pi$ and KK. 	Hadrons

The nature of the phase transition is a very interesting point to study. If the phase transition is first-order, an entropy jump would be expected at the critical temperature (T_C), because internal degrees of freedom of the constituents increase by a factor of about 12 from the pionic gas to the quark-gluon gas. If one plots temperature (T) as a function of the energy density (ε) of the system, the value of T increases with ε in the phase of a pionic gas until it reaches T_C . The system stays at this temperature, even as ε increases, until a sufficient energy density is accumulated to allow the system to complete the phase transition into a quark-gluon gas. At this point, the value of T starts to rise again as a function of ε . Because T is closely related to the average transverse momentum ($\langle p_T \rangle$), the second rise in $\langle p_T \rangle$ as a function of ε is expected [6].

Such an entropy rise induces an increase of pressure and, thus, induces a hydrodynamical flow [7]. The flow effect is stronger for nucleons than for pions due to the mass difference. In

order to study the flow effect, combined with the effect of an entropy rise, measurements of "identified" charged hadrons are extremely important. We plan to measure identified hadron spectra in the midrapidity region.

If the phase transition is second-order, "fluctuation" measurements could prove important, as pointed out recently by Wilczek [8]. We plan to measure the ratio $\pi^0/(\pi^++\pi^-)$ from the detection of both photons and hadrons.

Table 3.1 lists other important physics issues, such as strangeness and charm enhancement [9], jet quenching [10], and space-time evolution [11]. Here, we emphasize that simultaneous measurements of dielectrons, dimuons, photons, and identified hadrons are important and necessary to pin down the formation of quark-gluon plasma at RHIC, because while all the theoretical predictions listed in Table 3.1 are very attractive, they all contain many weak points and ambiguities.

We plan to measure all the variables listed in Table 3.1 as a function of a reasonably well-defined experimental quantity that is proportional to the energy density. Because the energy density is given by [12]

$$\epsilon = \frac{1}{\pi R_{\perp}^2} \frac{dE_T}{dy}, \quad (3.1)$$

the geometry of the collision (namely, $R_{\perp}$) will be fixed first by multiplicity measurements. A wide coverage for the multiplicity measurement is very important, because a possible fluctuation effect must be smeared out for the determination of the collision geometry. By using an electromagnetic calorimeter or a subset of the multiplicity detector, the quantity of dE_T/dy will be measured to probe the local energy density at that rapidity, including the effect of the fluctuation.

For each physics observable, short comments on how to detect these observables are described in what follows.

Dielectron Measurements

Three detectors are used for electron identification: RICH, TEC (for dE/dx), and EMCal. Initial test results for a subset of the actual RICH detector, using pion and electron beams at KEK, clearly demonstrates that a Cherenkov photon ring is observed by a photomultiplier read-out method, and it shows that the pion rejection of 10^{-4} is possible for single pions and electrons. In the high multiplicity environment, the pion rejection power will be reduced to the 10^{-3} level.

For the TEC in the 1 GeV/c region, the value of dE/dx in a gas is approximately 50% higher for electrons than for pions. This feature is used for electron identification. Test results show a clear separation between electrons and pions, and we confirmed that this method allows an additional 10^{-2} level pion rejection up to 2 GeV/c. The EMCal further helps pion rejection by at least another one order of magnitude. Our design criterion of the pion rejection below 10^{-4} can thus be attained.

We also paid special attention to achieve a high mass resolution in the tracking system. Our Monte Carlo calculations have shown that the mass resolution of approximately 4 MeV for ϕ -mesons can be achieved for the e^+e^- channel and less than 1 MeV for the K^+K^- channel (see Section 3.2.3.).

Photon Measurements

In order to measure direct photons, the first step is to learn the background associated with π^0 and η . Invariant mass spectra for 2γ have the following general tendency: at low p_T the background is high, and as p_T increases the peak is more pronounced. Thus, the yield ambiguity increases as p_T decreases (see Fig. 3.7).

If π^0 and η peaks are found in the 2γ spectrum, the next step is to estimate the differential cross sections of π^0 and η and, then to evaluate the spectrum of single photons by using these π^0 and η . If the reconstruction procedure is perfect, the reconstructed spectrum of single photons must be identical to the actual single photon spectrum recorded on the detector. This comparison is described later in Section 3.3. Our conclusion is that the direct photon measurement is possible at $p_T > 1$ GeV/c, if an enhancement of the direct photon is well above a few percent of the single photons from π^0 .

Photon measurements also allow us to study π^0 spectrum, as well as two- π^0 interferometry [13].

Charged Hadron Measurements

A recent test of a prototype time-of-flight (TOF) counter ($1.2 \times 1.2 \times 100$ cm³) showed that the time resolution of $\sigma < 100$ ps is possible, even if the light guide was bent by 180°. With this resolution, a pion to kaon separation up to 2.4 GeV/c is possible.

The first interesting subject is to measure single particle spectra of identified hadrons, as discussed above. Charged hadron measurements further allow us to study $\phi \rightarrow K^+K^-$ and two-boson correlations. We can measure these correlations for both $\pi\pi$ and KK to determine the source radius up to 20 fm. For a source radius larger than 10 fm, the detector bin size of relative momentum has to be corrected rather carefully, in particular for KK correlations, because the resolution of relative momentum, Δq , is proportional to γm and, thus, $\Delta q(\pi\pi) < \Delta q(KK)$.

The decay kinematics of $\phi \rightarrow K^+K^-$, together with the measurement of two-boson correlations, determine the solid angle coverage (0.36 sr) of the TOF wall in the baseline detector configuration.

Dimuon Measurements

Muon identification will be performed by a standard technology of streamer tubes, sandwiched by several layers of iron (Fe) absorbers. In contrast to $p + p$ collision experiments, it is not possible to track muons before the first absorber, because the multiplicity is too high in heavy ion collisions. We use the central magnet yoke as a first, thick absorber and the tracking system starts after the yoke, which functions here as a pion attenuator.

The muon arm mass resolution is relatively modest at about 85 MeV for the J/ψ (see Fig. 10.30 in the CDR) as compared to a resolution of about 18 MeV for this resonance via the dielectron channel in the central arm. Nonetheless, the muon arm mass resolution should enable a clear separation of the J/ψ and the ψ' , thus being able to measure the possible differential suppression of these resonances as mentioned in beginning of this chapter.

A clear advantage at forward angles is a gain in the yield. Although the muon arm covers only 1 sr, which is half the solid angle of the electron arm, the value of $\Delta y \Delta \phi$ is 2.4π radians

for the muon arm, while it is 0.7π radians for the electron arm. Because the pair yield is proportional to $[\Delta y \Delta \phi]^2$, the yield of J/ψ is, for example, about a factor of 10 higher in the muon arm than in the electron arm. For one year of RHIC running, we expect to have 5700 ψ' and 1200 Υ for 10% central collisions at the luminosity of 2×10^{26} .

Because the central magnet yoke imposes a relatively high momentum threshold (about 2.5 GeV/c) this tends to limit the y and p_T acceptance for low mass pairs. The two front muon identifier steel layers have been specially graded to compensate for this effect based on recent trigger simulations.

Electron-Muon Coincidence

Once both the electron and muon arms are installed, the measurement of the $e\mu$ coincidence is possible. The main leptonic decay mode of the D-meson is an emission of an electron or muon, like the β decay of the neutron. Therefore, the unlike-sign $e\mu$ pairs subtracted by the like-sign $e\mu$ pairs are primarily from $D\bar{D}$, where the unlike-sign pairs are mainly from combinatorial background pairs [14].

In the CDR, it was clearly demonstrated that the charm contribution is dominant in the mass spectrum for $e\mu$ pairs at $m_{e\mu} > 3-4$ GeV (see Section 3.3.3 in the CDR). If the charm production is enhanced, one may probe $D\bar{D}$ at a lower mass region.

Multiplicity and Vertex

We plan to install both a silicon multiplicity-vertex detector (MVD) and the two sets of quartz Cherenkov timing counter arrays (beam-beam counters).

Tests for the beam-beam (BB) counters showed 35 ps time resolution in the magnetic field of 3 kG. These counters are used as a time reference for the TOF counters. Also, they are used for the on-line vertex determination at the level of $\Delta z = 2$ cm.

The MVD is used to determine the event multiplicity and the vertex position. With this device, the vertex can be determined to within 100–200 μm .

3.2 Simulations: Strategy and New Results

In the CDR an extensive series of integrated simulations was undertaken in order to establish the feasibility of the dielectron physics program in the PHENIX central arm (see Section 3.3.1 in the CDR). However, because these simulations were the first proof-of-principle task, they were conducted under idealized conditions. Most importantly, the background contributions originating from secondary reactions in the bulk of the central arm's structural material were largely ignored in the signal-to-noise figures. The objective at that time was to confirm, to the best of our knowledge of how each detector subsystem would perform, that the irreducible Dalitz decay and photon conversion backgrounds, as well as the limits of the predicted e/π discrimination, would not fatally obscure the vector resonance mass signatures. This primary objective was convincingly satisfied in our estimation by this first round of simulation studies.

Nonetheless, we did recognize (see Section 3.4 in the CDR) that secondary particle backgrounds from the large quantity of material in and near the central arm acceptance could

pose problems. This is a special concern for the performance of the RICH, which is the cornerstone of the e/π discrimination hardware. Indeed, disappointing results, reported recently at the QM '93 conference, regarding the performance and backgrounds afflicting the double-RICH CERES/NA45 experiment at CERN have made us acutely sensitive to this issue. Moreover, design changes in the PHENIX central arm magnet yokes have brought this steel even closer to the acceptance region than in the CDR. Not surprisingly, this change has increased the secondary particle backgrounds into all central arm detectors.

Consequently, in the last several months we have engaged in a systematic effort to ascertain the deterioration, if any, in the overall performance of the dielectron program of measurements. This new round of simulations has led us to suggest certain changes in the designs of central arm subsystems (see the updates on the MVD and the drift chambers). Given these modifications, our calculated results to date still leave us confident that the dielectron measurements will be successfully accomplished. This is an ongoing process and we accept that there is still a lot of work to be done in order to verify that none of the subsystems will be compromised by secondary particle flux.

We discuss first the results of new simulations on the predicted performance of the RICH, where special attention is paid to the effects of the secondary particle backgrounds present in the central arm. In particular, we present an in-depth study of the effect of photon conversions by the MVD detector system in comparison with other electron background sources. After that study, we describe new simulations on the tracking performance, as regards the vector meson mass resolutions, and indicate how the track reconstruction performance can be integrated into studies of other detector subsystems, such as the RICH and the MVD. A brief discussion of new acceptance calculations by the muon arm is mentioned at the end of this section.

3.2.1 Secondary Electrons on RICH

One of the most severe backgrounds expected to be present in the Ring Imaging Cherenkov (RICH) detector subsystem is that of secondary electrons and positrons. These are produced by photon conversions at the magnet pole pieces, by the inner tracking chambers (including the chambers themselves and their support structures), and in the RICH gas radiator. Because they are "real" electrons and positrons, they produce Cherenkov light and will increase the noise level in RICH system substantially.

We have studied the secondary electron background using PISA, the PHENIX integrated simulation program. Included in the simulation are the magnet pole pieces, the beam pipe, the inner tracking chambers, with a somewhat simplified (if not pessimistic) support frame, and the RICH, with a C_2H_6 radiator gas. The "nose cone" and the MVD detector were not in this particular simulation. Related results on a modified version of the MVD and the nosecone are presented in other sections of this chapter.

The support frames of the inner tracking chambers (DC1, DC2, and PC1) was represented by a 2-cm thick aluminum plate, which was thicker in terms of radiation length than the intended support frames to be made of a light composite material. The simulated frame was about 10 cm away from the active region of the detector to avoid having converted electrons directly entering the acceptance of the RICH. The PC1 chamber was taken to be about

Figure 3.1: Hit pattern in the RICH PMT array for a typical Au + Au central collision event (HIJET). The PMTs that detect at least one Cherenkov photon are plotted as squares. The four panels of this figure corresponds to the four sectors of the PMT array. Each sector has $80(\phi) \times 20(z)$ segmentation. The horizontal axis is the ϕ position of the PMTs and the vertical axis is the z position.

0.6% radiation length. We used single density HIJET Au + Au central collision events, and simulated all hadronic and electromagnetic shower development in $10^\circ < \theta < 170^\circ$ in polar angle and 2π in azimuth. The magnetic field map for the single outer coil with full field strength in the new yoke design was used in the calculation. The position of the reaction vertex was fixed at $z = 0$. Additional studies at $z \neq 0$ are being analyzed but, based on preliminary indications, these results are not expected to change our conclusions significantly.

Figure 3.1 shows a hit pattern in the RICH for a typical event. The four panels of the figure corresponds to the four sectors of RICH (upstream and downstream sectors of east-arm and west-arm). Here we assume that the photon detector will be an array of 6400 PMTs, each PMT having a Winston cone to collect Cherenkov light. Because the Winston cones limit the angular acceptance of the PMTs, only those Cherenkov photons that are accepted

Table 3.2: Summary of RICH Ring Sources

Electron Source	Number
Dalitz	1.0
Beam Pipe Conversion	0.4
Magnet Pole, Decay in Flight (Kaon)	0.5
DC Active Region	0.7
PC1 Active Region	3.7
DC + PC1 Support Frame	1.9
RICH Gas Volume	1.9

by the cones are plotted as in the figure. There are about ten Cherenkov rings and 100 PMT hits per event. Most of these rings are indeed produced by the secondary background.

Figure 3.2 shows the production vertex of electrons and positrons that produce at least five PMT hits in the RICH system. A good electron track typically produces ten PMT hits, and this requirement of five PMT hits yields about 90% acceptance for real electron signal. It also required that the timing of PMT hit be within 2 ns of proper timing for good electron signal. Because the PMTs has good timing resolution (≈ 250 ps in *rms*), this timing cut is very conservative. The timing cut reduces the amount of background by 20%. From this figure, we can identify the major sources of background:

1. A concentration at the collision vertex ($r = 0, z = 0$).
2. A small amount of conversion in magnet pole surface reaching the RICH.
3. Conversion in the DC and PC1 region; in particular, in the support structure.
4. Conversion in RICH gas volume.

There are about 10 Cherenkov clusters per event on average, of which only one cluster is caused by a Dalitz electron from the collision vertex. All others are caused by the conversions. The reason the number of Dalitz electrons is so small is that many of them are cut off by magnetic field. The minimum p_T at $z = 0$ is about 92 MeV/c, and particles with less than that momentum at the collision vertex ($r = 0$) never reach the RICH detector. This curl-up momentum value becomes progressively smaller for $r > 0$, where structural material is located. Moreover, even if a Dalitz electron from the initial vertex is above the 92 MeV/c curl-up momentum value, it may still not be registered in the RICH because its trajectory will have been bent too much relative to the Winston cone axis. It is only when the electron's momentum exceeds 150–200 MeV/c that it will fire the RICH. With these facts in mind, we quantify the sources of those conversions reaching the RICH in Table 3.2 .

We also studied the secondary electron background for $z = 30$ cm and for $z = 0$ cm at half field strength, and have obtained similar results. Within statistical error, the amount of

Figure 3.2: A scatter plot of the production vertex of electrons and positrons that produce more than five hits in RICH. Thirty events of single HIJET Au + Au central collisions are accumulated. The region occupied by the central magnet, the inner tracking detectors (DC + PC1), and the RICH gas volume are drawn in dashed lines to elucidate the figure.

background does not depend on the z -vertex position. In the half field case, the number of Cherenkov clusters increases. This is because more low momentum electron and positrons that had been cutoff by the magnet in full field strength now reach the RICH detector. The background from the other sources did not show any significant increase. We also studied a case when the DC frame was moved another 5 cm away from the acceptance. This change reduced the background from the frame by half to about 1.0 per event.

The secondary electron background will have a substantial effect on the level-1 and level-2 trigger. In particular, it appears from these simulation results that it will be very unlikely to have any level-1 trigger rejection capability for dielectron events in central Au + Au collisions. Moreover, the level-1 DAQ hardware will have to contend with processing an average of ten electron tracks per central collision. The strategy appropriate for the level-2 trigger is now being studied. However, the effect on the e/π rejection and the final dielectron

measurement will be much smaller. This is because many of the Cherenkov rings that are produced by the secondary background do not have good reconstructed tracks pointing back to the collision vertex. Hence they will not contribute to the background in the final offline analysis. Even if a secondary electron has a well reconstructed track apparently coming from the vertex, it will generally have a large discrepancy between its real energy (typically less than a few hundred MeV) and its apparent momentum (a few GeV/c or larger). This is because these secondaries are produced just before the first tracking chamber in a region of weak magnetic field, and therefore they are not bent much by the field. Such spurious electrons will be readily removed by the means of the electromagnetic calorimeter energy and momentum matching condition.

All of these expectations are borne out by the most advanced PISA simulations done to date. Namely, these new off-line type calculations do the complete pattern recognition including the full set of primary and the background particles (40% of the primary particle rate). The pattern recognition "starts from scratch" (see the discussion of Fig. 3.5 in Section 3.2.2) in that it is based on the simulated TDC responses of the drift chambers, including multiple hits, and is seeded with the expected primary vertex reconstruction from the MVD and the BB. Then the momentum reconstruction is done in PISA self-consistently based on this track recovery information. In turn, the track recovery information is furnished to the RICH ring-finder algorithm, as described in the previous paragraphs. The net result from this sophisticated simulation is that the simulated e/π rejection ratio and the dielectron invariant mass spectrum signal-to-noise ratios should be largely unaltered from what was shown and discussed in Figs. 3.3-3.5 of the CDR. At most (see the next section) there could be a 50% increase in dielectron combinatoric background depending upon the particle momentum threshold taken for invariant mass pairs.

We cannot emphasize strongly enough the difference between the planned use of the RICH in PHENIX and that attempted by CERES/NA45. In the latter experiment, the dual RICH system was designed to function *both* as a tracker and as a particle identifier. Preliminary results indicate that while that system functions well in peripheral collisions with low particle multiplicity, it is plagued by ghost tracks and suspect rings in high background central collisions. By contrast the PHENIX RICH, albeit θ -sensitive, will be operated primarily in a particle identifier mode. The track pointing information, as determined principally by PC1 and PC2, will lock in the ring-finder search. Without that independent tracking information, our own simulation results shown here indicate that the RICH information would be hopelessly confused in central collisions. Indeed, it is our understanding from the QM '93 conference that just such an independent tracking capability will be added to CERES/NA45 in future.

3.2.2 MVD Induced Conversion Electrons

The previous section was devoted to a study of the secondary electron backgrounds generated by the magnet yoke steel and the inner tracking chambers, including their support structures. That section focussed directly on the number of rings expected to be found by the RICH from these sources that would pass the level-1 trigger cuts. In the present section, we report on the results of a separate study that was performed to determine the effects of the

multiplicity-vertex detector (MVD). Unlike those electrons coming from the magnet steel and the tracking chamber frame support, the secondary electron backgrounds generated by the MVD cannot be rejected by tracking cuts. At best the track pointing error from the tracking system, on track-by-track basis, is simulated in PISA to be of the order of a few centimeters (see Fig. 3.5) and that is not good enough to reject conversions occurring in the MVD silicon layers.

Because the MVD can be optimized, a special study has been made using PISA to ascertain the effect of the $e^\pm$ ("electron") background in the RICH from the MVD relative to the so-called "irreducible" background from all other sources. This study aided in the optimization of the design of a new MVD configuration by highlighting the trade-offs and compromises to be made, in addition to emphasizing the need for and value of an integrated simulation program. The new design of the MVD, which resulted from this simulation study, has removed two-thirds of the material in the outer layer. More complete details of the updated design are contained in Chapter 5. Here we present the integrated simulation results using the redesigned MVD.

For this special study the input to PISA was the hadrons produced in Au + Au Hijing events with impact parameters less than 0.12 fm and was thus similar to the RICH ring-finding study conducted above for non-MVD sources. Dalitz decay of the π^0 to $e^+e^-\gamma$ with a 1.2% branching ratio is one source of electron background. Conversion of the photons from the dominant $\pi^0 \rightarrow \gamma\gamma$ decay branch into $e^\pm$ pairs in the material around the interaction region is another important source of background. Other hadrons contribute to the electron background mainly via a two step process. A hadronic interaction produces π^0 's, which then make electrons via the same mechanisms as primary π^0 's. This two-step channel does not give rise to a significant background component in the relatively thin beam pipe or MVD, but it is important in the magnet iron and in the muon nosecone. All of the primary and secondary (tertiary and so on) particles are tracked with PISA through the PHENIX detector. The MVD, the combination of drift chambers 1 and 2 (DC1/DC2) and pad chamber 1 (PC1), and the RICH detectors were included in the simulation. The calculation also included the beam pipe (1-mm Be), the central magnet steel (and the field produced by it), the muon absorber nosecone, and the muon magnet. Information about all electrons entering the RICH detectors was recorded, including the location of the production vertex. This study did not focus on the actual level-1 rings produced by the RICH, but other studies have shown that about one-third to one-half, depending upon the source position, of the electrons entering the RICH will produce observable rings according to the criteria mentioned in the previous section.

The average number of background electrons entering the RICH ($\langle N(e^\pm) \rangle$) from various sources is summarized in Table 3.3. The row labeled "magnet" includes the magnet poleface and iron plus the muon absorber nosecone. The row labeled "other" includes conversions and other interactions in air along with decay products from particles other than π^0 's. The column labeled "no cuts" includes all electrons entering the RICH detectors, without any further cuts. The background sources near the beamline have a p_T cut of approximately 100 MeV/c introduced by the magnetic field in the central region; particles at lower p_T values are unable to leave the central region and hit the RICH. The next column includes

Table 3.3: Average Number of Background $e^\pm$ Entering the RICH

Source	$\langle N(e^\pm) \rangle$ (no cuts)	$\langle N(e^\pm) \rangle$ ($p_T > 200 \text{ MeV}/c$)	$\langle N(e^\pm) \rangle$ ($\delta z_{vtx} < 5 \text{ cm}$)
Dalitz Decay	2.46 ± 0.13	0.86 ± 0.08	2.46 ± 0.13
Beam Pipe	0.59 ± 0.06	0.27 ± 0.04	0.52 ± 0.06
MVD Barrel Si	0.82 ± 0.07	0.21 ± 0.04	0.80 ± 0.07
MVD Electronics	0.50 ± 0.07	0.00 ± 0.00	0.49 ± 0.07
DC1/DC2/PC1	4.26 ± 0.17	0.39 ± 0.05	2.42 ± 0.25
Magnet	9.55 ± 0.25	0.88 ± 0.08	0.93 ± 0.13
Other	0.87 ± 0.08	0.15 ± 0.03	0.19 ± 0.04
Total	19.05 ± 0.35	2.76 ± 0.13	7.81 ± 0.33
MVD Fraction	$7 \pm 1\%$	$8 \pm 2\%$	$17 \pm 1\%$

only those electrons with $p_T > 200 \text{ MeV}/c$. The angular acceptance (for Cherenkov light) of the Winston cones in front of the RICH photomultiplier tubes introduces a p_T threshold of approximately $150 \text{ MeV}/c$. The p_T cut used in Table 3.3 was chosen to be slightly above this threshold.

Figure 3.3 presents some of the information in Table 3.3. Almost all ($\approx 90\%$) background from the beam pipe and the MVD is produced by conversions of photons from π^0 decay. The top plot shows $\langle N(e^\pm) \rangle$ from various sources in the inner region: Dalitz decay, the beam pipe, the MVD barrel silicon detector wafers, and the MVD electronics. The numbers with and without the $200 \text{ MeV}/c$ p_T cut are shown. The second plot shows the number of radiation lengths of material (X_0) in each of the elements. For the beam pipe and the MVD silicon, this is straightforward. For the MVD electronics, an average number of radiation lengths, calculated by assuming the material in the electronics was distributed in a uniform cylindrical shell at the appropriate radius is shown. The real design will place all of the electronics at the bottom of the detector, out of the azimuthal coverage of the detector arms. The number of “radiation lengths” associated with Dalitz decay is defined as approximately $1.2\%/2 = 0.6\%$, where 1.2% is the branching ratio for π^0 Dalitz decay. This is the number of radiation lengths of material which would be required to produce the same number of $e^\pm$ conversions from π^0 decay photons. The bottom part of Fig. 3.3 shows the ratio $\langle N(e^\pm) \rangle / X_0$. With or without a p_T cut, the beam pipe gives a few more electrons per radiation length than the the MVD silicon. Without a p_T cut, the contribution of the MVD electronics is similar to the MVD silicon, but slightly lower. With a $200\text{-MeV}/c$ p_T cut, the contribution of the MVD electronics is completely eliminated. The slightly reduced number of electrons per radiation length of material from the MVD silicon, which has a hole in the upper two-thirds of the outer silicon layer, and the more substantially reduced ratio from the MVD electronics demonstrates the clear advantage in putting as much of the MVD material out of the direct path to the two electron arms as possible. The contribution of Dalitz decay is larger than predicted based on the number of “radiation lengths”. The

Figure 3.3: The average number of $e^\pm$ per event from the beam pipe, MVD silicon, and MVD electronics (top), the number of “radiation lengths” associated with each of these sources (middle), and the ratio of the two (bottom) are shown. “Radiation lengths” for Dalitz decay is defined in the text.

number of these “radiation lengths” was defined such that the same number of electrons was produced. However, the average electron from Dalitz decay has half of the energy of the original π^0 compared to a quarter for the average conversion electron. This energy difference is reflected in their transverse momenta. The approximately 100-MeV/c p_T cutoff imposed by the central magnetic field eliminates a smaller fraction of the electrons from Dalitz decay than from conversions, enhancing their numbers relative to the estimated number of “radiation lengths”.

As noted in the previous section, the first problem caused by background electrons is at the trigger level. Large numbers of electrons entering the RICH represent a significant problem for the trigger. The second column of Table 3.3 shows the number of electrons

entering the RICH from the different sources. The largest contributor is electrons produced in the magnet, followed by the the contribution from the combination of DC1/DC2/PC1. The triggering challenge is dominated by the need to diminish where possible electrons from these major background sources, but the trigger is not predicted to be limited by the MVD source.

After the data have been analyzed (perhaps including the online analysis in the later stages of the trigger), some sources of background should be easily eliminated. In particular, most electrons originating in the magnets and muon absorber should be easily separated from those originating near the vertex. This is illustrated by comparing columns 2 and 4 of Table 3.3. Column 2 shows the number of electrons entering the RICH without cuts, while column 4 shows the number which also trace back to a z position within 5 cm of the vertex. This cut eliminates most of the background from the magnet. Of the total "magnet" contribution about 44% of the electrons in the "no cuts" column are associated with primary π^0 's, falling to 29% when the 200-MeV/c p_T cut is applied. However, of the small fraction of electrons from the magnet which satisfy the vertex cut, about 90% are associated with primary π^0 's. The vertex cut eliminates slightly less than half of the background from the DC1/DC2/PC1 combination. About 90% of the DC1/DC2/PC1 background which survives the vertex cut comes from conversion of π^0 decay photons. Without cuts (column 2), only 62% of the electrons from the DC1/DC2/PC1 combination are from conversion of photons from primary π^0 's. The two step process by which particles other than π^0 's produce background can give rise to numerous electrons, but most of these electrons can be eliminated via a vertex cut.

On the other hand, the vertex cut has no effect on the the background from conversion of π^0 decay photons in the beam pipe and MVD; the small number of particles eliminated by these cuts were associated with interactions of primary particles other than π^0 's. No reasonable vertex cut would be able to eliminate the conversion electrons from the inner region of the spectrometer. For practical purposes, the photons from the π^0 decay comes from the vertex ($c\tau = 25$ nm) and the pair production process produces electrons which travel, to a good approximation, in the direction of the incident photon and therefore also appear to have come from the target (to within ≈ 100 – 300 μm for conversion in the MVD and beam pipe).

Because the background from material near the vertex is so difficult to eliminate, its impact on the pair mass spectrum must be understood. The MVD provides measurements that are important in the overall goals of PHENIX (e.g., $dN/d\eta$, $dN/d\eta d\phi$, multiplicity for the trigger, and the vertex position), but these positive contributions must be compared to the increase in electron background it causes. When considering the impact of the MVD on the background, the appropriate scale for comparison is the irreducible background that comes from Dalitz decay (unavoidable) and conversions in the beam pipe, which is as thin and light as possible (1-mm Be). There are two p_T cutoffs which apply to the inner detectors. First, is the approximately 100 MeV/c cutoff due to the magnetic field and second is the nearly 150 MeV/c p_T cutoff introduced by the RICH detector. The comparison of the MVD background to the irreducible background Table 3.3 used a third, slightly higher, p_T cutoff of 200 MeV/c. With this p_T cut, the ratio (MVD electrons)/(Dalitz+beam pipe) is 0.20 ± 0.04 .

Figure 3.4: Simulations of the invariant mass resolutions of the vector mesons to be detected in the PHENIX central arm.

When constructing a pair mass spectrum for electrons, the contribution of these background particles depends on the square of the total number produced, so the MVD in its current design is expected to increase the background in the pair mass spectrum by approximately $1.20^2 = 1.42 \pm 0.09$. About 10% of the material in the MVD barrel has been neglected in this calculation. The increase in background in the pair mass spectrum, compared to that from Dalitz decays and the beam pipe would be about $1.27^2 = 1.61 \pm 0.10$, with a 150-MeV/c p_T cut.

3.2.3 Mass Resolution in the Central Arm

We report that the simulated resolutions of the vector mesons detected in the central arm continue to be within acceptable limits. These new calculations were performed for a magnetic field that is slightly stronger (0.78 T-m vs. 0.70 T-m) than in the previous design. The detector tracking performance (position and angular resolutions) has been modeled here with better accuracy than before. Figure 3.4 is a plot of the mass resolution of the PHENIX central tracker, shown as a function of p_T , for the detection of these vector mesons. The drift chambers and the pad chambers were the only detectors used in these calculations. Also included in this simulation were the effects of the beam pipe, the intervening helium filler gas, and a realistically thick helium bag envelope. While not included, the MVD has been shown in other simulations to have no significant effect on the predicted vector meson

resolutions.

The resolutions were calculated in a single particle mode where only one particle of the species of interest was detected within the spectrometer acceptance. The results shown here indicate that in the K^+K^- decay channel the ϕ can be measured with a resolution of about 0.07% for a pair momentum below 1.0 GeV. For decays via electrons, the corresponding resolution is about 0.5%. The J/Ψ can be measured with a resolution of about 0.7%. These results are very similar to those presented in the CDR. Additional information about these calculations is given in Chapter 6.

The PISA/PISORP simulations now include a detailed description of the expected response of the drift chambers, and that response in turn is coupled to a modelling of the RICH ring finder procedure. An example of this calculation for the analysis of simulated ϕ -meson data is shown in Fig. 3.5.

The top panel presents the electron spectrum from a ϕ -resonance event generator [15] as determined by the tracking system, and then as accepted by the RICH electron identification algorithm. This algorithm takes as its starting input the tracking information provided by the drift chambers in order to confirm the presence of an electron ring. According to these calculations, the RICH should be 91% efficient with respect to the number of tracks recovered by the drift chamber analysis.

The middle panel shows the invariant mass spectrum as deduced from the opposite sign pairs tracked by the drift chambers and confirmed as electrons by the RICH system. As before, these are single meson event type simulations but the pattern recognition conditions still have to be fulfilled. In this middle panel are listed all the details of the drift chamber modeling, essentially the intrinsic noise effects that we can assume will detract from perfect performance. As usual, the tracking and momentum analysis will be done in a "multipass" fashion for which in the first pass one is completely ignorant as to the identity and momentum of the particle being tracked. In the second, and possibly higher, pass one can mitigate some of the imperfections of the response function that are known to be momentum or particle dependent, and in that manner achieve an improved mass resolution. The mass resolutions quoted in this figure are essentially the resolutions shown in Fig. 3.4, integrated over this event generator's transverse momentum distribution.

Finally, in the lower panel is shown the track pointing accuracy expected from the drift chamber system. In this simulation, all the electrons started from $z = 0$, so the width of the deduced z origin is a measure of the tracking system's vertex resolution on a track-by-track basis. It should be pointed out that this spectrum was obtained without the use of the PC1 data, and that data should improve the vertex resolution somewhat. Nonetheless, one would not expect that the final resolution would be much better than 3 cm, and that resolution, as has been mentioned in the previous section, is insufficient to tag effectively photon conversion electrons originating from an MVD source. Hence, our effort to minimize the MVD material available for photon conversion into the central arm acceptance.

3.2.4 Dimuon Measurements

In the CDR, the analysis of the acceptance and coverage of the muon arm was carried out assuming "frameless" chambers. Since that time, an analysis of the acceptance was carried

Figure 3.5: Simulations of the data quality expected to be available from an analysis of the ϕ -mesons thrown by a standard event generator available in PISA. The top panel shows the electron momentum spectrum before and after gating by the RICH ring-finder algorithm. The middle panel shows the predicted invariant mass resolution with a variety of DC noise contributions included. The bottom panel shows the spectrum of vertex z position determinations for the electron tracks using only the data from the UV wires of the DCs.

out that used realistic chamber frame widths. The analysis required that the muons from the vector mesons penetrate to the third layer of the muon identification system and not hit any frames in the tracking system. The fraction of events that survive those cuts are 16 % for the ϕ , 47 % for the J/ψ , and 66 % for the Υ . Drell-Yan and $D\bar{D} \rightarrow \mu^+ + \mu^- + X$ rates will be reduced by about the same fraction as for a resonance of the same mass. There is a rapidity dependence to those acceptance losses, with larger losses occurring at higher rapidity. More detail can be found in Chapter 10. The chamber design is still being optimized, and it is felt that a significant fraction of the above losses are recoverable.

3.3 Revision for Photon Simulations

Invariant mass spectra for two photons under the condition of a realistic case of $dN_{\pi^0}/dy = 375$ is shown in Fig. 3.6. A similar figure is shown as Fig. 3.10 on page 3-26 of the CDR. A mistake was made in creating the previous figure, in which it was assumed that $dN_{\pi^0}/dy = 95$. Correspondingly, the yield uncertainty when one estimates cross sections of π^0 and η from the two-photon spectrum is higher than that presented in the CDR. The correct figure on the yield uncertainty is shown in Fig. 3.7. This figure is the replacement of Fig. 3.11 on page 3-27 in the CDR. Yield extraction uncertainty is defined as statistical uncertainty only; systematic errors are not taken into account.

If π^0 and η peaks are identified in the two-photon spectrum, the next step is to evaluate the spectrum of single photons by using these π^0 and η . This spectrum is shown in Fig. 3.8b. On the other hand, the actual photon spectrum recorded on the detector is given in Fig. 3.8a. If the reconstruction procedure is perfect, Fig. 3.8a must be identical to Fig 3.8b. In fact, these two seem identical.

However, if one takes the ratio between the two spectra, Figs. 3.8a and 3.8b, the ratio deviates from 1 at $p_T < 1$ GeV/c, as seen in Fig. 3.8c. This deviation would originate for two reasons: (1) the signal-to-noise ratio of π^0 becomes worse as p_T decreases and, therefore, the yield ambiguity of single photons from π^0 increases, and (2) the opening angle of 2γ from π^0 increases as p_T decreases, so that the contribution from π^0 's produced in the rapidity region outside the detector coverage becomes stronger as p_T decreases.

As described in Section 3.1, our conclusion is that the direct photon measurement is possible at $p_T > 1$ GeV/c, if an enhancement of the direct photon is well above a few percent of the single photons from π^0 .

3.4 QED Background Studies in the BB

The beam-beam counter (BB) determines the time zero (T_0) when an interaction occurs in a beam bunch crossing (2-ns traversal time) and it provides the first estimate of the vertex point (z_0) along the beam axis. The BB consists of two sets of identical counter arrays installed on both sides of the interaction region. One set is in the muon plug on the muon detector side, and the other is in free space. By measuring the pair of arrival times (T_1 and T_2) of the fastest particle that intercepts the counter arrays on each side, the time zero T_0

Figure 3.6: Two photon invariant mass spectra at various photon pair transverse momenta using $dN_{\pi^0}/dy = 375$. The results are shown for the full PHENIX electromagnetic calorimeter detector acceptance for photons with energy greater than 250 MeV.

Figure 3.7: PHENIX electromagnetic calorimeter detector performance for $dN_{\pi^0}/dy = 95$ (squares) and $dN_{\pi^0}/dy = 375$ (circles), the multiplicity anticipated at RHIC. Top: π^0 (open points) and η (solid points) measurement showing the transverse momentum dependence of the peak-to-total ratio (Q). An asymmetry cut of $\alpha \leq 0.4$ is applied. Bottom: The statistical uncertainties in the meson yield extraction. Solid points: π^0 , open points: η . Systematic errors are not taken into account.

Figure 3.8: Measured single photons as compared with the reconstructed single photons from π^0 and η for two days runs at RHIC for central Au + Au collisions. An electromagnetic calorimeter resolution of 7% at 1 GeV is assumed.

and the interaction point z_0 can be obtained as the average and the difference of T_1 and T_2 , respectively.

If there were no backgrounds nor any slow particles, the resolutions of the extracted T_0 and the z_0 would be affected only by the intrinsic time resolution of the counter elements. However, all interactions do produce slow particles, even at very forward angles, and there does exist a significant amount of background particles scattered from the central spectrometer structures and from the muon detector. Therefore an examination of the BB performance with a Monte Carlo simulation is mandated. The specific results of this simulation of the intrinsic BB performance is detailed in Chapter 5. We choose to show here the first results of a study of the *physics* background into the BB, namely that coming from quantum electrodynamic (QED) pair production.

Figure 3.9: The p_T cutoff momentum of electrons at the BB.

In a relativistic heavy ion beam crossing, there is production of energetic electron and positron pairs due to QED processes. In this environment, the BB subsystem risks being exposed to an enormous amount of QED electrons and positrons. These events occur independently of the nucleus-nucleus strong interaction. The QED backgrounds could affect the performance of the T_0 and the z_0 determinations in the nucleus-nucleus interaction if both of the QED and strong nucleus-nucleus processes occur in the same beam bunch crossing. The majority of such QED electrons and positrons are expected to be of low total momentum and would be curled up in the axial magnetic field, absorbed in the beam pipe or an absorber in front of the BB, and will not be detected. However, some energetic electrons or positrons could reach the detector and produce misleading TDC data.

A simulation was performed to evaluate the yield of QED backgrounds that might contaminate the BB measurement. This study used the PISA/PISORP program chain and is described more fully in Chapter 5. Briefly, 1000 electrons with various P_T momenta were thrown at a fixed angle $\theta = 4^\circ$ from the $z = 0$ position, and the number of events with any hit in the BB array was counted. Most of the hits originated in showers created at the beam pipe or absorbers. The result is shown in Fig. 3.9. We found the P_T cutoff momentum of electrons to be around 4 MeV/c. A theoretical P_T spectrum of the QED pair creation [16]

indicates that roughly 98%¹ of QED backgrounds will be cutoff. A more quantitative study is in progress with a specially developed QED pair event generator from which results should be available by year's end.

¹The P_T spectrum at 4° was not available, therefore used the total P_T spectrum integrated over all the kinematical region to get this value.

Bibliography

- [1] T. Matsui and H. Satz, Phys. Lett. **B178**, 416 (1986).
- [2] C. Lourenco (for NA38), Talk at QM'93 (to be published).
- [3] D. Lissauer and E.V. Shuryak, Phys. Lett. **B253**, 15 (1991).
- [4] K. Kajantie, J. Kapusta, L. McLerran, and A. Mekjian, Phys. Rev. **D34**, 2746 (1986).
- [5] E. Shuryak and L. Xiong, Phys. Rev. Lett. **70**, 2241 (1993); K. J. Eskola and M. Gyulassy, Phys. Rev. **C47**, 2239 (1993).
- [6] L. Van Hove, Phys. Lett. **B118**, 138 (1982).
- [7] M. Kataja, P. V. Ruuskanen, L. McLerran, and H. von Gersdorf, Phys. Rev. **D34**, 2755 (1986).
- [8] F. Wilczek, Talk presented at QM'93 (to be published).
- [9] B. Müller and X.-N. Wang, Phys. Rev. Lett. **68**, 2437 (1992).
- [10] M. Gyulassy and M. Plümer, Phys. Lett. **B243**, 432 (1990).
- [11] S. Pratt, Phys. Rev. **C33**, 1314 (1986); G. F. Bertsch, M. Tohyama and M. Gong, Phys. Rev. **C37**, 1896 (1988).
- [12] J. D. Bjorken, Phys. Rev. **D27**, 140 (1983).
- [13] D. K. Srivastava, J. I. Kapusta, NUC-MINN-93/3-T; also Talk presented at QM'93 (to be published).
- [14] A. Chilingarov *et al.*, Phys. Lett. **B83**, 136 (1979); K. Kodama *et al.*, Phys. Lett. **B263**, 579 (1991).
- [15] R. Vogt, Atom. Nucl. Data Tabl. **50**, 343 (1992).
- [16] J.-S. Wu, M.J. Rhoades-Brown, C. Bottcher and M.R. Strayer RHIC Detector Note 4. May 1992

Chapter 4

Magnets

4.1 Overview

Two magnets, a central magnet (CM) and a muon magnet (MM), comprise the PHENIX magnet subsystem. Although the muon identifier steel (μ ID) is part of the muon arm (see Chapter 10), it is integrated with the magnet components for design, structural, and scheduling reasons. The design of these magnets has increased in complexity and sophistication since the PHENIX CDR was submitted and the muon identifier section is completely different than in the CDR.

4.2 Changes

The major changes in the magnet subsystem since the CDR include:

1. The CM is permanently mounted on a transport system utilizing rollers and tracks.
2. The Inner Coil has a square cross-section as compared to the stepped configuration shown in the PHENIX CDR.
3. The region near the RICH phototubes has additional shielding to reduce stray magnetic fields.
4. The CM and MM steel designs have been optimized based on the manufacturing capabilities of the Russian enterprises that will build them.
5. The μ ID section now uses steel walls instead of concrete to shower hadrons and aid in muon identification.
6. The notch in the muon magnet piston which accommodates tracking chamber station two has moved forward to improve the performance of the system.
7. The bore hole in the MM piston has been increased to a radius of 11 cm.

Figure 4.1: Isometric view of the PHENIX magnet subsystem.

8. The forward third of the MM lampshade (the "Teacup"), is a separate semi-permanent structure that will be used to mount muon tracking chamber stations one and two.

These modifications are shown in Figures 4.1, 4.2, and 4.3. Figure 4.1 shows an isometric view of the central and muon magnets with the μ ID absorber walls in place. Figure 4.2 shows a cross-sectional view of the CM and MM.

The CM transport frame shown in figure 4.1 is approximately 30 inches tall and provides a hardened steel bed for the CM to roll on. The tracks are temporary, except those immediately under the CM, and will be built in short sections that can be laid out on the floor of the detector hall to allow the CM to be pulled straight back to the far wall. Alternatively, it can be turned 90 degrees and rolled out into the assembly hall. The CM will have a hydraulic drive mechanism that can pull it 1.5 meters along the track into a maintenance mode configuration. The detectors are designed so that no major disassembly or uncabling is required to get the CM into this position. For major maintenance operations, we are

Figure 4.2: Configuration control drawing showing the CM and MM in cross-section.

designing the transport system so it can move the CM anywhere in the hall in a single 8 hour shift.

Figure 4.2 shows the cross-section of the new CM inner coil. The assembly will be square in cross-section and cylindrical when viewed from its front face. This geometry requires less power than the previous design and generates less heat in the vicinity of the central detectors. This geometry was made possible by recent changes to the back side of the CM pole face. Previously, this region was shaped in a complex pattern to minimize the dE/dX of the high angle muons passing through the pole. Now, this is no longer necessary and so both the design of the steel pole and the Neutron and gamma ray absorbers are greatly simplified and the amount of steel behind the inner coil has increased. This is good because this region was heavily saturated in the previous (CDR) design and required more power to generate the desired field. With this design, we have succeeded in keeping the temperature in this region below 30 degrees Celsius and the new coil design is simpler and cheaper to build.

Figure 4.2 shows a detail view, labelled CM profile, of the region where the phototubes for the Ring Imaging Cherenkov will be mounted. The new feature is the 20 cm tall "lip" on the outer radius of the CM pole face. This modification decreases the residual magnetic

Figure 4.3: Cut away view of the CM Steel illustrating the split pole pieces.

field near the first few rows of tubes from 150 gauss to 20 gauss. We anticipate that normal mu metal shielding around the tubes will decrease the field an additional factor of 10 to 20 which should be satisfactory for normal phototube operation.

Figure 4.3 shows the CM steel cut away to illustrate its components. The design of the CM pole face has been modified to accommodate the maximum forging sizes possible at the Azura Steel plant in St. Petersburg, Russia. In addition to the changes on the back side of the pole, mentioned above, the radius of the CM has been reduced and the CM pole has been split into three pieces to decrease the size of the raw ingots of steel. This is necessary to prevent the steel from taking too long to cool down and thus forming large crystals that would have poor chemical and magnetic properties, as well as interfere with ultrasonic detection of voids and cavities. The split occurs between the inner coil and outer coils and was placed in a region that is not magnetically saturated.

Figure 4.4: Cut away view of the MM Steel illustrating the MM coil power leads.

The flux return yokes have also been sized to match the forging capabilities at Azura. Instead of being built of one large piece, we now propose to build each out of three pieces.

Figures 4.1 and 4.4 show the major modifications to the Muon Magnet. Principally, we now have a steel muon identifier. It is composed of layers of steel and proportional tubes to cause hadrons to shower but will identify the straight through tracks of the muons. The first layer of the identifier is the 30 cm backplate of the muon magnet. Then the nominal design calls for two 10 cm thick layers of steel followed by three 20 cm thick layers of steel; each separated by a 15 to 30 cm wide gap.

The bore of the MM piston has been increased from a radius of 6 cm to a radius of 11 cm to allow both p-A and A-p running through the MM. The larger bore is needed because the heavier beam sweeps a larger arc in the final DX focussing magnet and causes an asymmetry in the position of the beams as they converge to the interaction point. Thus the DX magnet must have a wide bore and the MM must have a large bore to accommodate beams in both

directions. The large bore is also useful for vacuum fittings, valves, and heater jackets.

We have also made changes to the design of the lampshade which surrounds the MM piston. Previously the lampshade consisted of 8 trapezoidal panels that extended the full distance from the backplate to the front of the MM piston. Since accessibility to the area under the forward part of the lampshade has not proved useful, we have adopted a new design where the forward third of the lampshade is bolted semi-permanently in place. It can be rigged and lifted out of the way for major maintenance but will not be moved for any other reason. The partition line between the forward and back sections coincides with the back edge of the support fin for the MM piston. This makes it possible to use the "teacup" as a robust assembly for mounting the tracking chambers at stations one and two. The chambers will be bolted directly to the teacup at both locations.

Figure 4.4 shows the plans for bussing the power from the floor of the hall to the MM coils. The water cooled buss leads will extend upwards directly from the floor but will be in the shadow of one of the detector octants at 22.5 degrees from the vertical. The hydraulic lines for cooling the coils will go horizontally to the backplate and then go to the floor in the gap between the last detector and the backplate.

Finally, Figure 4.4 shows the new location of the notch for positioning tracking chamber number two. It has moved forward to accommodate the new lampshade design and to provide a simple and robust location for mounting the chamber. This change improves the resolution of the tracking system slightly although the precise location of station two is not critical. There is a broad region near the half way point between stations one and three where the resolving power of the tracking system is maximized. Station two now sits near the peak of this region for all trajectories between 15 and 35 degrees. See figure 4.5.

In the following sections, we describe the specific details of each of the major components of the CM and MM subsystems. We will not discuss the CM or MM detectors any further except to point out that the residual field near the CM detectors is quite low; below 10 gauss in the vicinity of the EM Calorimeter. (See figure 4.6.)

Figure 4.5: Sagitta of tracks for 10 GeV muons in the muon tracking system.

4.3 Central Magnet

The Central Magnet (CM) uses two pairs of resistive coils. Each pair is mounted in concentric, circular, annular counter bores in the poles of the magnet. The coils use conventional square, hollow, copper conductor. (The inner coil will not be installed for the first day of running at RHIC. It's cost has been deferred.) The poles, as well as the upper and lower flux return steel, are made from low carbon steel. The CM is designed so that it can be moved as a complete assembly along the beam axis (z-direction) and horizontally transverse to the beam axis (x-direction) into the assembly area of the detector hall. This movement provides access to the MM and CM interface area for maintenance and new construction.

Figure 4.6: Total B field in the Central Magnet at Various Radii

4.3.1 Inner and Outer Coils

Description of the Inner Coil

The inner coils (2 required) are made of 6 bifilar wound double pancake coils which are vacuum epoxy impregnated in three sets of two double pancake assemblies. All three pancake assemblies are identical in size and configuration. The inside diameter (ID) of the coil is 1084.7 mm which leaves a 6 mm radial clearance between the ID of the coil and the inside surface of the counterbore (to allow for coil installation onto the steel poles). The outside diameter (OD) of the coil is 1553.5 mm and the coil is 300 mm wide. When installed, the three sets of two double pancake assemblies are stacked together, side by side, and rest on the annular counterbore machined into the iron pole pieces of the Central Magnet.

The inner coils are designed to operate at 293,000 amp-turns. They utilize 21.5 mm square, hollow, copper magnet conductor with a 13.7 mm hole. The coil configuration is a 120 turn conductor package (40 turns per assembly, times 3 assemblies). The total length of conductor required is 497 meters per coil resulting in a coil weight of 1408 kg. Twelve water circuits corresponding to 4 circuits per potted pancake subassembly (bifilar wound) provide

a pressure drop of 60 psig with a maximum water temperature of 26.2 degrees (inlet water temperature is assumed to be 20 degrees) and this corresponds to an average conductor temperature of 23.1 degrees. The copper conductor will first be insulated with a "half-lap" of mylar tape, then enclosed in a continuous dacron sheath, wound on a coil winding form, ground wrapped with wide dacron cloth tape and epoxy impregnated under vacuum in a potting mold. Both before and after vacuum impregnation, the coil will be electrically checked for shorts and hydrostatically pressure checked and flow checked through each of the 12 water circuits. The required water flow rate is approximately 375 liter per minute for each coil assembly.

The inner coil requires 2442 amps at 67.1 volts using 164 kwatts (293,000 amp-turns) per coil. The coil can theoretically achieve 883,924 amp-turns (195% increase) by allowing the outlet water temperature to rise to 80 degrees using a power supply of 1573 kwatts (7199 amps at 218.4 volts) per coil.

Description of the Outer Coil

The outer coils (2 required) are conceptually identical in design to the inner coils. They are made of 6 identical (except for the electrical flag connections) bifilar wound double pancake coils with each double pancake assembly being individually vacuum epoxy impregnated. They have an inside diameter of 3200 mm and an outside diameter of 3729 mm and a width of 300 mm. There is a 6 mm clearance between the ID of the coil and the inside surface of the counterbore to allow for coil installation onto the steel poles. As in the inner coil, the six double pancake assemblies that make up the outer coil are nested together side by side and rest on the circular counterbore machined into the iron pole pieces of the Central Magnet.

The outer coils are designed to operate at 247,500 amp-turns but they utilize a copper conductor which is 20.3 mm square with a 12.8 mm hole. The coil stack is a 144 turn conductor package (24 turns per double pancake, times 6 double pancakes). Total conductor length is 1578 meters per coil resulting in a coil weight of 4003 kg. Maximum outlet water temperature rise is 45.6 degrees and this corresponds to an average conductor temperature of 32.8 degrees. The required water flow rate is 166 liters per minute per coil.

The outer coil consumes 1719 amps at 174 volts for a total power consumption of 300 kwatts (247,500 amp-turns) per coil. The coil can theoretically achieve 367,200 amp-turns (48% increase) by allowing the outlet water temperature to rise to 80 degrees and using a power supply of 699 kwatts (2550 amps at 274 volts) per coil.

The key physical parameters of the inner and outer coils for the central magnet are listed in Table 4.1.

Fabrication

The fabrication of the central coils is *build to print* from drawings provided by the PHENIX collaboration. The winding mandrel and the potting fixture are specified by the collaboration, but the detailed design of this hardware is expected to be the responsibility of the coil vendor, with the design being approved by the collaboration's magnet design group. The

Table 4.1: Physical and Operating Parameters of the PHENIX Central Magnet Coils

Parameter	Inner Coils	Outer Coils
Number of Coils	2	2
Amp-Turns	293,000	247,500
Configuration	6 Dbl. Pancakes	6 Dbl. Pancakes
Conductor Material	Copper	Copper
Inside Diameter (m)	1.08	3.20
Outside Diameter (m)	1.55	3.73
Number of Turns	120	144
Conductor Size (mm)	21.5 (square)	20.3 (square)
Conductor Hole (mm)	13.7 (diameter)	12.8 (diameter)
Conductor Length (m)	497	1578
Current (A)	2442	1719
Voltage (V)	67	174
Power (kW)	164	300
Flow Rate (L/min)	375	166
Weight (kg)	1408	4003
Average Coil Temp. (°C)	23	33

construction of the inner and outer coils is straightforward. The design utilizes standard coil winding and vacuum potting techniques.

4.3.2 Pole Pieces and Support Structure

Description

The CM steel structure consist of two pole pieces, two large flux returns and a support structure. Each pole piece will be made of AISI 1010 steel (or equivalent). The large size of these poles exceeds world fabrication capacity, so each is split into three concentric parts. (See figure 4.3.) The outermost part, called the pole ring or "eyebrow", has a finished weight of 60 tons and interfaces to the flux returns. The "pole piece center" fits into the pole ring and weighs 61 tons. The center element of the pole, the "pole piece plug", surrounds the RHIC beam pipe and is made from thick rolled plate and weighs seven tons. Both pole pieces have machined circular grooves to accept the inner and outer coils and are separated by a 1.2-m gap.

The two facing poles are magnetically and structurally tied together by two large flux return yokes ($2 \times 1.6 \text{ m}^2$ in total cross section), which attach to the top and bottom of the poles. The large size of these return yokes (90 tons) dictates that they be split into a yoke subassembly consisting of two pole keys and a smaller yoke.

The entire pole assembly is supported by four outriggers that bolt to the bottom return yoke. These outriggers rest permanently on Hilman rollers sitting on a guide track. This

structure is designed to allow for transport of the CM over a short distance (1.5 m) along the beam axis (z -direction) and horizontally transverse to the beam axis (x -direction) several meters into the assembly area of the detector hall. A transporter consisting of Hilman rollers, hydraulic jacks, and a steel structure is planned for this purpose.

The transporter assembly will be capable of lifting the entire CM and a hydraulic control system will ensure that all four jacks share the load equally. Once lifted, the assembly is rolled on transport tracks using the Hilman rollers. These transport tracks, consisting of W24 $\times$ 162 wide flange beams, are laid directly on shims on the concrete floor. A pair of these beams running from the experimental area to the high bay roll up door will be shimmed and stabilized by cross beams. The use of beams rather than flat plates is needed to spread the concentrated load of the relatively small rollers over a larger area of the concrete floor. Each beam has section gussets approximately every 60 cm to stabilize the web from buckling under the concentrated load.

Fabrication

Based on vendor estimates, it will take approximately 15 months to fabricate, machine and transport the CM steel to BNL. The fabrication cost estimate for the pole pieces is provided in the form of a vendor estimate made for budgetary purposes. The balance of the magnetic steel and steel support structure is based on current market prices for forged and plate steel, and generally accepted fabrication practices.

The entire magnet assembly, excluding the transport system, will be fabricated and prefit in Russia. The design will accommodate shipping constraints, and allow 100% bolt up assembly at BNL. A strict quality assurance program will be implemented in Russia to minimize risk to the program. An installation study will be performed by an experienced US rigging company to determine equipment, manpower, space, time, and logistics requirements. And a contract will be negotiated with Long Island Lighting Company (LILCO) for use of their docking facility at Shoreham Nuclear Power Station, which is in close proximity to BNL. This facility will be used to off load major steel components from barges to be transported over local roads using multiaxis, self jacking trailers.

4.3.3 Assembly and Installation

The CM will be fabricated outside the United States. The CM outer coils will be fabricated in Japan and transported to BNL as completed double pancake units. It will take approximately 5 man-months to fully install the coils on the magnet pole pieces and check them out.

The CM steel will be fabricated in Russia. Assembly and installation of the two CM pole pieces includes all magnetic flux return steel, the entire steel support structure, and the Central Magnet transport system. It includes all structural alignment and fit-up of the magnet, and doing all subsystem integration checkouts. Once the magnet assembly is complete, the magnet will be rolled into place and installed in its operating position on the RHIC beamline.

Assembly of the CM will require 10 technicians 2 months to construct and assemble the support structure, assemble the two pole pieces onto the support structure, and align and

mate all hardware. This work includes welders, riggers and general technicians as well as their own supervisory staff. This time includes the moving of the subassembly from the high bay area (where assembly is taking place) into the beamline itself. In addition, 2 alignment surveyors will be needed for 3 months. Their effort will be shared equally between the Muon Piston Magnet so the amount charged will be a total of 3 man-months. Total manpower required for the central magnet is 23 man-months.

4.3.4 Testing

The Coil tests include running the CM outer coils at full rated current, magnetically mapping the magnetic field produced in the gap and other key locations around the magnet as required, and checking all instrumentation and system performance.

It is estimated that it will take 3 technicians 1 month to test and map a coil pair (including reverse polarity).

Estimated cost for the magnetic mapper, data acquisition system and miscellaneous hardware is \$100k but this cost is shared equally between the Central Magnet Outer Coils and the Muon Piston Coils.

4.3.5 Cost and Schedule

The cost estimate and schedule for the central magnet is presented in Table 4.2. In this table the CM outer coil (5.1.1), inner coil (5.1.2), and steel (WBS 5.1.3) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

The contingency on the magnet cost will be calculated on the basis of the costs we expect to incur (ie. based on the cost of fabrication in Russia) and on portions of the cost that are not deferred and not covered by Japan. The breakdown of all costs into DOE, Japanese, Russian, other non-DOE, and deferred components is shown in the summary cost estimate in Chapter 15.

The proposed timeline for the CM delineates the time required for engineering, design, analysis, procurement, fabrication, assembly, installation, and testing of the CM. It is important to note that the critical path for the CM is the production of the ingots and final machining of the open die forgings for the large pole and yoke assemblies (~15 months). The CM is installed after the MM and the Mu ID steel is in place in the detector hall. This precludes the CM from being identified as the critical path item in the magnet subsystem; however, it is on the critical path in terms of the entire detector schedule. Both the CM and the MM must be installed and tested in the detector hall before any detector subsystem is installed.

Table 4.2: Cost Estimate and Schedule for the PHENIX Central Magnet

WBS	Description	Total Cost	1992			1993				1994				1995				1996				1997				1998		
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	
5.1.1	Central magnet outer coil	\$926,669.78																										
5.1.1.1	CM outer coil EDIA	\$203,129.78																										
5.1.1.2	CM Outer Coil Prototype	\$0.00																										
5.1.1.3	CM Outer Coil Fabrication	\$625,250.00																										
5.1.1.4	CM Outer Coil Assembly	\$28,000.00																										
5.1.1.5	CM Outer Coil Test	\$70,290.00																										

WBS	Description	Total Cost	1992			1993			1994				1995				1996				1997				1998	
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	
5.1.2	Central magnet inner coil	\$470,122.62																								
5.1.2.1	CM inner coil EDIA	\$111,222.62																								
5.1.2.2	CM Inner Coil Prototype	\$0.00																								
5.1.2.3	CM Inner Coil Fabrication	\$324,600.00																								
5.1.2.4	CM Inner Coil Assembly	\$14,000.00																								
5.1.2.5	CM Inner Coil Test	\$20,300.00																								

WBS	Description	Total Cost	1992			1993				1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	
5.1.3	Central magnet steel	\$1,613,653.07																											
5.1.3.1	CM Steel EDIA	\$336,676.91																											
5.1.3.2	CM Steel Prototype	\$0.00																											
5.1.3.3	CM Steel Fabrication	\$1,077,075.36																											
5.1.3.4	CM Steel Assembly	\$201,100.80																											
5.1.3.5	CM Steel Test	\$0.00																											

4.4 Muon Magnet

4.4.1 Coil

Description

The muon piston coil (1 required) is a bifilar wound solenoidal coil (having two layers) which will be vacuum epoxy impregnated onto its epoxy-fiberglass winding form. It is made up of two identical sub-coils with the only difference being that one is slightly smaller in diameter than the other. Each coil is 676 mm in overall length. The small coil ID is 1624 mm and the OD is 1740 mm. The large coil ID is 1881 mm and the OD is 1997 mm. (See figure 4.4.)

The muon piston coil operates at 300,000 amp-turns. It utilizes a hollow copper magnet conductor which is 24.1 mm square with a 15.5 mm hole. The coil configuration is a 102 turn conductor package comprised of two layers of solenoidal type windings. Each sub-coil has 51 turns and the total conductor length required is 581 meters for both coils. The overall coil assembly weight is 2047 kg; the small coil weighs 951 kg and the large coil weighs 1096 kg. Each bifilar wound conductor segment has its own individual water circuit in addition to each coil layer having its own parallel water supply giving rise to a total of 8 individual water circuits. The LCW system provides cooling water to achieve a pressure drop of 60 psig with a maximum water temperature rise of 35 degrees which corresponds to an average conductor temperature of 25.6 degrees for the smaller coil and 27.1 degrees for the larger coil. The copper conductor will be insulated with a "half-lap" of mylar tape enclosed in a continuous dacron sheath, wound on a coil winding form, ground wrapped with wide dacron cloth tape and epoxy impregnated under vacuum in a potting mold. Both before and after vacuum impregnation, the coils will be electrically checked for shorts and hydrostatically pressure checked and flow checked through each of the 8 water circuits. The required water flow rate is approximately 257 liters per minute for the entire coil assembly; 134 lpm for the small coil and 123 lpm for the large coil.

The muon coil requires 2941 amps at 76.6 volts (35.5 volts for the small coil and 41.1 volts for the large coil). The total power consumption is 225 kwatts (300,000 amp-turns). The coil can theoretically achieve 587,739 amp-turns by allowing the outlet water temperature to rise to 80 degrees and using a power supply of 944 kwatts (5762 amps at 164 volts).

The key physical parameters of the MM coil are listed in Table 4.3.

Fabrication

The fabrication of the MM coil is *build to print* from drawings provided by the PHENIX collaboration. The winding mandrel and the potting fixture are specified by the collaboration, but the detailed design of this hardware is the responsibility of the coil vendor, with the design being approved by the collaboration's magnet design group.

Table 4.3: Physical and Operating Parameters of the PHENIX Muon Magnet Coil

Parameter	Muon Coil
Number of Coils	2
Amp-Turns	300,000
Configuration	Solenoid
Conductor Material	Copper
Large Coil Diameter (m)	2.00/1.88
Small Coil Diameter (m)	1.74/1.62
Coil Length, each (m)	0.68
Number of Turns, each	51
Conductor Size (mm)	24.1 (square)
Conductor Hole (mm)	15.5 (diameter)
Total Conductor Length (m)	581
Total Current (A)	2941
Voltage (V)	77
Total Power (kW)	225
Total Flow Rate (L/min)	257
Total Weight (kg)	2047
Average Coil Temp. (°C)	27
Field Integral at $\theta = 15^\circ$ (T·m)	0.72
Piston Length (m)	4.3
Coil Thickness (cm)	6.0
Lampshade Thickness (cm)	8.0

4.4.2 Pole Pieces and Support Structure

Description

The basic MM magnet steel structure consists of a large tapered iron core bolted to a flux return back plate which is 30 cm thick. Eight trapezoidal steel plates, 8 cm thick, are arranged in an octagon "lampshade" around the iron core. These plates provide the return flux path for the muon piston magnet. The bottom three plates bolt together to form a permanent element of the overall structure. The top five plates are split across the global Z axis into a permanent "teacup" section and five independently removable lampshade pieces. The tapered iron core will be forged as a single piece and weigh approximately 60 tons. The back plate will be constructed from three sections bolted together. The "lampshade" flux return will be fabricated from rolled plate. All magnetic material will be ASTM 1010, or equivalent.

The iron core, back plate, and lampshade plate assemblies form a rigid steel structure. This structure is designed to allow for access to the three sets of muon chambers via access doors in the side of the "teacup" or by removing the rear lampshade pieces. The muon

piston magnet assembly is a stationary structure which is stabilized by two support fins during detector operation. These fins can be safely removed to allow for installation of the MM coil. The endplate deflects a total of 0.8 mm under these conditions but will elastically restore its position when the fin is re-installed.

Fabrication

The fabrication of the Muon Magnet steel and structure will be a build to print from PHENIX provided drawings. All components will be made from pieces as large as can be fabricated by shops in St. Petersburg, Russia. For example, the muon piston core will be forged from a single ingot while the back plate will be made from 3 major slab sections.

The Entire MM magnet assembly will be fabricated and prefit in Russia. A strict quality assurance program will be implemented in Russia to minimize risk to the program. And installation will be performed by an experienced US rigging company. A contract will be negotiated with Long Island Lighting Company (LILCO) for use of their docking facility at Shoreham Nuclear Power Station which is in close proximity to BNL. This facility will be used to off load major steel components from barges to be transported over local roads using multiaxis, self jacking trailers.

The piston of the MM is the critical path item of the magnet system. Based on vendor estimates, it will take approximately 15 months to fabricate, machine, and transport the piston to BNL.

4.4.3 Assembly and Installation

The MM coils will require special handling and installation fixturing to prevent possible damage of the coils. The cost of this fixture is included in the coil fabrication section above. It will take approximately 3 technicians 1 month to completely assemble and install the coils onto the piston. This includes the fixturing and preparation required for the installation.

The assembly and installation schedule and the cost estimate for the MM steel include the piston, back plate, the eight steel plate elements that make up the lampshade flux return, teacup, steel support structure, installation of the solenoidal coil onto the piston, all associated electrical and plumbing hook-ups in the detector hall, all structural alignment and fit-up of the magnet, and magnet subsystem integration and checkout.

Magnet assembly and installation will require 6 technicians 3 months to construct and assemble the magnet support structure, completely assemble the piston onto the back plate, assemble the lampshade and teacup, and align and mate all hardware. This work includes welders, riggers and general technicians as well as their own supervisory staff. In addition, 2 alignment surveyors will be needed for 3 months. Therefore, the total manpower required for the muon piston magnet is 21 man-months.

4.4.4 Testing

The magnet steel will be certified for chemical composition, magnetic permeability, and percent void fraction through a series of certification tests and nondestructive testing techniques.

Dimensions will be checked through all phases of fabrication and assembly.

Field mapping of the MM will be performed with a Hall probe after assembly is completed in the detector hall. A field map with better than 1% precision is required.

4.4.5 Cost and Schedule

The cost estimate and schedule for the muon magnet is presented in Table 4.4. In this table the MM coil (5.1.4), and steel (WBS 5.1.5) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

The contingency on the magnet cost will be calculated on the basis of the costs we expect to incur (i.e., based on the cost of fabrication in Russia) and on portions of the cost that are not deferred and not covered by Japan. The breakdown of all costs into DOE, Japanese, Russian, other non-DOE, and deferred components is shown in the summary cost estimate in Chapter 15.

Russian steel fabrication costs are based on an agreement between PNPI and BNL (see Technical Attachment 2). Trans-atlantic shipping has been estimated using a Russian cost of \$175k for the CM and MM together, or \$87.5k each. This trans-atlantic shipping costs carry a 100% contingency, however, due to the uncertainty of future Russian shipping costs.

The proposed timeline for the MM delineates the time required for engineering, design, analysis, procurement, fabrication, assembly, installation, and testing of the MM. The PHENIX collaboration is planning to install the MM steel during RHIC facility construction because installation at a later date would incur additional costs. The schedule delineates the time required for engineering, design, analysis, procurement, fabrication, assembly, installation, and testing of the MM. It is important to note that the MM is the critical path item of the magnet subsystem. In particular, the procurement and fabrication of the large iron piston (~15 months) is the critical path component of the subsystem. As mentioned previously, the MM is assembled in-situ in the detector hall. It is the first major subsystem installed in the detector hall after the μ ID steel and it must be completed prior to installation of the CM or any of the detector subsystems.

4.5 Muon Identifier Steel

4.5.1 Description

The muon identifier (μ ID) will consist of several layers of low carbon steel plate and proportional counter panels located behind the the Muon Magnet assembly. The design calls for five layers of steel after the 30 cm thick endplate of the MM. The layer thicknesses are 10, 10, 20, 20, and 20 cm thick. There is also the possibility of adding a fourth 20 cm layer onto the back of the μ ID stack. Other configurations are possible and the final design will depend on the availability of surplus steel at the time of manufacture. In addition, bracing and anchoring structures must be designed to accommodate safety requirements. The assembly must also allow for the μ ID counter panels to be installed, secured, and aligned in 15 cm a wide gap between the steel layers. See Chapter 10 for a discussion on the μ ID instrumentation.

4.5.2 Fabrication

The fabrication of the muon identifier steel and structure will be a build to print from PHENIX provided drawings. The large slab pieces will be produced from low grade rolled plate generally rejected for other purposes because of its chemical composition. No specification on mechanical properties or chemical composition will be given in the purchase order other than minimum iron content, uniformity, and continuity. Domestic estimates have been obtained based on current stock availability at vendor's locations, with some projections on future availability. Since this is mill reject material, there is no guarantee on pricing or stock inventory until actual time of RFQ.

4.5.3 Assembly and Installation

The μ ID steel must be installed in the detector hall before all other magnets and detectors. The primary reason the steel goes in the hall first is because of its size and location in the detector hall. It is a system consisting of hundreds of tons of steel that is best assembled on the floor of the detector hall and then tilted up into its final position in order to accommodate shipping size and weight limitations and to minimize cost.

If the muon subsystem is installed after RHIC commissioning and operation, the ring will be shut down during the MM construction period. Installed detectors, which are vulnerable to damage, will have to be removed and isolated. The CM assembly and final focussing quadrupoles may also need to be partially removed. These activities will delay the start of MM installation by at least a month. Also, cable trays, piping, and electronics racks could take another month to remove and two months to re-install and test. We estimate that the additional cost of installation of the MM after RHIC commissioning could add 10 man-months of labor (\$100K) and an additional three months of lost time beyond what it takes to install the MM steel on Day-1.

Chapter 5

Inner Detectors

5.1 Silicon System Overview

The PHENIX multiplicity and vertex detector (MVD) provides event characterization, a centrality trigger, the collision vertex position along the beams, and the measurement of fluctuations in the charged particle distributions. The system has silicon strip detectors arranged in two concentric barrels around the beam pipe in the center of PHENIX, and endcaps consisting of a single layer of silicon pad detectors. The total area of silicon is approximately 0.5 m^2 . Both pad and strip detectors are single sided and are read out with identical front end electronics mounted on hybrid circuits. The pulse height from each strip or pad is transferred to the data acquisition system via optical fiber links. The total silicon channel count is 34,816.

The central barrels have radii of 5 and 8 cm and lengths of 64 cm. The single-event coverage about midrapidity is five units in η for the inner barrel layer and four units for the outer layer; however, the η coverage for event-averaged quantities is extended due to the variation of the vertex position along the direction of the colliding beams. The inner and outer barrels are composed of silicon strip detector wafers, with an area of $5 \times 7 \text{ cm}^2$ and $5 \times 7.5 \text{ cm}^2$, respectively. The strips have a pitch of $200 \text{ }\mu\text{m}$ and are oriented perpendicularly to the beam. Each barrel layer has six azimuthal sectors, and the two barrel layers are mounted on a common support structure made of Rohacell foam.

The disk shaped end caps, which comprise pads ranging in size from approximately 2×2 to $5 \times 5 \text{ mm}^2$, are positioned in z at -35 and $+35 \text{ cm}$. The pseudorapidity coverage is 2.7 for a single event occurring at $z = 0$. The η segmentation is more than adequate for fluctuation studies, and pulse height analysis allows the unfolding of multiple hits in one pad. The azimuthal segmentation, along with the full azimuthal coverage, makes it possible to look for fluctuations in two dimensions. The highest channel occupancies for central events are 0.35–0.46 for $dN_{ch}/d\eta = 1000$.

The vertex position can be located using hits in the barrel for z position between -40 and $+40 \text{ cm}$, which covers the 2σ length of the interaction region. The multiplicity trigger is formed using the sum of silicon channels (both pads and strips) above a threshold corresponding to approximately 0.25–0.33 of a minimum ionizing particle (mip). This threshold

will be tuned to optimize the trigger performance, and a rather low threshold will be used for zero suppression in the data stream.

The front-end electronics are optimized for minimum size, power dissipation and cost, while the design allows for the unfolding of multiple particle hits in each silicon pixel. The dynamic range is 8 mips, and the signal to noise better than 10:1 for a single mip. The shaping time of the preamplifier is 50–60 ns, limited by the bunch crossing time to avoid pileup in $p + p$ collisions and allow use of the silicon in the LVL-1 trigger. Although the multiplicity in a single Au + Au collision is very high, the expected luminosity at RHIC is such that special considerations of radiation hardness for either the detectors or the front end electronics is not required [1].

5.2 Silicon System Changes

Design work after the CDR has resulted in a number of changes in the silicon detector coverage, mechanical structures, cooling system, and front end electronics, which include:

1. Removal of a large portion of the central silicon detectors in the upper two thirds of the outer barrel layer.
2. Progress in the design of the front-end electronics and hybrid circuit for the MVD, making them more similar to the rest of PHENIX, lowering the cost, and improving the applicability of the design to other experiments.
3. Relocation of the LVL-1 trigger to outside the silicon enclosure and concentration of the electronics below the detector.
4. Extensive studies and tests of air cooling the MVD, which would lower the cost and minimize material in the path of the particles, and would eliminate the possibility of any coolant spills in the region of the silicon.

The most significant changes have been driven by detailed studies of electrons arising from photon conversions in the MVD. These studies have been made using PISA to integrate the MVD with the magnetic field, the central arms, and other sources of conversion electrons in PHENIX. Details of this simulation are presented in Chapter 3. This approach allowed a careful study of the tradeoff between large silicon coverage to maximize the vertex finding efficiency in low multiplicity events, and reduction of the material to minimize photon conversions in high multiplicity events. The optimum solution to these competing requirements was achieved by removing a large portion of the central silicon detectors from the upper two-thirds of the outer barrel layer, and concentrating the electronics as far from the MVD as possible at the bottom of the detector support carriages. Removal of the silicon from only one barrel layer preserves the triggering capability, and the coverage for the multiplicity and fluctuation measurement.

The updated geometry is shown in Fig. 5.1. The consequences of these changes to channel counts, mechanical design, and cooling requirements are detailed in the mechanical design

Figure 5.1: Layout of silicon detectors in the MVD.

section below; the simulation results quantifying the tradeoff and leading to these design changes are presented in the performance section; and the improved design for the front-end electronics and hybrid circuit are described in the electronics section.

5.3 Mechanical

5.3.1 Silicon detectors

The inner barrel consists of silicon detector wafers with an area of $5.3 \times 5.0 \text{ cm}^2$, and the outer barrel uses two different wafers with an area of $5.3 \times 7.0 \text{ cm}^2$ and $5.3 \times 7.5 \text{ cm}^2$. The maximum size possible varies somewhat with the detector supplier; we will make the choice of supplier and finalize this geometry issue in the next year. The strip pitch on both detector sizes is $200 \text{ }\mu\text{m}$, yielding 256 channels per detector. The silicon detectors are arranged in six azimuthal sectors consisting of 12 detector wafers along the beam direction. The PISA studies described in more detail below resulted in the removal of the 8 central wafers in top four sectors of the outer barrel. This is shown in Fig. 5.1.

Each pad endcap is still composed of 12 “sectors”, with each “sector” being manufactured from one 4-in. wafer. The inner radius is 5 cm and the outer radius is 12 cm. There have been no changes made to the silicon pad detectors or layouts. We are, however, currently

Table 5.1: Silicon Detector Channel Count

Detector Type	Size (cm ²)	ϕ sectors	z sectors	Total Channels
Inner Barrel	5.3×5.0	6	12	18,432
Outer Barrel	5.3×7.5	2	12	6144
Outer Barrel	5.3×7.5	4	4	4096
End Cap	6.3×7.0	12	2	6144
Total				34,816

investigating the possibility of having detectors with double metallization layers. This would bring traces from the pads out to the edge of the detector, and allow wire bonding directly to the front end electronics with a much simpler kapton cable.

The detector types and the updated number of wafers, and channel counts are summarized in Table 5.1.

5.3.2 Support Structures

The support structure for the MVD must use a minimum of material thickness to prevent background from photon conversions, but it must also be rigid to support and protect the silicon inside. The MVD will reside inside an enclosure, which opens in a “clam-shell” fashion to allow installation around the beam pipe. Each half-cylinder consists of aluminum endplates, with two graphite composite struts to connect and support them. The endplates mount to the magnet with graphite composite struts, and upon installation are brought together to form a complete cylinder. The cylinder is supported from one magnet pole face and held with a spring-loaded bar to the other pole face to provide vibrational isolation. The endplates are separate from the motherboards for the silicon pad detectors, which are located inside of the support endplates. This arrangement is shown in Fig. 5.2.

The half-cylinders have an enclosure of Rohacell foam, covered on each side with aluminized mylar, to provide environmental isolation of the MVD and mechanical protection. Such an enclosure allows the electronics to be cooled with air below the ambient temperature of the experimental hall, if necessary. The enclosure also electrically shields the detectors and front end electronics from the other elements of PHENIX to reduce pickup of noise. The aluminized mylar layers can also be held at a fixed potential, should additional shielding be necessary.

The silicon pad detector endcaps are mounted to a G-10, or polyimide/fiber circuit board for better CTE match, motherboard that supports both the detectors and the hybrid circuits and provides DC power. The motherboards are mounted to the endplates, and have cutouts to allow standoff bars for mounting the central barrel support structure. These structural elements are shown in Fig. 5.2.

The barrel detector wafers are supported on Rohacell cages as described in the CDR. We are investigating the machining properties of the Rohacell foam to determine what fraction

Figure 5.2: Cutaway view of the silicon support and enclosure half-cylinders.

of an azimuthal sector may be machined in one piece. We would like to fabricate an entire sector to eliminate the need for glue joints in the Rohacell. Finite element analysis of the azimuthal sector structure has shown that the graphite epoxy stiffener bars described in the CDR are unnecessary due to the structural strength provided by the silicon detectors glued to the Rohacell. We will investigate the effect of omitting some of the outer silicon detectors on this conclusion.

The azimuthal sectors are mounted to an adapter plate bolted to standoff bars on the enclosure endplate. This is a thin plate with large cutouts, as shown in Fig. 5.3, fabricated from G-10 or a graphite epoxy sheet. The design of this structure is constrained by the requirement of minimizing the amount of material in front of the silicon pads. The exact geometry and thickness of this piece is still being optimized.

The PISA studies described below show that the background of conversion electrons from the MVD is smaller when the hybrid circuits are located as far from the electron arms as possible in azimuth. To enter the electron arm, the path of the electrons must be significantly curved by the magnetic field. For a larger distance in azimuth the required bend becomes larger, implying softer electrons captured into the electron arm acceptance. Maximizing this distance results in accepting electrons which fall mostly below the angular acceptance cutoff of the ring imaging Cherenkov counter (RICH). Consequently, we have modified the mounting scheme of the electronics to concentrate them at the bottom of the barrel, as shown in Figs. 5.1, 5.4. The electronics are radially as far from the vertex as possible within the MVD radial distance allotment. This helps with rejection of background electrons and allows use of the graphite composite enclosure struts to help support the electronics packages (Fig. 5.5). We have determined that the resulting length increase of the kapton cables is acceptable; the capacitance is still low enough for good preamplifier signal to noise performance.

Figure 5.3: Silicon enclosure with adapter plate for central assembly.

Figure 5.4: Cutaway view of the silicon enclosure, showing the central barrel assembly and plenum enclosure. Cables are omitted for clarity.

Figure 5.5: End view of the silicon enclosure showing the endcap, silicon pad motherboard with detectors and hybrid assemblies, and central barrel assembly with cables and plena.

The hybrid circuits are mounted to a support made of low-Z material to minimize the radiation length; we currently foresee using Rohacell foam for this. Besides supporting the hybrids, the support piece also forms a plenum to guide the cooling air. These plena provide the correct air flow geometry to allow near-laminar flow and isolate the cooling air from the rest of the MVD to minimize vibration effects. This is an important function, as the position stability of the hybrids is less critical than of the detectors. In this design, the air required to cool the electronics does not pass along the detectors. The detectors are only coupled to air flow-induced motion via the flexible kapton cables. Nevertheless, we are seeking to minimize the cooling air flow rates, and are performing R&D to measure the stability of the proposed structures with realistic air flows.

The plena, with hybrids attached, will be grouped together in each half-cylinder and enclosed to simplify the support and further channel the cooling air. A cross-section of this assembly is visible in Fig. 5.5. The assemblies are supported at both ends from the same support structure that holds the barrel azimuthal sectors. The geometry of this support

scheme is illustrated in Figs. 5.4,5.3. If necessary, additional structural stability can be provided by tying the electronics package to the enclosure struts.

5.3.3 Cabling

The silicon strip and pad detectors are attached to the front end electronics on the hybrid circuits via custom designed kapton cables. These cables allow a physical distance between the silicon wafers and hybrids, providing thermal and mechanical isolation as described above. The updated electronics location is consistent with cable lengths of 20 cm or less, which have been shown to be acceptable by the L3 microvertex detector group. The cabling scheme is shown in Fig. 5.5.

With the new electronics location, particles pass through two cables in part of the electron arm acceptance. Consequently, it is crucial to minimize the radiation length of the cable. The kapton is 50 μm thick and aluminum, rather than copper, traces will be used. The STAR SVT group has had such cables fabricated, and found their performance adequate for analog signals [2]. The performance problems noted in those tests concerned high voltage on the cables, which is not an issue for the PHENIX MVD. We will perform extensive tests on prototype cables of this kind.

The cabling for the silicon pad wafers is more complicated, dictated by the geometry of the detector. The current conceptual design, as described in the CDR, involves a cable covering the entire detector surface, with holes for solder connections to each pad. An attractive alternative, which we are investigating with suppliers of silicon detectors, would use a double metallization layer to provide traces connecting each detector pixel to bond pads on one end of the detector wafer. This would preclude the need for the special kapton cable and allow wirebonding directly to the front-end electronics. We believe that the incremental cost of fabricating the pad detectors would be offset by savings in cable design and fabrication, and fewer wire bonds in the assembly stage.

5.3.4 Cooling

We have made extensive studies of air cooling the MVD, to take advantage of the low cost and minimal material in the path of the particles as well as eliminating potential coolant spills in the region of the silicon. The results of these studies are summarized in this subsection, along with progress in the conceptual design of the cooling system.

Each row of 12 barrel hybrid circuits is mounted to a structure forming a plenum to provide a path for the cooling air. The cross section of each plenum is 2.3-cm wide by 64-cm long, dictated by the size of 12 adjacent hybrids. The endcap hybrids are enclosed by a curved plenum, similar in aspect ratio. Each endcap plenum encloses six hybrids. The plenum depth is 6 mm, to maintain a plenum aspect ratio appropriate for maintaining laminar flow. Laminar flow is desirable to minimize vibration effects. However, our mounting scheme isolates the detectors from the cooling air and may allow us to relax this requirement in exchange for the improved cooling efficiency when the air flow is no longer laminar. This is being extensively investigated in the lab by measuring the cooling effectiveness and movement

Figure 5.6: Cutaway view of the silicon support and enclosure showing the cooling air channels and the entry and exit path.

induced by various air flows in prototype plena. Air flows along one side of the hybrids only, as the other side is covered by polyimide, which has poor thermal conductivity.

Figure 5.6 shows the distribution system for the cooling air inside each enclosure half to supply the plena for the barrel and endcap silicon. The air is driven into the enclosure halves by a fan and heat exchanger located outside (we expect that this will sit at the bottom of the central magnet, below the detector active area). Air enters from one side, at the bottom of the enclosure, and is split into two streams. One stream flows through the group of plena holding the barrel hybrids, and the other passes through the half-circular plenum covering the endcap hybrids. This air then flows to the other side of the enclosure to cool the hybrids of the other endcap. The path from one side of the enclosure to the other is inside one of the graphite composite tubes serving as a structural strut, minimizing the material in the cooling system. After passing along the hybrids of the other endcap, the air then joins the barrel cooling air and exits the enclosure to be cooled and recycled, or vented outside of PHENIX.

The electronics on the hybrid circuits will dissipate 5 mW per channel, and there are 256 channels per hybrid; the total power dissipated is about 150 W. First tests of the prototype preamplifier lend confidence to this estimate of the power dissipation. The operating temperature must be maintained below 50 °C, where components fabricated in CMOS become temperature sensitive. The silicon wafers do not need cooling since they are physically separated from the electronics by the cable. The MVD outer enclosure provides humidity control to avoid any condensation problems on the detector wafers or other elements. We envision a low flow rate of dry nitrogen through the outer enclosure to provide this. The hu-

midity control and isolation of the cooling air inside the plena make it possible to use chilled cooling air. This possibility gives a safety margin should the electronics power dissipation prove larger than currently estimated, or should we identify problems with the air flow rates described below.

The concept of air cooling the MVD was tested in the laboratory with a plenum made of phenolic (a paper/resin composite plus insulating foam). The thermal load of the electronics was provided by flat resistive heaters from Minco arranged in the geometry planned for the hybrids. A range of 2.5–18V powered the heaters to simulate heat loads around the expected 15W per plenum. Eight temperature sensors recorded temperature changes. Air was blown through the plenum with a fan, and the velocity measured with a flow meter.

These tests showed that an air flow rate of 14 mi/h keeps all the hybrids in a plenum at a temperature below 40 °C. The temperature increase from the first to last hybrid is only a few degrees Celsius, while the temperature increase of the cooling air is somewhat larger, but less than 10 °C. The temperature profile measured by the sensors along the plenum agrees reasonably well (within 10%) with calculations and indicates that the air flow is in the transition region between laminar and turbulent flow. The isolation of the cooling air inside the small plena, and the physical separation between the hybrids and detectors (linked only by the flexible kapton cable), mitigate the problems of detector vibrations from the air flow. The ruggedness of the hybrid fabrication implies that some motion of the hybrids should have no adverse effect on their performance. The magnitude of the hybrid motion is now being investigated using microaccelerometers.

We have also investigated the cooling of 10mW per channel, which results in 30 W per plenum. In this case, an air flow of 22.5 mi/h maintains the hybrids at 40–42 °C operating temperature. Another option, which we are currently studying in the lab, is to use chilled cooling air.

Another change from the CDR is that the elements of the first level trigger are now located outside of the silicon enclosure. The current sums from the hybrids can be driven far enough to allow this and significantly simplify the MVD cooling problem.

5.3.5 Assembly

Subassemblies are constructed by gluing and tab- or wire-bonding the kapton cables to the strip detectors. The detector/cable system is subsequently tested. The hybrid circuits should be independently mounted, and then connected to the cables. We are investigating the use of a connector consisting of a clamp and an elastomer strip which conducts in one direction only. Such elastomers have been successfully used at both LANL and BNL. The radiation length of the clamp required is of concern, so we are also looking into the possibility of using z-axis adhesive film for this connection.

The subassemblies are glued to the Rohacell, in subunits of the barrel. We are investigating azimuthal sectors which would be straightforward to install into the enclosure, but have the disadvantage of containing many permanently fixed silicon detectors. Another option is a half wheel-shaped subunit, which contains only 6 silicon detectors, but presents larger installation challenges. The endcaps will have the pad detectors and hybrids fixed to the motherboard, which is a much simpler assembly task.

It is clear that many fixtures will be required for alignment, assembly, testing and storage of detector/cable subassemblies and barrel subunits, as well as for the completed enclosures.

5.4 Electronics

5.4.1 Front end

Front end electronics design for the silicon detector systems is underway at Oak Ridge National Laboratory, supported by PHENIX R&D funds. A preamplifier/shaper chip has been fabricated using a CMOS process and has been shown to function. Several prototype chips have been shipped to LANL for wirebonding to a silicon strip detector and further testing. This chip will ultimately have 32 channels and a shaping time of 50-60 ns, though the prototype chip has five channels only.

As the silicon detectors provide a pulse height from each strip or pad, adequate pulse-height resolution is required to separate multiple charged particle hits on a single element. The dynamic range required is 4 mip and signal to noise of better than 10:1 for 1 mip. We are designing a circuit with a dynamic range of 8 mip to provide an adequate safety margin. A preliminary measurement of the electronic noise in the prototype preamp/shaper gives a signal to noise better than 12:1 for one mip with the capacitance of a silicon strip or pad.

The silicon detectors will provide information on the charged particle multiplicity for triggering via a sum of channels above threshold. The next version of the preamplifier (the true prototype) will provide a signal proportional to this. A current sum of the discriminator outputs from the eight preamplifier chips on each hybrid will be driven off the hybrid to the level-1 (LVL-1) trigger electronics for further manipulation. It has been shown that channels on a single silicon die in an ASIC are well matched in gain [3].

The analog memory unit (AMU) is a switched capacitor array with simultaneous read/write capability to pipeline the data pending a low level trigger decision. The MVD requires eight bits of resolution in the ADC to enable fitting the pulse height spectrum of the silicon pads. The MVD will use a variant of the AMU/ADC circuit under development for all of the PHENIX subsystems, modified for lower resolution and smaller power dissipation. Each channel has its own ADC circuit and a FIFO of at least two registers to store the data for readout. When an event is triggered, the information from the preceding beam crossing (which is very unlikely to contain a high multiplicity collision or particle hitting the same pixel) is used to set a reference level. The signal from the triggered event is compared to this reference to remove baseline shift and then digitized.

Control of the AMU, ADC and FIFO/readout is provided by a controller referred to in PHENIX as a "heap manager". The heap manager for the MVD will be a variant of the one used by all PHENIX subsystems. Currently, it is foreseen that this function can be provided by a field programmable gate array. However, the MVD must use the heap manager as a bare die, whereas other subsystems can use the commonly available packaged modules. There will be one heap manager per hybrid to provide all necessary controls.

Figure 5.7: Layout of the hybrid circuit. The various components are labeled.

5.4.2 Hybrid circuit

Each detector wafer has 256 channels and is connected via a kapton cable glued and bonded to one hybrid circuit. This hybrid circuit has been simplified since the writing of the CDR by removing the zero-suppression and data packing functions. These will now be done on boards outside of the target region. All the information from a triggered event will be transferred off the hybrid via an optical link. A updated layout of the hybrid circuit is shown in Fig. 5.7. The signals enter the preamplifier/shaper/summer and then the AMU/ADC which both have 32 channels per die. Each hybrid also drives the current sum of the discriminator signals from the eight preamp/shaper chips to the LVL-1 trigger electronics. The 32 channels on each individual chip are summed on the chip. We expect that the driver will consist of surface mount components on the hybrid circuit and will not require an ASIC. The current hybrid circuit is 48-mm wide and 20-mm deep, fitting within the nominal mechanical constraints of PHENIX and other RHIC experiments as well.

The hybrid circuits will probably be fabricated using a process from Texas Instruments called High Density Interconnect (HDI). This process appears to be the best choice for the trace densities required in our circuit. Alumina substrates are used, with milled wells for the silicon dies. A polyimide coating is laser drilled at each die contact and then metallized to provide electrical contact and interconnect traces. This is a proven process that allows pitches of less than 100 μm . The high interconnect density is driven by the bus providing control signals (16×2 signals) from the heap manager to the AMU/ADC chips. The HDI process also offers the advantage of excellent thermal contact between the substrate and the

Table 5.2: MVD Material Thickness †

Source	$\langle X_0 \rangle$ (PISA)	$\langle X_0 \rangle$ (CDR Update)
Dalitz	0.61%‡	0.61%‡
Beam Pipe	0.28%	0.28%
MVD Layer 1	0.32%	0.32%
MVD Layer 2	0.32%	0.32%
MVD Electronics	0.52%	0.52%
Rohacell Support	0.003%	0.01%
Enclosure	—	0.08%
Kapton Cables	—	0.04%
Total	2.06%	2.18%
MVD/(Dalitz + Beam Pipe)	1.31	1.45
MVD/(Dalitz + Beam Pipe) - without electronics	0.73	0.87

† This table includes a full outer barrel of silicon detectors.

‡ Dalitz decay of π^0 expressed as an equivalent in radiation lengths contributing to background.

chips, to simplify the heat removal.

5.5 Silicon System Performance Issues

We have made detailed studies of the effect of the MVD on the rest of the PHENIX experiment using the PISA program as described in Chapter 3. The number of background electrons arising from the MVD and the effect of those on the electron trigger and offline electron analysis was discussed there. In this section we will describe studies specific to the layout and design of the MVD. Table 5.2 shows the radiation length of material in the current MVD design, and compares to what was included in the PISA calculations shown. About 10% of the material was not accounted for in the simulations, but this omission should not alter any of the conclusions.

5.5.1 Electronics Location

PISA was used to investigate the difference between the current MVD with the barrel electronics on the bottom and the CDR design with the electronics oriented radially, in three sections separated by 120° in azimuth. The new arrangement both reduces the background in the RICH and simplifies the cooling of the electronics. Table 5.3 Comparison of $e^\pm$ background from the barrel electronics in the CDR configuration and with all of the electronics on

Table 5.3: Comparison of the $e^\pm$ Background from the Barrel Electronics

Electronics Configuration	N/event (no cuts)	N/event ($p_T > 200$ MeV/c)
CDR	0.78 ± 0.12	0.38 ± 0.09
Bottom Only	0.63 ± 0.08	0.00 ± 0.00

the bottom of the detector; both calculations include a full outer silicon layer. Backgrounds are calculated using PISA, based on photons from π^0 decay in central Hijet events.

The calculations summarized in Table 5.3 assume a complete barrel detector without a hole in the upper two-thirds of the outer layer. With the CDR electronics configuration 0.78 ± 0.12 electrons originating in the electronics enter the RICH per event. This number falls slightly to 0.63 ± 0.08 per event with the electronics on the bottom. For electrons with $p_T > 0.2$ GeV/c, the difference is significant. In 100 simulated events, no electrons reach the RICH from the electronics when they are on the bottom. With the electronics configured as in the CDR 0.38 ± 0.09 electrons per event with $p_T > 0.2$ GeV/c reach the RICH. It is easy to understand why the p_T cutoff is so successful in eliminating background when the electronics are on the bottom. When a photon converts to an $e^\pm$ pair, the produced particles go approximately in the direction of the original photon, in this case, downward, away from the detector arms. Electrons with momenta significantly greater than the 100 MeV/c p_T kick from the central magnet cannot be bent the approximately 60° or more that is required for them to enter the RICH.

5.5.2 Silicon Coverage

The issues influencing the extent of silicon coverage in the MVD are electron background from photon conversions in the MVD material, and the competing requirements for good efficiency to find the vertex in low multiplicity events. It is clear that minimizing the background implies minimum material in the MVD, and maximizing vertex finding efficiency requires maximum coverage with two layers of silicon. The rapid changeover time between A + A and p + A running prohibits the solution proposed in the CDR of modifying the number of barrel azimuthal sectors in the MVD. We have used the PISA calculations to find an optimal solution between these two extremes.

To do this, we used the PISA simulation to study the effect of omitting silicon strip detector wafers from the outer barrel layer, along with the associated cabling and electronics. Making a hole in only one of the two layers preserves the possibility of triggering on multiplicity in the barrel in an unbiased way, and the ability to search for multiplicity fluctuations at midrapidity. Furthermore, this approach allows the remaining outer silicon to be useful for vertex finding. The vertex algorithm uses tracks which hit both silicon layers, but the overlap in η becomes small as the two layers become short in z , causing the number of useful tracks to fall drastically. However, leaving one layer complete in z means that whatever

Figure 5.8: The electron background from MVD in central Au + Au collisions and the probability of finding the vertex in p + Au and p + p collisions, as a function of the amount of silicon removed from the top two-thirds of the outer barrel layer.

remains of the other layer contributes to the vertexing capability. This compromise avoids relying on the lower one-third of the barrel alone, which has been shown to yield only 50% vertex reconstruction efficiency for p + p collisions. One of the major issues addressed is whether the electron background is sufficiently low with this approach.

Figure 5.8 shows the results of this study. The electron background and vertex finding

efficiency are plotted as a function of the size of the hole along the z direction created by omitting elements of the top two-thirds of the barrel outer layer. The top left shows the number of conversion electrons per event from the MVD which enter the electron arms. The filled circles show the total number of electrons, while the open circles give the number of electrons with $p_t > 0.2$ GeV/c. This p_T cut is roughly equivalent to the threshold imposed by the angular acceptance of the Winston cones in the RICH. Below this threshold, the electrons will not be identified by the RICH. It is clear that the majority of the MVD conversion electrons in the electron arm acceptance are very soft and are below this p_T threshold. The top right gives the fraction of MVD conversion electrons out of the total background electrons entering the RICH. This illustrates the magnitude of the MVD contribution to electrons at the trigger level. This fraction is quite small, independent of the p_T threshold. The bottom left shows the ratio of the MVD conversion electrons to those arising from Dalitz decays and conversions in the beam pipe (the irreducible electron background in the baseline PHENIX detector). This plot shows the increase in the background from the MVD expected in the offline analysis, and allows evaluation of the increase in the combinatorial background caused by material in the MVD. The right bottom shows the probability of finding the vertex position in $p + \text{Au}$ and $p + p$ events, also as a function of the hole size. It should be noted that this is for events which have at least two charged particles produced in the η range of the barrel.

By comparing the bottom two sections of Fig. 5.8 the tradeoff between background and vertex efficiency can be quantified. The background above p_T of 0.2 GeV/c is the relevant quantity for this comparison. The left side shows that with the entire outer silicon layer, the MVD increases the single electron background by one-third, and the dielectron combinatorial background by a factor of 1.75. An MVD with only one silicon layer in the top two-thirds increases the single electron background by about 20% and the combinatorial background by 40%. The collaboration agrees that we do not want to increase the combinatorial background by nearly two, but a 40% increase is acceptable. This implies a hole size of 40 cm or more in z . Inspection of the right bottom plot shows that the probability of finding the vertex is only 50% with no outer layer in the top two-thirds of the MVD, but improves to 65% with a 40-cm hole in that layer. Though this is not ideal for analysis of $p + p$ collisions, the collaboration has decided to accept this as truly low multiplicity events have a high cross section and the rate loss should not preclude collection of sufficient statistics. The measurement of high mass dileptons, which have lower production rates, will be hurt by this efficiency loss if the dileptons arise from low multiplicity events or are accompanied by jets which can miss the two layer section of the vertex detector.

5.6 Cost and Schedule

The cost estimate and schedule for the MVD is presented in Table 5.4. In this table the MVD mechanical (WBS 5.2.1) and front-end electronics (WBS 5.3.1) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

Table 5.4: Cost Estimate and Schedule for the PHENIX MVD

WBS	Description	Total Cost	1992			1993				1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	
5.2.1	MVD	\$1,262,296.00																											
5.2.1.1	Engineering & Design	\$291,024.00																											
5.2.1.2	MVD Prototypes	\$233,476.00																											
5.2.1.3	MVD Fabrication	\$647,478.00																											
5.2.1.4	MVD Assembly	\$42,640.00																											
5.2.1.5	MVD Testing	\$47,680.00																											

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998					
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4		
5.3.1	MVD FEE	\$697,926.00																								
5.3.1.1	Engineering & Design	\$229,216.00																								
5.3.1.2	MVD FEE Prototypes	\$157,310.00																								
5.3.1.3	MVD Electronic Fabrication	\$311,400.00																								

5.7 Beam-beam Counter Overview

The PHENIX beam-beam counter (BB) provides information for triggering on beam-beam collisions, the rough collision vertex position along the beam axis (z) that helps to find correlation between tracks in various detector elements, and the time origin of the collisions. The information must be available at the trigger level to discriminate backgrounds, and the timing information is also used to perform the time-of-flight measurement for particle identification combined with the information from the TOF detectors.

The BB consists of two arrays of counter elements, one each placed upstream and downstream of the beam crossing point. Each array consists of 66 modules of counter elements surrounding the beam pipe, and covers the forward and the backward rapidity regions. The average and the difference of the arrival times of fast leading particles from beam-beam collisions to the two BB counter arrays provide the time origin and the vertex position of the collisions, respectively.

Each counter element is constructed from a Cherenkov radiator of fused quartz and a 1-in. mesh-dynode photomultiplier made by Hamamatsu (Hamamatsu H5082). The fused quartz radiator is glued to the photomultiplier window. With such a combination of radiator and the photomultiplier, an intrinsic timing resolution of 44 ps was achieved for 1.6 GeV/c

pions in a beam test at KEK. Good operation of the BB is expected in the magnetic field environment of the BB location due to the choice of photomultiplier.

The BB electronics chain consists of discriminators, shaping amplifiers, time-to-voltage converters (TVC), flash ADCs (FADC) and buffer memories. The timing and pulse height information of the BB elements are digitized in every beam crossing by the TVC and FADC. The digitized data are provided as a set of input of the LVL-1 trigger, and are stored at the same time in the buffer memory for the event building. The discriminators are designed so that the walk-free timing information from the BB is available in the LVL-1 trigger.

5.8 Beam-beam Counter Changes

Since the CDR, there has been considerable progress in the BB design. Based largely on the results of a beam test of counter elements, there are several changes and refinements, which include:

1. The selection of a hexagonal face, rather than circular face, as the shape for the front of the BB counter elements.
2. The change in design of the counter elements of the BB to minimize radiation damage.
3. The development of a new method for maintenance access to the BB detector on the muon arm side to reduce interference with access to the muon arm and central magnet.
4. The investigation and testing of a candidate for the BB discriminator, based on the double threshold method.

Performances of Cherenkov radiators with a hexagonal and a round cross-section for the shape of the BB counter elements were compared in a beam test at KEK. Round Cherenkov radiators with different lengths and diameters were also tested to optimize the radiator size. The shape (hexagonal) and size ($29 \times 30 \text{ mm}^2$) of the Cherenkov radiator for the BB counter elements were finalized by the beam test.

The BB location is a high radiation environment in the PHENIX detector. Though the estimation of the radiation level at the BB location is still under way, it is obvious that the BB detector requires good radiation hardness. The radiation hardness of photomultipliers and Cherenkov radiator materials against neutrons and γ -rays was considered, and the design of the counter elements of the BB was changed to minimize the damage of the radiation. Furthermore, the BB location is also a high counting rate environment. Together with the radiation damage, a deterioration of photomultiplier tubes by an over current in the dynode chain, caused by the high rate counting, possibly contributes to the damage of the BB counter elements. An averaged current in the dynode chain was estimated and was compared with the critical current of the photomultiplier tubes.

The BB counter array at the muon spectrometer side will be installed into a space between the central magnet and the muon tracking chambers. The maintenance access to the BB detector is not straightforward. The access inevitably interferes with the central magnet and

the muon chambers. The method of the maintenance access to the BB detector at the muon spectrometer side was changed to reduce the interference.

An important requirement of the BB counter is to provide timing information of good quality to the LVL-1 trigger. For the requirement, the discriminator of the BB counter must provide walk-free timing information in the high multiplicity environment. A candidate of the discriminator based on the double threshold method was tested.

5.9 Mechanical

5.9.1 Detector Elements

Each element of the BB counter is a flashlight-type Cherenkov counter, with a UV-transparent solid Cherenkov radiator mounted directly on a mesh-dynode photomultiplier tube.

There were two options for the shape of the front face of the Cherenkov radiator: a hexagonal or a circular cross-section. Performance of the two Cherenkov radiator shapes was studied in a beam test at KEK. In the test, radiators of both shapes had comparable intrinsic timing resolutions. Consequently, the choice of shape is not determined by the intrinsic timing resolution. The hexagonal shape was chosen for PHENIX to minimize insensitive gaps between counter elements. The outer diameter of the hexagonal cylinder is 29 mm, which matches the 1-in. photomultiplier tube, and the length is 30 mm, which is the minimum length giving enough Cherenkov light to achieve a 50 ps intrinsic timing resolution.

It is known that the major radiation damage of photomultiplier tubes by neutrons and γ -rays is deterioration in the optical transparency of the window material, especially at short wave lengths. Such deterioration at short wave lengths is, of course, critical for Cherenkov counters. Deterioration in the transparency is lower for fused quartz than for normal and UV glasses. The fused quartz is accordingly a good choice for the Cherenkov radiator and window material. In the previous design of the BB counter elements, the radiator and photomultiplier were glued together, however glues have notoriously poor radiation hardness qualities. Therefore, a mesh-dynode photomultiplier with a thick fused quartz window, which is the Cherenkov radiator at the same time, of a hexagonal cylinder shape is newly designed by Hamamatsu, as shown in Fig.5.9. A 1-in. mesh-dynode photomultiplier based on the Hamamatsu H5082 is used. Similar to the Hamamatsu H5082, this new mesh-dynode photomultiplier is expected to work without a magnetic shield in the magnetic field environment at the BB location (see Fig.5.10).

Another factor of deterioration in the counter performance is an over current in the dynode chain of the photomultiplier tube. A damage by the over current lowers the current gain of the photomultiplier. In a high rate and high multiplicity environment, the deterioration determines the life time of the photomultiplier and the cycle of maintenance access needed to keep the BB performance. The deterioration comes from a thermal damage of the dynodes. The critical anode current is given as an averaged current over a few seconds, and its catalog value is 10 μ A for the Hamamatsu H5082 photomultiplier. An estimated anode current I_{anode} of the photomultiplier tube in the Au + Au collisions at a luminosity of $2.2 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$

Figure 5.9: A 1-in. mesh-dynode photomultiplier newly designed by Hamamatsu based on the Hamamatsu H5082. The photomultiplier has a fused quartz window or radiator of hexagonal cylinder shape, a 29-mm outer diameter and a 30-mm thickness.

is:

$$I_{\text{anode}} = 1 \times 10^{-10} \text{C/mip} \times 2 \text{mips/event} \times 14 \times 10^3 \text{events/s} = 2.8 \mu\text{A}$$

In performing this calculation, we assumed that the integrated charge of the photomultiplier tube signal by a minimum ionizing particle (mip) is $1 \times 10^{-10} \text{C/mip}$, that the average number of mips in a counter module per event is 2 mips/event, and that the event rate is $14 \times 10^3 \text{events/s}$. The resulting average anode current I_{anode} is below the critical value. At much higher luminosity, I_{anode} is likely comparable with the critical current. A reduction of I_{anode} by a factor of two is possible by reducing current gain of the photomultiplier tube without significant growth in the timing resolution. The study on effects of the over current in a realistic condition is still under way.

5.9.2 BB Location

The BB detector consists of two arrays of the counter elements. Each BB array has 66 modules of the counter element. One BB array at the muon spectrometer side is located in the copper plug attached to the central magnet. Another BB array of identical assembly is mounted on a support stand at the other side.

In the previous design, the copper plug at the muon spectrometer side consisted of three parts, one cylindrical part covering the beam pipe and two tapered parts covering the BB. The copper plug was attached to the iron piston of the muon magnet. The maintenance access to the BB was possible only after removing several muon tracking chambers and the tapered part of the copper plug.

In the present design, the copper plug has a simple tapered shape and a maintenance access to the BB is possible from the rear side (see Fig.5.11). The access can be made by

Figure 5.10: The dependence of the signal pulse height (filled circle) and the intrinsic timing resolution (open circle) on an axial magnetic field strength is shown for a combination of a hexagonal cylinder shape radiator and the Hamamatsu H5082 photomultiplier tube. The magnetic field strength at the BB location is expected to be less than 3 kG.

moving the central magnet away from the muon spectrometer. The BB, the copper plug, and the γ -ray absorber for the muon spectrometer are moved with the central magnet. The BB array is accessed from the rear side by a person standing at a 150-cm gap between the γ -ray absorber and the muon tracking chamber station #1, after the CM is moved away.

In accordance with the change of the copper plug design, cabling and cooling of the BB array at the muon spectrometer side was changed. Sixty-six coaxial signal cables, 11 high voltage cables, and 66 optical fiber cables, for counter calibrations, are laid through a 5-cm gap between the γ -ray absorber and the muon tracking chamber. The number of high voltage cables was increased from nine to 11 to keep the total current per high voltage cable less than 1 mA. The main heat source of the counter array is the bleeder circuit of the photomultiplier tubes. Total power consumption of the photomultiplier tubes in each BB array is $0.37 \times 66 = 24.4\text{W}$. The heat production of the BB is small enough to be cooled by fans. Furthermore, the circulation of air between the inside of the copper plug and the atmosphere is improved.

5.9.3 Electronics

In central Au + Au collisions, each counter element of BB is likely to be hit by multiple particles. The timing signal arises from the first particle entering the element, but the measured pulse height includes all particles, so traditional slewing correction is not applicable

Figure 5.11: One BB counter array in the muon spectrometer side is installed to the inside of the copper plug. Another array is located at the other side of the central magnet.

at all. We will address this problem by using either a discriminator based on the double threshold method or a constant fraction discriminator (CFD).

The double threshold method uses two leading-edge discriminators with different thresholds. Two timing information T_1 and T_2 from the two discriminators are combined to get a walk-free timing information T_0 as follows.

$$T_0 = 2 \times T_1 - T_2$$

The two timings T_1 and T_2 are determined by the pulse shape at the leading-edge, at which the contribution of the first particle is dominant, and no pulse height information is needed to obtain the walk-free timing T_0 . A merit of the double threshold method is that the method is based on the well established leading-edge discriminators. Another merit is that the on-line correction of the timing by the double threshold method is simple and is possible with no help of CPU. The feasibility of the double threshold method with ordinary discriminators was tested with 1.6 GeV/c pion beams at KEK, and we achieved 47 ps intrinsic timing resolution, which was comparable with a timing resolution by a leading-edge discriminator with applying the slewing correction.

The CFD might also provide walk-free timing information for the multiple particles. The

development and testing of the CFD for PHENIX is still under way at LANL, and, at present, the final choice of the discriminator scheme is still open.

5.10 Beam-beam Counter Performance Issues

As indicated previously in Chapter 3, the beam-beam counter determines the time zero (T_0) when an interaction occurs in a beam bunch crossing (2 ns traversal time) and provides the first estimate of the vertex point (z_0) along the beam axis. The BB consists of two sets of identical counter arrays installed on both sides of the interaction region. One set is in the muon plug on the muon detector side, and the other is in free space just behind the central magnet yoke. These counter arrays intercept particles emitted at very forward angles in the rapidity range from three to four at $z = \pm 150$ cm. Charged particles with a velocity faster than $\beta = (1/1.458)^{1/2} = 0.69$ will be detected by the counters. By measuring a pair of arrival times T_1 and T_2 of the fastest particle that intercepts the counter array on each side, T_0 and z_0 will be obtained as the average and the difference of T_1 and T_2 , respectively.

If there were no backgrounds nor any slow particles, the resolutions of the extracted T_0 and z_0 would be affected only by the intrinsic time resolution of the counter elements. However, all interactions do produce slow particles, even at very forward angles, and there does exist a significant amount of background particles scattered from the central spectrometer structures and from the Muon detector. Therefore an examination of the BB performance with a Monte Carlo simulation is mandated.

The first results of the effects of the QED pair background have already been discussed in Chapter 3. Here we focus on evaluating the intrinsic performance of the BB counter arrays, in particular the simulation of the T_0 measurement. Those simulations have been carried out with an input of the geometry of the BB, as described in the PHENIX baseline design, and taking the intrinsic time resolution of each counter element to be 50 ps, and then using this information in the PISA/PISORP program chain.

The spectrometer components along the beam axis (i.e., the central magnet, beam pipe, nose cone, muon plug, and the muon piston) as well as the BB were installed in the PISA simulation to study the performance of the T_0 measurement under the realistic condition of background particles scattered from those materials in a magnetic field. Subsets of the PISA UA1-type events were thrown from the $z = 0$ position into a forward cone of polar angle 10° . The BB acceptance covers from $2^\circ - 6^\circ$. The mean multiplicity of particles thrown within the cone was 100, and the BB detected about ten charged particles in average per event.

The PISA simulation returns several kinds of variables relating to the hit time, position, and direction of each charged particle which entered the sensitive volume of a given detector element. The simulation program also provides the amount of energy loss and the path length for each particle whether it has stopped in the detector element or entered and exited. The event contains not only leading particles from the vertex but also shower particles created by the leading particles both inside and outside the sensitive volume. All of this information is retained in the PISA output file for subsequent diagnostic analysis.

¹The refractive index of quartz, which is the sensitive material of the counter elements.

Figure 5.12: Simulations of ADC distributions of the BB detector system for 1000 UA1 type events. The left figure is for the detector elements on the muon arm side, and the right figure is for the detector elements on the free space side.

This output file is “played back” by the PISORP program in order to obtain the arrival hit times of all the particles in each event. The earliest time in each detector was defined to be the TDC data of the detector. The energy loss itself does not directly afford any signal information by the detector because it is a Cherenkov device. Instead, and in a manner somewhat akin to the simulation of the RICH counter, the ADC data for each BB detector element was taken to be the sum of the track lengths of all particles which have incident velocities above the Cherenkov threshold and whose directions are within a specified acceptance angle of 45° . (The RICH simulation goes one step further in that for tracking and timing purposes it generates Cherenkov photons and follows through their trajectories.) The ADC spectrum in Fig. 5.12 has a sharp peak at 3 cm, since the current design has a sensitive volume 3-cm thick and no smearing effects were included for the ADC response. The peak has a small tail on the high side due to secondary particle production inside the sensitive volume. Double hits in a single detector element result in a peak at 6 cm. The ADC simulation data show a large and continuous component over the full range. The region below the first peak represents events for which there were no leading particle hits but there was only a soft background of particles scattered from material outside the detector. The left figure is of the detector residing in the muon plug on the muon arm side, while the right figure is of the detector situated in free space. No noticeable yield differences in the backgrounds are apparent.

The two left figures in Fig. 5.13 show distributions of TDC data of all tracks recorded by the counter elements for 1000 events. The top figures are of the detector in the muon plug on the muon detector side, and the bottom are of the detector in the free space. The width of the both distributions is around 60 ps, and is wider than the intrinsic time resolution of the counter elements. This broadening occurs because the distribution contains TDC data of

Figure 5.13: TDC distributions (left) of all tracks recorded by counters for about 1000 UA1 sub-events and the T_0 distributions (right) determined for each event. Top figures are of the BB in the muon plug on the muon detector side, and the bottom figures are of the BB in the free space on the opposite side. No significant differences are seen, although the two background environments are quite different.

Bibliography

- [1] J.P. Sullivan, et. al, Proceedings of the Symposium on RHIC Detector R&D, BNL-52321, p.107 (1992).
- [2] R. Bellwied, private communication.
- [3] C.L. Britton, Jr., G.T. Alley, M.L. Simpson, A.L. Wintenberg, R.J. Yarema, T. Zimmerman, J. Boissevain, J. Simon-Gillo, J.P. Sullivan, B.V. Jacak, and N. Lockyer, submitted to IEEE Nuclear Science Transactions.

Chapter 6

Tracking System

6.1 Overview

The PHENIX tracking system measures the p_T and p_{total} of charged particles from RHIC collisions, reconstructs invariant masses from electron and kaon pairs, contributes to the particle identification, and provides position information to the LVL-2 trigger. Three tracking subsystems are optimized for different functions. Low mass, multiwire focussing drift chambers (DC) provide high resolution p_T measurements. Three interpolating pad chambers (PC1, PC2, PC3) provide three dimensional space points for the charged tracks, determine p_z/p_T , and provide space points to the LVL-2 trigger. A time expansion chamber (TEC) tracks charged particles between the RICH and the EMCal, identifies tracks originating from outside the fiducial volume, and provides e/π separation in the momentum range of 250 MeV/c to 2.5 GeV/c using dE/dx information. All three subsystems work together to perform the tracking pattern recognition in the high multiplicity RHIC environment.

The design of the PHENIX tracking system is driven by unique features of both heavy ion collisions and the PHENIX experiment. Central Au + Au collisions at RHIC produce charged particle multiplicities of $dN_c/dy = 1500$ with a spectrum of soft particles peaking at 200–300 MeV/c. Low momentum complications such as multiple scattering and particle decay make tracking difficult. In addition, the PHENIX axial magnetic field includes a focussing component that creates complex track paths through the detector for low p_T particles. The split in central-arm azimuthal coverage results in a large number of particles entering the detector's active region without passing through the drift chambers. This effect makes a compact tracking system impractical and has helped to motivate the distribution of tracking detectors.

6.2 Changes

The changes in the tracking subsystem since the CDR are:

1. PC2, TEC, and PC3 are now four segments per arm instead of two.

2. Chevron pad is the PC technology choice for the PHENIX baseline detector. The Pixel detector R&D program will continue at Lund.
3. New prototype results are being used to determine tracking detector parameters.
4. A modified TMC chip developed at KEK for SDC will be used for the drift chamber TDC.
5. A new track reconstruction, pattern recognition, and detector simulation code is now being used.
6. There is progress on detailed mechanical design of all tracking systems.

6.3 Drift Chambers

There is considerable progress on the design of the DC subsystem, which is based, in part, on results from several prototype drift cells fabricated and tested at PNPI. A PNPI variant of the original CDR drift cell configuration gives the best overall results and confirms that the drift chambers will provide excellent single and two-track resolutions. In addition, the engineering effort at PNPI is augmented by a local industrial firm, Azimut. Much of the current mechanical design is based on suggestions from Azimut engineers. Finally, the primary components of the front-end electronics are identified, including a new pipeline TDC chip designed specifically for use with the proposed ASD-8 amplifier/shaper/discriminator. A preamplifier/shaper hybrid circuit development is also underway at PNPI.

6.3.1 Mechanical Design

Each arm of the drift chamber subsystem is a single arch-shaped frame supporting individual X, U and V modules. This design minimizes the mass of support material in the fiducial volume and facilitates design, fabrication, transportation and assembly.

There are 2560 sense wires in each of the two central arms. Each wire is electrically split at the center ($z_{\text{beam}} = 0$ cm) and both ends are read out. The total number of electronics channels is 10,240 (2560×2 arms $\times 2$ ends). Within each arm, there are four layers of drift cells. The two layers closest to the beam are X cells with 12 sense wires per cell (6 left and 6 right wires); these two layers are separated radially by approximately 20 cm. The final two layers of U and V cells, each with 4 sense wires, are positioned radially between the second X layer and the output windows.

Each layer is divided into 8 separate modules with each module subtending 11.25° in ϕ . This gives 320 sense wires in the radially overlapping X, X, U and V modules. Within each module there are 10 drift cells independent of layer. This means that the drift distance is not identical in each layer (since the cells all subtend the same ϕ) but varies radially outward from 2.0 to 2.4 cm.

There are 32 sense wires in the four overlapping X, X, U and V cells: $12 \text{ X} + 12 \text{ X} + 4 \text{ U} + 4 \text{ V}$. Therefore the total number of channels can be computed as $32 \text{ wires/cell} \times 2 \text{ channels/wire} \times 10 \text{ cells/module} \times 8 \text{ modules/arm} \times 2 \text{ arms}$ equals 10,240 channels.

Figure 6.1: Schematic of the ceramic wire support plates showing the metallic layers, wire bonding locations, and the location of one of the reference roads.

There are 64 separate modules: 32 identical X modules, 16 identical U modules and 16 identical V modules. The final module and arch framework designs are underway. Among the remaining issues to be resolved are whether the wire tension should be supported within the modules themselves or from the arch framework. A self-supporting module would necessarily have some mass in the fiducial volume; however, it is unclear at this point whether a single arch frame can be designed to withstand the total wire tension without some cross support.

There is considerable progress on the individual cell and module designs. The cell wire configuration was selected based on further electrostatics calculations and preliminary prototype results. The configuration is very close to the original CDR design. The guard and channel wires lie in two planes offset 3.0 mm on either side of the sense wire planes (the guard wires were originally closer to the sense wires). This was done to simplify the wire support mechanism and to facilitate the ease of construction.

Particular attention is focussed on attaching and locating the wires precisely. The plan is to fix the wires onto ceramic plates (Al_2O_3), which can be fabricated in Russia with extremely high precision. Two reference roads on the plates are positioned with an accuracy of a few microns. Metallic layers are evaporated onto the plates to form the conducting surfaces on which the wires are attached and the signals are removed. The wires are bonded to the metallic surface with three separate procedures, which in sequence are: (1) laser soldering at several points, (2) standard soldering, and (3) high resistive epoxy. Figure 6.1 shows a schematic of two ceramic plates with the locations of one reference road and several of the metallic surfaces to which wires are attached.

Figure 6.2 shows schematically the location of the ceramic plates within each module. Special cages designed to maintain precise wire position are used to transfer the plates to the module frame once the wires are attached. It is important to maintain the proper tension throughout the fabrication process.

Figure 6.2: Schematic of a proposed module assembly showing DC end plate, module support and alignment.

To electrically isolate the two halves of a single sense wire, without adding additional tension support at the center, we use a 10 mg “glass” bead, which has a hole for the sense wire. A laser beam is used to cut the wire. The wire can be attached to the bead by welding, using epoxy, or both. Tests will determine which bonding method is most suitable. Since the sense wires are 25- μm gold-plated tungsten with 70 g of tension, the mass of the glass bead is chosen to provide just enough additional gravitational sag to match the 200 g of tension of the 90- μm copper-beryllium field, drift, guard, and channel wires.

Additional calculations show that there should be no electrostatic wire staggering at the nominal tensions, and that the sag of the sense wires caused by electrostatic repulsion is of order 200 μm . The electrostatic sag causes some loss in gain (estimated at 18%), which is not crucial, because high gains were achieved in the prototype cells (see Section 6.3.3).

6.3.2 Electronics

A total of 320 motherboards with 32 readout channels per board comprise the drift chamber electronics. The primary components of the front-end electronics readout chain comprise a preamplifier/shaper/discriminator circuit mated with a pipeline TDC and at least 4 μs of onboard data buffering. Initial studies indicate that the ASD-8 chip described in the CDR

would be a good choice for the preamplifier/shaper/discriminator. Only minor modifications are required to match the input impedance, optimize tail cancellation and possibly increase the shaping time. It is expected that these modifications can be made by the original designers at a reasonable cost. A second viable option for the preamplifier/shaper/discriminator circuit is a hybrid chip developed by the PNPI electronics group. This hybrid is based on the circuit currently under development for the full scale driftchamber prototype testing. The hybrid option may offer a significant cost savings by eliminating the reengineering and setup costs for producing the ASD-8.

A new TDC chip is now identified. This CMOS chip, the TMC4004 [1], was designed at KEK by Y. Arai to operate with the ASD-8 for the SDC straw-tube-tracker subsystem. The TMC is a time memory cell with 1-ns resolution, no dead time, excellent temperature stability and linearity and power dissipation of only about 7 mW/channel. Minor modifications are required for the TMC to function as the pipeline TDC for the drift chamber FEE. In particular, we need to record the leading and trailing edges of all pulses and to slightly improve the two pulse resolution. However, it may be possible to put all of the output buffering onto the TMC chip itself. This mitigates the need for a separate DMU chip. Preliminary discussions with Arai indicate that the required modifications to reengineer the TMC can be made at KEK.

A leading solution for the drift chamber FEE is the ASD-8 amplifier/shaper/discriminator chip and TMC TDC, because these now form a working circuit for the SDC detector. However, the hybrid option may be selected after further effort to match the output to the TMC input. This work should begin in FY 1994.

6.3.3 Prototype Operation and Performance

Prototype measurements in the lab using both a beta source and cosmic rays, as well as in test beam at the PNPI synchrocyclotron (1 GeV/c protons), demonstrate that the drift chamber cell described above meets all of the performance criteria listed in the CDR. In addition, these measurements show that the electrostatic behavior of the cell is well described by GARFIELD calculations including the predicted gas gain. Gas gains up to 50,000 were measured with the nominal 50/50 argon-ethane gas mixture. The single-wire efficiencies for one of the prototype chambers are listed in Table 6.1 for the voltage configuration $V_d = -4.5$, $V_f = -2.65$, $V_c = 1.6$, and $V_g = 0.85$ kV. The even numbered wires are protected by guard wires where little or no signal is expected, while the open odd wires should show an efficiency in excess of 98%. With a 10-mV threshold, this is indeed the case; however, the guarded sense wires also show some small efficiency. By optimizing the operating voltage configuration and threshold, we expect to further minimize the guarded-wire efficiencies without compromising the excellent efficiency of the open wires.

The single-wire-position resolution was found to be approximately $100 \pm 20 \mu\text{mm}$ using a beta source and only slightly higher using cosmic rays. The two-track resolution has been estimated by combining the signals from two sense wires and then inclining the particle tracks (using both the source and cosmic rays) until the two signals can be resolved. Based on these measurements, the two-track resolution is estimated to be approximately $1300 \mu\text{m}$, which is well within the desired resolution.

Table 6.1: Prototype Single-wire Efficiencies in Percent for Two Different Threshold Levels

Wire	10 mV	20 mV	Wire	10 mV	20 mV
1	98.3±0.4	94.7±0.9	2	1.1±0.4	0.4±0.3
3	98.4±0.3	95.4±0.8	4	12.6±0.0	0.7±0.4
5	98.8±0.3	93.2±0.0	6	15.0±0.1	1.4±0.5
7	98.7±0.3	94.2±0.9	8	2.5±0.5	0.7±0.4
9	99.0±0.3	94.6±0.9	10	5.5±0.7	0.6±0.3
11	98.7±0.3	94.0±0.9	12	6.6±0.8	0.6±0.3

6.3.4 Cost and Schedule

The cost estimate and schedule for the MVD is presented in Table 6.2. In this table the MVD mechanical (WBS 5.2.4) and front-end electronics (WBS 5.3.4) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book. The cost of the drift chamber mechanical and FEE have not changed significantly since the CDR was submitted. The schedule accommodates the commissioning of the PHENIX detector, and no critical schedule conflicts are anticipated for the completion of the drift chamber mechanical and FEE construction.

Table 6.2: Cost Estimate and Schedule for the PHENIX Drift Chambers

WBS	Description	Total Cost	1994				1995				1996				1997				1998				1999		
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3
5.2.4	DC	\$1,258,422.64																							
5.2.4.1	DC design	\$164,280.00																							
5.2.4.2	Prototype Fabrication and Testing	\$111,368.40																							
5.2.4.3	DC Fabrication	\$808,264.64																							
5.2.4.4	Drift Chamber Assembly and Testing	\$174,489.60																							

WBS	Description	Total Cost	1994				1995				1996				1997				1998				19	
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2
5.3.4	DC FEE	\$922,901.12																						
5.3.4.1	DC FEE design	\$301,544.00																						
5.3.4.2	DC FEE prototyping	\$63,516.80																						
5.3.4.3	Procurement of DC FEE	\$447,040.32																						
5.3.4.4	DC FEE test	\$34,400.00																						
5.3.4.5	Assemble DE FEE	\$56,400.00																						

6.4 Pad Chambers

During the past few months, considerable attention was devoted to the question of chevron-pad versus pixel-pad cathode readout scheme for the pad chambers. After careful deliberation, the chevron interpolating-pad scheme was adopted. Although the pixel-pad scheme offers potential advantages in both occupancy and double-track resolution, further R&D work is required. The University of Lund group is encouraged to continue to develop the pixel-pad system as a possible future upgrade to the pad chambers. There are a number of significant new developments in the conceptual design of all the pad chambers (PC1, PC2, and PC3), both in mechanical structure and electronic readout scheme.

6.4.1 Mechanical Design

The PC1, which occupies the radial space between inscribed radii 247.50 and 252.50 cm, was changed from a nine-sector to an eight-sector per detector arm design, as illustrated in Fig. 6.3. The eight sectors are completely modularized and attached to each other by detachable mechanical locks. Adjacent sectors are separated by a 25- μ m aluminized mylar gas seal on each side, and along the seam between two wire planes there is a one-millimeter dead space. The cathode panels, which have areal dimensions $490.00 \times 1887.02 \text{ mm}^2$, are to be fabricated with rohacell sandwiched between 250- μ m FR4 skins. The chevron cathode panel has 25.40 mm of rohacell, and the panel heads are made of FR4 rectangular pipes of 6-mm wall thickness and outer dimensions of $25.40 \times 35.00 \times 490.00 \text{ mm}^2$. The chevron pattern is etched on the FR4 skin facing the wire plane. For the plain cathode, the FR4 skin is metalized and the rohacell is 19.05-mm thick, and the panel heads are FR4 with dimensions $19.05 \times 35.00 \times 490.00 \text{ mm}^2$. The hollow panel heads of the chevron cathode also serve as gas inlet and outlet manifolds such that gas flows only through the region between the cathodes. The FR4 spacers between the cathodes define the separation between the wire plane and the cathodes. The sector detectors, which are mechanically locked to each other, are supported by an external frame consisting two parallel 90° sectorial arches interconnected at the ends. The arches are made up of rectangular carbon fiber pipes. Table 6.3 lists the revised characteristics of PC1.

The design of PC2 and PC3 has been changed from two sectors to four sectors per detector arm, similar to the TEC geometry. The PC2 occupies the radial space defined between inscribed radii 415.00 and 421.00 cm, and PC3 sits between 491.00 and 498.00 cm. The mechanical structure of PC2 and PC3 are the same, except that the dimensions are different. Figures. 6.4 and 6.5 show isometric and sectional views of PC2 and PC3. The basic design concept of each sector detector is similar to that used in PC1. except that these detectors are much larger in area and each sector detector requires a carbon fiber frame around its boundary for support. The panel dimensions are $1659.40 \times 3069.80 \text{ mm}^2$ for PC2 and $1960.70 \times 3610.85 \text{ mm}^2$ for PC3. The larger panel area requires more material for strength. For PC2, the rohacell thickness for the chevron cathode is 31.75 mm, and for the plain cathode it is 25.40 mm. For PC3, the rohacell is 38.10 mm thick for the chevron cathode and 25.40 mm for the plain cathode. However, the FR4 skin for the cathode panels remains 250 μ m thick for PC2, but must be increased to 500 μ m for PC3. The rectangular

Figure 6.3: Isometric and sectional views of PC1. All measurements are in millimeters.

Figure 6.4: Isometric and sectional views of PC2. All measurements are in millimeters.

Table 6.3: Pad Chamber Characteristics

Item	PC1	PC2	PC3
No. of sectors/arm	8	4	4
No. of anode (A) wires	122	238	243
A-wire diameter	20 μm	20 μm	20 μm
A-wire tension	30 g	30 g	30 g
A-A separation (2-A-Cathode spacing)	4.0 mm	6.8 mm	7.9 mm
No. of field (F) wires	123	239	244
F-wire diameter	100 μm	100 μm	100 μm
F-wire tension	100 g	100 g	100 g
Chevron width	3.0 mm	5.2 mm	6.3 mm
Ground strip width	0.8 mm	1.4 mm	1.4 mm
Space line width	0.2 mm	0.2 mm	0.2 mm
Active wire length	1777.02 mm	2979.80 mm	3520.85 mm
Panel width	490.0 mm	1619.40 mm	1920.70 mm
Chevron pad length (apex-apex)	246.47 mm	417.92 mm	495.20 mm
F value	1.20	1.12	1.10
No. of readout nodes/wire	8	8	8
Readout nodes/panel	976	1904	1944
Readout nodes/arm	7808	7616	7776
Total readout nodes	15,616	15,232	15,552

hollow pipe FR4 panel head for the chevron cathode has a wall thickness of 6.00 mm and cross-sectional dimensions of $31.75 \times 50.00 \text{ mm}^2$ for PC2 and $38.10 \times 50.00 \text{ mm}^2$ for PC3. For the plain cathode, the panel heads have cross-sectional dimensions $25.40 \times 50.00 \text{ mm}^2$ for both PC2 and PC3. The carbon fiber plates surrounding each sector detector are 20.00-mm thick with a cross-sectional dimension the same as the sector detectors. The sectors, which are mechanically locked to each other, form a 90° four-sector arc, and are supported externally by a pair of parallel sectorial arches (similar to PC1) via adjustable screw bolts. The arch frame is made of rectangular carbon fiber pipes. Table 6.3 also lists the revised specifications for PC2 and PC3.

6.4.2 Electronics

The front-end electronics readout system for the pad chamber comprises a preamplifier, shaper, AMU, and ADC. The detailed designs of these units are being specified. Figure 6.6 illustrates the readout chain for each channel. The preamplifier and the shaper is housed in one eight-channel chip. This preamplifier/shaper chip is a modified version of the unit

Figure 6.5: Isometric and sectional views of PC3. All measurements are in millimeters.

designed at BNL for GEM to handle the large detector capacitance (80–200 pF) with an optimum shaping time of 400 ns. The shaper, which is equipped with pole-zero cancellation, returns the signal to within 0.1% of the baseline in four times the peaking time. The output stage drives a 2-m long low impedance line.

The original plan was to create a readout “mother board” on the back plane of the chevron cathode to house the preamplifier/shaper chip; however, there is a relatively large coupling between the shaper output lines and the chevrons on the opposite side. Therefore, it was necessary to isolate the shaper output from the chevron board in order to minimize these feedback effects. Figure 6.6 illustrates this scheme. Because a multilayer board isolating the chevrons with a ground plane from the signal lines would introduce too much capacitance to the chevrons, their charge signals are brought together in groups of eight on the back of the chevron-cathode plane. These grouped signals are then routed through the rohacell using flexible printed circuit strips (kapton cables) to the outer side of the outer FR4 skin of the cathode panel. An electronic mother board for these signals is then created. The inner surface of the mother board will be metalized to form a ground plane shielding the chevrons from external interferences.

The illustration at the bottom of Fig. 6.6 shows that the AMU, ADC, serial-interface, and heap-manager chips are all located at the edge of the mother board, outside of the active volume of the detector. These chips are being developed at ORNL. The AMU is an eight-channel chip with 64 cells per channel. AMU input signals are to be sampled at three points at a rate of 10 MHz. Each output of the AMU is digitized by an 11-bit ADC. The choice of 11-bits accommodates the full charge range and preserves the 1% spatial resolution of the chamber.

Figure 6.6: Electronic readout scheme: (top) FEE for one channel and (bottom) top view of back plane of chevron cathode board.

6.4.3 Calculations

The spatial resolutions of the pad chambers were simulated in the CDR. In those calculations, it was assumed that the space lines separating the chevrons from each other and from the ground strips in the chevron cathode were uniformly $60\text{-}\mu\text{m}$ wide. However, due to technical reasons, this relatively narrow linewidth can not be preserved in the fabrication of large cathode boards. Industrial assurance can only be obtained for a linewidth of $250\text{ }\mu\text{m}$. With this information, we recalculated the achievable spatial resolution of all three chambers

Figure 6.7: Simulations of spatial resolution of PC1, using space linewidth 250μ . Calculation parameters were taken from Table 6.2.

by optimizing the chevron form factor (F) value and using a new linearization technique. Figures 6.7–6.9 display the results. The resultant resolutions are still within the design specifications in the z coordinate of $\sigma_z = 2, 3$, and 4 mm , respectively. IHEP (Beijing) has facilities to fabricate large cathode boards with an assured uniform-space linewidth of $200 \mu\text{m}$, which would provide further improvement of chamber performance.

6.4.4 Cost and Schedule

The cost estimate and schedule for the pad chambers is presented in Table 6.4. In this table the MVD mechanical (WBS 5.2.5) and front-end electronics (WBS 5.3.5) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

PC2 single standard chevron F=1.12 Noise=0.1%

Figure 6.8: Simulations of spatial resolution of PC2, using space linewidth $250\text{ }\mu\text{m}$. Calculation parameters were taken from Table 6.2.

Table 6.4: Cost Estimate and Schedule for the PHENIX Pad Chambers

WBS	Description	Total Cost	1994				1995				1996				1997				1998				1999			
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
5.2.5	PC	\$1,502,455.40																								
5.2.5.1	PC design	\$127,389.68																								
5.2.5.2	PC full scale prototyping	\$36,545.00																								
5.2.5.3	PC fabrication	\$1,324,979.92																								
5.2.5.4	PC assembly	\$13,540.80																								
5.2.5.5	PC test	\$0.00																								

[illegible]

PC3 single standard chevron $F=1.10$ Noise=0.1%

Figure 6.9: Simulations of spatial resolution of PC3, with space linewidths of $250 \mu\text{m}$. Calculation parameters were taken from Table 6.2.

6.4.5 Pixel Pad Cathode Option

The update of the pixel pad solution concentrates on the studies and modifications done after the completion of the CDR. These include:

1. Test measurements on a small chamber.
2. Simulations based on the test results.
3. Modified mechanical construction based on the mechanical engineering performed for the chevron pads.
4. Cost estimates for IC development, fabrication and chip-on-board assembly.

The pixel-pad solution has many advantages compared to the chevron solution. In particular the fine granularity ensures reliable pattern recognition at extreme multiplicities. The major uncertainty with the pixel system is that it makes use of the chip-on-board (COB) construction technique, which is not yet used in HEP experiments. Therefore, an evaluation of a full scale system is needed to determine whether this new technique is reliable enough to be used in an essential detector system.

As a consequence, the pixel pad solution is not currently a part of the baseline detector, instead it is considered to be an upgrade with the special status that it could potentially have lower cost than the baseline version once additional resources are considered. The goal is to complete a full scale prototype (one sector of PC1) on a time scale that would allow the upgrade option to be implemented in the baseline detector. The development of the pixel readout electronics will be performed by the Lund University group.

Mechanical Design

As described in Section 6.4.2, the back plane will be fabricated from rohacell backed on both sides by 250 μm PC-board material. One of these FR4 plates will be the cathode board. The signals are brought through the rohacell to the top PC board where the preamp/shaper chips are located. This board also has the printed leads bringing the signals to the sides of the detector.

For the pixel readout we anticipate a similar solution; however, the integrated circuits and the leads to the sides are to be on a separate 120- μm thin PC-board fabricated with the COB technique. This board should connect directly to the signal feedthru through the rohacell. The upper skin would then be without a printed pattern.

This solution allows the mechanical design of the pad chambers to be fixed, but keeps open the possibility of subsequently adopting the pixel-pad solution. The method used to bring the signals through the sandwich structure, which is of common interest to both readout solutions, needs to be developed for large scale solutions.

With mounted IC and gloptop the additional readout PC-boards will add about 0.08% of a radiation length to the material in PC1, and a factor four less in PC3 due to the larger area served by each IC. This would add less than 0.01%, if the signal feedthru is made of a 5-mm wide 120- μm PC-board material. We estimate the additional material from the pixel readout electronics to be 0.1% of a radiation length.

Electronics

For this type of readout solution with 900,000 channels it is essential that the cost per channel be kept low (less than one dollar per channel as shown in the WBS) and that the electronics can be mounted on the detector backplane. These goals can be achieved with a high degree of integration and the COB mounting technique.

Since signals are processed locally, little signal driving is required. Together with moderate demands on resolution this ensures low power consumption of the order of a few milliwatts per channel. We plan to develop new integrated circuits for the full readout chain and to place the analog/discriminator function and the digital pipeline in two separate integrated circuits in order to minimize the risk of crosstalk.

We plan to fabricate the IC analog in a 1.2- μm BiCMOS process. It will contain 16 channels of preamp/shaper plus discriminator. The shaping time will be about 100 ns. The output of the discriminator will cover two beam crossings in order to make sure that the time slice corresponding to the event will have the information. Therefore only one time slice needs to be transferred to the DSP. The IC area is estimated to be $3 \times 3 \text{ mm}^2$.

The digital IC will include a 64-cell deep digital pipeline (DMU) timed with the beam clock. It will be followed by a five-events-deep derandomizer and bus output. The width of the DMU will probably be 48 or 64 channels. When transported on a 16-bit wide (10 MHz) data bus, 4096 channels can be read out within the limits set by the average LVL-1 trigger rate (30 μs). The chip area will be about $4 \times 4 \text{ mm}^2$.

Both IC will be mounted as bare silicon wire bonded down to a 120-micron thick PC-board. After the wire bonding the IC and wire bonds are embedded in an epoxy gloptop. This gloptop can be covered by a conductive layer applied by screen printing. The digital and analog parts of the readout system will be shielded from each other by a ground plane.

Once assembled and tested the COB mounting should provide a very reliable and stable system. The difficulty is that the electronics must be tested thoroughly between the stages of bonding and sealing with the gloptop. These "burn-in" tests should be an integral part of the production line. The Lund group has a close collaboration with a COB company in Sweden, which would afford the possibility of performing such tests. The mounting technique is most suitable for electronics with moderate resolution, because lengthy tests of individual chips would not be required.

Simulations and Testing

The basic functional principles of the pixel-pad readout were studied in bench tests performed in Lund. A MWPC with the same 4-mm wire spacing as PC1 was equipped with a cathode plane (2-mm distance to wires) segmented in small rectangular pads as shown in Fig. 6.10. When operated with P10 gas, the chamber was thin enough to allow electrons from a Sr source to pass through the chamber as well as the cathode plane and the entrance/exit windows. The only electrons producing a trigger were the ones in the collimated (1 mm circular) beam from the source. These had high enough kinetic energy to produce a coincidence between a thin (0.1 mm) and a thick scintillator on the opposite side of the chamber. Thus, the electrons producing a readout trigger were simulating the dE/dx of a

Figure 6.10: Schematic of the pad configuration in the Pixel test chamber.

minimum-ionizing particle (mip).

Each pad was read out and software simulated the connection of pads to form the pad structure of the pixel-pad system (i.e., software added Cl + C + Cr in Fig. 6.10). Software was also used to apply a threshold, which simulated the discriminator level.

Figure 6.11 shows the pulse-height spectra of each individual pad when the source is positioned in the center of pad C. The requirement that three pads should always be fired is obviously well fulfilled, since the distributions are shifted away from zero for the three pads closest to the source. It is evident that it is possible to define a threshold which sets the three middle pads true and all others false. In this case the hit would be classified as being in the center of the pixel defined by the three fired pads. In the few cases when a neighbor pad along the wire becomes true, the maximal error would be made and the hit would be reconstructed with an offset of 2 mm. When there is crosstalk to pads under the neighboring anode wire, the hit would be treated as logically incorrect. In reality, such crosstalk events will be accepted as valid hits with a special flag indicating that this may be a hit right between two neighboring pixels (across the wires).

The least favorable case occurs when a hit is between two pads along the wire. When this happens, the charge is shared among a maximum number of pads producing the minimal pulse height, which determines the threshold setting.

Figure 6.12 summarizes the position reconstruction when the source has been moved in steps of 0.5 mm from the center of one pad to the center of the next pad. The figure shows the difference between the reconstructed position and the position of the collimated source. Evidently the hit is always reconstructed to better than ± 2 mm and the width of the distribution is about 1 mm. The position resolution is actually even better, because the

Figure 6.11: Pad pulse height spectra created with source placed at center of pad C.

collimator has an opening diameter of 1 mm.

The reconstruction efficiency was tested for different thresholds and found to be stable for variations of the threshold by at least a factor of three. The absolute efficiency is difficult to measure with a beta source, but the results indicate an efficiency close to 100%. This implies that the system should perform well with moderate demands on resolution in the electronics readout, and that a rather large variation in gas gain over the area of a chamber can be tolerated. The test results are very promising. The position resolution should be substantially better than ± 2 mm, which was the design criterion for PC1.

We plan to fabricate a full size prototype, one octant of PC1, to evaluate the performance and reliability of the system. Within a year we hope to construct a full-length readout chain equipped with prototype ICs. Construction of a full sector would require actual fabrication of integrated circuits and would need about another half year of work. The prototype studies can be made with the existing ICs. This would allow the prototyping to run considerably faster, but the cost and technical feasibility of this option must be carefully evaluated.

A detailed Monte Carlo study was begun to develop the pixel-system software and to investigate the two-track separation. The pulse height and charge-distribution results were tuned to the test measurements. The first studies indicate that when the distance between the hits is 4 mm, a double-hit is identified in 50% of the cases. At 10 mm the two-track

Figure 6.12: Pad position resolution. A collimated source has been moved in 0.5 mm steps between two pads. Reconstructed position minus source position has been plotted and a resolution of ~ 1 mm *rms* has been obtained.

reconstruction is almost perfect, as long as all particles are minimum ionizing. A realistic dE/dx mixture from particles with different energies gives very little difference in the reconstructive efficiency and position resolution. There is only a small increase in the number of hits flagged as possible double hits.

For the LVL-2 trigger it should be possible to identify the specific pixel hit (4×4 mm² square in PC1) and to flag it as either definitely a single hit or possibly a double-hit. For the LVL-1 trigger it should be possible to provide coarse position information, which could be combined with the EMCal electron trigger and possibly with the RICH signal. For LVL-1 the position resolution achievable (several cm in PC3) is limited by the maximum possible number of detector leads.

6.5 Time Expansion Chamber

The main functions of the PHENIX time expansion chamber (TEC) are: (1) to track all charged particles passing through the region between the RICH and the EMCal, (2) to identify tracks entering from outside the fiducial volume, and (3) to determine particle species

using dE/dx information. TEC track measurements are used to produce direction vectors, which allow track reconstruction by matching the tracking information from the drift chambers and pad chambers. Because approximately 25% of the tracks that enter the EMCal do not pass through the drift chambers, the TEC is critical for proper track reconstruction in PHENIX. Particle energy loss in the TEC detector gas is measured using the flash ADC (FADC) electronics. Information derived from these measurements is used to separate electrons from pions. A challenging aspect of the TEC is the design of large area wire chambers, which must maintain a consistent anode-cathode separation, while keeping all frame and support material to a minimum. In addition, the detector electronics parameters must be optimized for the best performance of the various detector functions, which are: (1) track reconstruction, (2) particle identification using dE/dx , and (3) particle identification using transition radiation detection (after the TRD upgrade). Recent progress in the project is concentrated in the design of the detector mechanics and electronics.

6.5.1 Mechanical Design

Four segments in each arm of the PHENIX detector comprise the TEC. Each segment covers 22.5° in ϕ and ± 0.35 in η ($\pm 19.7^\circ$ in θ). The TEC occupies the space between the inscribed radii of 4.217 and 4.91 m centered on the interaction point (Fig. 6.13). Each segment contains six planes of wire chambers spaced 10.8 cm apart. The wire chambers are constructed individually and then assembled into a self-supporting segment. The segments, which interlock with their neighbors through frames located at their boundaries, are to be constructed separately before assembly into a four-segment arm prior to installation into PHENIX.

The wire chambers are composed of an Al-mylar drift cathode, a 3-cm drift space, and a 6-mm thick proportional region containing two cathode wire planes surrounding a plane of $25\text{-}\mu\text{m}$ Au-W anode wires (Fig. 6.14). The stainless steel cathode wire planes are each spaced 3 mm from the anode plane. The anode wires are spaced 4 mm apart. All wires run in the z direction. The second plane of cathode wires is a modification of the TEC CDR design. The motivation was to overcome the difficulty of keeping the distance variation to a few tens of microns between a plane of anode wires and a 4 m^2 , $25\text{-}\mu\text{m}$ thick Al-mylar foil. Each anode wire is divided in z at a thin G-10 wire support structure running across the midpoint of each frame. The maximum length of a TEC anode wire is 1.7 m.

A 7.2-cm thick gap between each wire chamber is reserved for the installation of transition-radiation radiator foil packs for the eventual upgrade of the TEC to a TRD. The chambers have been designed to operate with either an Ar- or Xe-based gas mixture. We presently plan to begin day one operation with a P10 mixture and only 4 of the 6 wire planes instrumented.

The detector frames, made from a combination of G-10 and carbon composite material, allow a frame design that maintains rigidity under wire tension loading but has a minimum amount of material in the active region. Contacts with C-fiber manufacturers assure us that frame construction from carbon composite materials is both physically and fiscally realistic.

A recent modification in the TEC design is the change toward a more cylindrical geometry by going from two to four segments per arm. The change was motivated by the desire to keep the angle of each track through the TEC more closely correlated to its p_T and to reduce

Figure 6.13: Schematic of PHENIX Time Expansion Chamber showing radial dimensions and its octant structure.

the maximum track angle that would be found near the edges of each segment. Studies of the TEC frame design indicated that for frame stability, the two segments per arm design would require a strut running in the z direction to be located in the middle of each segment. Dividing the segment at this location, and changing from two to four segments, was a natural step in the design.

6.5.2 Electronics

The TEC electronics chain comprises a preamplifier, a unipolar shaping amplifier with 70 ns integration time, a 5-bit flash ADC designed to run at 36 MHz (i.e., one quarter of the time between RHIC crossings), and a digital memory chip. The TEC anode wires connect to traces on printed circuit boards that route the signals to detachable motherboards containing the front-end electronics. The preamplifier/shaper, FADC and DMU each exist on separate chips (Fig. 6.15).

The preamp/shaper chip is a modified version of an IC created by the BNL Instrumentation Division for the GEM cathode strip chambers. The IC shaper gain and integration time is to be optimized for the PHENIX TEC, and a pole-zero cancellation stage is to be added. The dE/dx measurement is very sensitive to tail cancellation. The shaping amplifier also has

Figure 6.14: Schematic of TEC wire chamber.

Figure 6.15: Block diagram of TEC electronics chain.

Table 6.5: TEC Detector and Electronics Parameters

Item	Specifications
Total channel count (Baseline)	28,032
Anode wire spacing	4 mm
Drift distance	3 cm
Shaping time	70 ns
Input capacitance	15 pf
Preamp input charge sensitivity	1 fC – 1 pC
Front-end ENC($C_{in} = 15pF$)	1200–1500 electrons
Hit-wire rate	200 Hz

a dual output with an additional gain stage on one of the output lines. Both lines feed into the FADC. This allows the coverage of a dynamic range including both minimum ionizing dE/dx signals and TR signals, while keeping the number of FADC bits to a minimum. The range from the minimum detectable signal from dE/dx and maximum signal from TR x-rays is a factor of 1000.

The 5-bit FADC chip is also being designed at BNL Instrumentation based on a 1.2- μm CMOS process. The design sets the 31 comparator levels, which are 28 in the low dE/dx range and three in the high TRD range. The 28 lower levels are set in a nonlinear distribution to respond to charge signals of 2–120 fC on the input of the preamp. More levels are concentrated in the regions around the charge from the mean and the most probable energy deposit for a mip. The TR comparator levels are set for 400, 600 and 800 fC of charge on the preamp input. The levels are chosen to optimize the particle identification capability of the TEC. The TEC electronics and chamber performance parameters are listed in table 6.5.

6.5.3 Operation and Performance

The TEC is to be operated initially with an argon-based gas mixture, but there is a planned upgrade to a xenon-based gas. With a P10 gas mixture the total drift time is expected to be 1.2 μs . For a minimum ionizing particle the most-probable (mp) energy deposited is $(dE/dx)_{\text{mp}} = 160 \text{ eV/mm}$. For a 70-ns shaping time, and an absolute gas gain of 1.5×10^4 , a minimum ionizing particle depositing $(dE/dx)_{\text{mp}}$ generates an effective anode charge of 10 fC (see Table 6.6). The upgrading of the TEC gas system for full recirculation will allow us to change the TEC gas mixture to 50%Xe + 45%He + 5%C₄H₁₀. For this gas mixture the total drift time is 2.0 μs and $(dE/dx)_{\text{mp}} = 200 \text{ eV/mm}$. For our shaping, time an absolute gas gain of 2.5×10^4 yields an effective anode charge of 10 fC.

6.6 System Performance

6.6.1 Momentum and Mass Resolution

The momentum and mass resolutions of the PHENIX tracking detector were calculated for the present DC detector configuration. The simulation now includes a realistic variation of the position resolution of the detector with distance from the anode wires. The magnetic field in these calculations is stronger than that used in the CDR (0.78 vs. 0.7 Tm), and a small magnetic field is present in the detector volume. The new drift chamber wire arrangement used in these calculations is described earlier in this section. The results are impacted by the geometry change, because the spacing between the front six X wires, and the next six X wires has been reduced from the previous value of 30 cm to the present value of 24 cm. This reduction in lever arm results in a loss of momentum reconstruction accuracy at high momenta. Tracks which pass close to the anode wires in the drift chamber are measured with less position resolution (600 μm vs. 150 μm). This effect was not included in the previous simulation. The result is a loss of the momentum resolution at high p_T . This loss is compensated by the increase in the magnetic field. For low p_T the resolution is dominated by multiple Coulomb scattering, so the results are similar to those obtained previously.

Figure 6.16 shows the momentum resolution of the tracking detector as a function of p_T for the detection of electrons, pions and kaons. The resolutions shown in this plot are very similar to the CDR results. These simulations required the integration of code from several members of the tracking group, and represented significant progress toward integrating all of the tracking-group software with the off-line and on-line computing efforts in PHENIX.

Figure 6.17 is a plot of the vertex resolution obtained from the reconstruction of tracks in the spectrometer. As discussed in the CDR, the vertex resolution in an event with a multiplicity of N_{ch} above a p_T of 0.5 GeV will be $1/\sqrt{N_{ch}}$ cm. Once again, this resolution is similar to that shown in the CDR.

Figure 6.18 is a plot of the mass resolution shown as functions of p_T for the detection of vector mesons. The momentum resolution at low momenta is dominated by multiple Coulomb scattering. At high momenta the resolutions are governed by the position resolution of the detector and by the length of the lever arm. The decays of the ϕ can be measured with a resolution of better than 0.1% in the kaon channel and 0.5% in the electron channel.

The drift chambers and the pad chambers were the only detectors used in these calculations. The resolutions were calculated in a single-particle mode where only one particle of the species of interest was detected within the spectrometer acceptance. The results indicate that the ϕ can be measured with a resolution of about 0.07% for a pair momentum below 1.0 GeV. For decays via electrons the corresponding resolution is about 0.5%. The J/ψ can be measured with a resolution of about 0.7 %. These results are very similar to those presented in the CDR. The detector performance (position and angular resolutions) have been simulated here with better accuracy than before.

Figure 6.16: Momentum resolution as a function of p_T
Vertex Resolution (single track)

Figure 6.17: Vertex resolution as a function of p_T

Figure 6.18: Mass resolution as a function of p_T for the ϕ and J/ψ

6.6.2 Pattern Recognition

We have continued to work on the development of pattern recognition for the PHENIX spectrometer. The most difficult aspects of the pattern recognition involve the identification of low momentum tracks. We have developed a multiloop pattern recognition procedure to enable the detection of low momentum tracks with high efficiency. First we set tight criteria on the "roads" used to look for good track candidates. Once the high p_T tracks are found, we remove those hits from consideration. We then increase the size of search windows, and repeat the track finding procedure to find tracks for intermediate p_T . In the third and last iteration, the tracks for the lowest-momenta particles are found. Figure 6.19 is a plot of the ratios of the numbers of tracks found to the number of true tracks. Plotted are ratios both for the tracks detected in an "unknown" track mode (a *realistic* situation) and for a "known" track mode in which the known hit information from the detectors was used in the reconstruction of the tracks. The "known" track mode represents the efficiency of the cuts for finding good tracks. The "unknown" track mode includes the ghost tracks. These calculations will be upgraded soon with the addition of new code for the pad chambers. Furthermore, the information from the TEC is presently not incorporated into the simulation. Its inclusion should aid in the reduction of the number of ghost tracks.

Figure 6.19 shows the present status of our pattern recognition calculation. In the earlier version of the code (data points connected by lines) the reconstruction of double HIJET events resulted in a large number of ghost tracks. The open symbols show results for the

Figure 6.19: Pattern recognition showing reconstruction efficiency and number of ghost tracks. The open symbols represent the newest results of the tracking code.

present multiloop pattern recognition calculations. For the same “known” track efficiency there is a lower ghost track fraction in the “unknown” track mode. Similar improvements are evident for the single point shown for a low p_T sample of pion events, when compared to the older calculations.

Bibliography

- [1] Y. Arai, IEEE J. of Solid-State Circuits, **27**, No. 3, 359-364, March 1992.

Chapter 7

Ring Imaging Cherenkov Detector

7.1 Overview

The Ring Imaging Cherenkov (RICH) detector subsystem serves as one of the primary devices for the identification of electrons. The scope of this chapter is to describe the progress in the design of the RICH detector subsystem since the PHENIX CDR.

Figure 7.1, which appeared in the CDR, is shown again to illustrate the principle of electron detection in the RICH. A typical electron track (400 MeV/c) is shown, together with Cherenkov photons emitted in the radiator gas; the photons are reflected and focused by the mirror.

Figure 7.1: Cherenkov photons emitted by a 400 MeV/c electron are reflected and focused by the mirror (from PISA simulation).

7.2 Changes

The major changes in the RICH subsystem since the CDR include:

1. The boundaries of the RICH subsystem were slightly changed.
2. Magnetic shielding rings (iron pieces) were added to the design of the magnet yoke, which will reduce the fringe field at the photon counters.
3. Changes in the optics caused by changes in the geometry are being investigated.
4. The mechanical design of the RICH gas vessel and windows is in progress.
5. The mirror segmentation has been determined.
6. The mirror support structure has now been designed and a method of aligning the mirror segments has been developed.
7. Photomultiplier tubes were chosen to be the focal plane detectors.
8. A production version of the Winstone cone was designed.
9. Magnetic shielding of the photon detector was investigated.
10. The conceptual design of the front-end electronics is being reevaluated.
11. A prototype HV power supply was designed and fabricated.

7.3 Mechanical

7.3.1 Global Geometry and Optics

The outer radial extension of the RICH subsystem was changed from 400 to 410 cm, which improves the optics and significantly increases the minimum path length of electrons in the radiator gas.

Magnetic shielding rings (iron pieces, 10-cm high and 5-cm thick) were added to the design of the outer edge of the magnet yoke. The purpose of the rings is to reduce the magnetic field strength at the photon counters. Without the shielding rings, the calculated field strength at the inner most photon detector array is approximately 200 G. With the shielding rings, a preliminary calculation shows a reduction in the magnetic field strength by more than a factor of two. Therefore, the addition of the iron pieces to the magnet yoke reduces significantly the magnetic shielding requirements for the individual PMTs that comprise the focal plane detector.

The addition of the magnetic shield ring to the magnet resulted in the inside z boundary of the photon detector recess being moved out to 137.3 cm, in turn pushing the photon detector to larger z . To maintain acceptable efficiency for all events, this forced an increase in the tilt of the RICH mirror by moving the center of curvature to a larger z . Even after much

Figure 7.2: Cross section of the RICH gas vessel.

discussion and debate about minimum radiator gas thickness versus RICH efficiency, the issue of how great the mirror tilt should be has not been entirely settled. The configuration shown in Fig. 7.2 has the mirror center of curvature at $z = 215$ cm, giving a minimum radiator gas thickness of 83.4 cm. Because of space limitations, the maximum possible radius of curvature for the mirror is 403 cm. This limitation is almost independent of the mirror tilt.

7.3.2 Radiator Gas Vessel

Each RICH arm consists of a self-supporting radiator gas vessel, in which two mirrors and two focal-plane photon detectors are mounted. The gas vessel is approximately 1.5-m deep in the r direction, 6-m wide in the z direction, and spans 90° in ϕ .

A finite-element analysis of the RICH gas vessel structure has been used to refine and to check the mechanical design. This structural analysis is now complete, and the design of the structure is close to being settled. We are in the process of gathering quotations for the components of the gas vessel.

A schematic view showing the cross section of the gas vessel appears in Fig. 7.2. The structure will consist of an aluminum frame with 0.25-in. thick aluminum sheet walls. The support beams making up the frame will be 0.25 in. wall rectangular aluminum tubes,

Figure 7.3: A schematic view showing the mechanical arrangement for supporting the Kapton exit window. It also shows the carbon-fiber/epoxy beams used to support the mirrors.

varying in dimensions from $1 \times 5 \text{ in.}^2$ (in locations where the wall thickness is critical) to $3 \times 3 \text{ in.}^2$. When the RICH gas vessel is installed in PHENIX, it will be supported along two lines at the lower end and at three points at the top end.

There is no metal within 23° of the center line of the detector. The entrance and exit windows are 0.003-in. thick Kapton film, supported against the 2.5 lb/ft^2 over pressure of the radiator gas by 0.063-in. wall thermoplastic pultrusions (tubes) of dimensions $0.5 \times 2 \text{ in.}^2$ for the entrance window, and $0.75 \times 3 \text{ in.}^2$ for the exit window. The arrangement is shown schematically for the exit window in Fig. 7.3.

Finally, a $48 \times 144 \times 2 \text{ in.}^3$ gas vessel was constructed at Florida State University in order to test methods for mounting, supporting, and sealing the entrance and exit windows. The gas vessel provides a platform for testing of a prototype bubbler, which is designed to maintain a gas pressure of 2.5 lb/ft^2 in the gas vessel at flow rates of up to 100 L/min .

7.3.3 Mirror System

The mirror segments will be supported by rectangular carbon-fiber/epoxy tubes formed on a Nonex honeycomb core. The tubes will have outside dimensions of $1 \times 3 \text{ in.}^2$, and the carbon-fiber/epoxy walls will be 0.050-in. thick. The mirror in each gas vessel will be divided into 48

Figure 7.4: A schematic view of the mirror segmentation.

segments, with the segment size chosen to minimize the overall mirror cost. Each segment will consist of a carbon-fiber/epoxy substrate, with a 3 nm *rms* surface roughness finish applied by replication. Fig. 7.4 shows the mirror segmentation. Note that the segments adjacent to the cusp of the mirror are slightly smaller than those on the outside.

There is still some discussion about whether the mirror substrate should be a somewhat flexible 0.050-in. thick carbon-fiber/epoxy laminate, similar to the small mirror that was tested successfully at the AGS in June, or a more rigid structure made by sandwiching approximately 1 cm of foam between two 0.020-in. thick carbon-fiber/epoxy laminates. The substrate and replication masters will be identical in either case, and the final decision on the structure of the substrate can wait until after we have made a full size substrate both ways. Although there are two different mirror segment sizes, the same masters can be used for making all of the substrates and for performing the replication. The mirrors will be trimmed to the correct sizes at the end of the manufacturing process.

The scheme that will be used to mount and align the mirror segments is illustrated in Fig. 7.5. Threaded plastic attachment points will be epoxied permanently to the back of each mirror substrate during the mirror manufacturing process. A plastic cup will be threaded tightly into each attachment point. Each mirror segment will be installed using a removable alignment fixture that will be attached to the carbon-fiber/epoxy mirror support beams by

RICH MIRROR MOUNTING DETAILS

Figure 7.5: A scheme to mount and align the mirror segments.

clamps, which are not shown in Fig. 7.5. The alignment fixture will have a clamp with a micrometer position adjustment at each attachment point on the mirror segment. Once the mirror segment is clamped to it, the micrometer adjustments on the alignment fixture will allow precise positioning of the mirror segment relative to the carbon-fiber support beams. Once the mirror segment has been aligned, a permanent epoxy connection will be made between the cup threaded into each mounting point and a plastic clamp on the nearest carbon fiber support beam. The alignment fixture will then be removed so that it can be used to mount another mirror segment. After alignment, the mirror may be removed and reinstalled without disturbing the alignment. If the mirror segment has to be realigned for some reason, the plastic cups that are threaded into the attachment points have to be removed and discarded, along with the part of the clamp that is epoxied permanently to each one. After new plastic cups are threaded into the attachment points, a new alignment can be carried out using the procedure outlined above.

7.3.4 Gas Handling System

The gas handling system has not yet been reevaluated in the light of the decision to use PMT technology rather than CsI chamber technology for photon detection in the RICH. However, some general remarks can be made.

The use of CsI chambers would have imposed very strict requirements on the radiator gas purity. In the wavelength range of 165 to 200 nm, where the CsI chambers operate, a significant fraction of the Cherenkov light can be absorbed by parts per million levels of water vapor and oxygen contaminating the radiator gas. Our decision to use PMTs as the RICH photon detectors eliminates these concerns. The 200 nm lower wavelength cutoff imposed on the PMTs by their UV glass windows is above the wavelength region where water vapor and oxygen absorption are significant.

The dramatic relaxation of the radiator gas purity requirements means that we no longer have to think in terms of flowing high purity radiator gas through the RICH at very high flow rates. It is likely that it will not be cost effective to recirculate and purify the radiator gas, because a rather small flow rate of not very pure (i.e., inexpensive) gas should be all that is required to maintain adequate levels of purity. This is being studied now.

The gas handling system will still have to be able to support very high flow rates (100 L/min) so that the RICH gas vessel can be flushed in a reasonable length of time when we need to open it for maintenance.

One consequence of adopting the PMT option is removal of heat generated by the HV resistor chain. Each PMT assembly produces about 0.3 W of heat, so the total heat production amounts to approximately 1 kW per RICH arm. Some cooling will have to be provided to compensate for this.

We are presently working on the conceptual design for a gas handling system that is tailored to the PMT version of the RICH.

7.3.5 Focal Plane Detector

The photomultiplier tube option was chosen as the focal plane detector technology for the baseline PHENIX RICH detector. The Cherenkov photons are detected by using arrays of UV-sensitive photomultiplier tubes. In order to reduce dead space on the focal plane, a light-collecting funnel (i.e., a Winston cone) is attached to each PMT. A total of 6400 PMTs are required to cover the four focal planes, which cover a total area of approximately 20 m².

Technology Choice

Two technologies were investigated for the photon detector: (1) a highly segmented array of photomultiplier tubes, and (2) a UV-sensitive wire chamber, using a CsI-coated photocathode.

As described in the CDR, the test of the PMT array prototype at KEK last year (10/92), as well as the test last June (6/93) at BNL, demonstrated that the PMT option can meet all the performance goals set for the baseline PHENIX RICH detector. A drawback of the PMT solution is its high cost per pixel. The CsI photocathode option can potentially lead to lower cost at a fixed segmentation or to higher segmentation at a fixed cost, as compared to the PMT version. The test done at BNL last June showed that a prototype CsI photocathode chamber performed almost perfectly and met performance expectations.

The number of photons produced with a 1.2-m thickness of ethane radiator was approximately 20 for the PMT option and approximately eight for the CsI photocathode option. The

smaller number of photons measured using the CsI photocathode option comes mainly from the wavelength cutoff of the quartz window at about 170 nm. A larger number is expected if a CaF_2 window is used, but this option was excluded due to the formidable cost of the material. The CsI photocathode option excludes the use of methane or nitrogen as a radiator gas.

We have serious concerns about whether the chamber performance can be maintained in the long term, since the technology is too new for us to have enough data on long-term stability. There is also a serious question as to how well adequate gas purity can be maintained in such a large gas volume. Since the PMT option utilizes only photons with wave length larger than 200 nm, absorption of Cherenkov light by oxygen and water vapor in the radiator gas is not important, and so the gas purity requirements are greatly relaxed over those for the CsI chamber, where ppm levels of oxygen and water vapor cause significant absorption of Cherenkov light. Therefore we adopt the PMT option as our technology choice, although we acknowledge that this choice is quite conservative.

Photomultiplier Tube

Our current choice for the PMT is the Hamamatsu H3171S (gain-selected H3171). As described in the CDR, this tube satisfies the performance requirements. Also being investigated is the possibility of using Russian phototubes, which would allow for a significant cost reduction, if they meet the performance requirements.

Winston Cone and Magnetic Shielding

A production version of the Winston cone has been designed. This is a modified version of the one used in the previous beam tests. The thickness of the wall at the entrance is reduced from 2 mm to 1 mm. The Winston cone accepts photons with incident angle less than $\sin \theta \leq 0.5$ away from the normal to the photocathode plane.

As was mentioned in the overview section, a calculation shows a fringe field having a maximum of several tens of Gauss exists at the region of the PMT array. Several shielding schemes have been tested. Our goal is to obtain a solution which guarantees less than 5% change in output signal under a magnetic field of up to 100 G normal to the photocathode plane. Two layers of soft iron with thickness of 1 mm is needed to fulfill this goal. Work on optimizing the shielding is still in progress.

PMT Array Supermodule

We are currently investigating how best to mount the phototubes in the focal plane. The tubes are grouped into modules, called supermodules, which represents the basic unit for testing, installation, readout, and maintenance unit. The identical supermodules are prefabricated, tested, and mounted on the focal planes.

7.4 Electronics

7.4.1 Front-end Electronics

Both analog and timing information from each phototube will be recorded. In addition, an analog sum of 25 PMTs will be used in the LVL-1 trigger. Specifications for the readout electronics is summarized here.

Analog: Each PMT will be operated at a gain of approximately 10^7 , which generates 1.6 pC of charge per photoelectron. Amplification by a factor of 10 gives signals of 16 pC. Splitting the signal into three gives about 5 pC per photoelectron. There will be at most several photoelectrons per PMT. Therefore, the maximum charge per PMT is approximately 25 pC. A 50 pC range with 250 bins (0.2 pC per bin) is sufficient; therefore, an 8 bit range is required.

Timing: Each PMT is expected to have an *rms* time resolution of less than 1 ns for a single photoelectron before a slewing correction is applied. A timing resolution of approximately 350 ps was obtained after a slewing correction was applied to a testbeam measurement. For the time measurement, 200 ps per bin is required, which is much less strict than the requirement for the counters in the TOF system. The spread in arrival times of the Cherenkov photons at the PMT's is estimated to be less than several nanoseconds. A practical minimum time range needed is about 30 ns. If 50 ns is used as a conservative estimate, the time range can be covered with an 8 bit range.

Analog Sum: The two signals discussed above are from individual PMTs. A sum of the output from 25 PMTs is required for the LVL-1 trigger. The number 25 was not chosen arbitrarily; a Cherenkov ring is nearly contained in a 5×5 PMT array. Because there are between 20-30 photoelectrons per electron track per Cherenkov ring, depending on the path of the electron tracks in the Cherenkov radiator, a maximum charge of about 150 pC (30×5 pC) is expected. An analog sum with the capability of processing 25 PMT outputs and with a prompt digital output having 3 bit resolution will be implemented.

A readout scheme very similar to that designed for the beam-beam counters is being considered for the RICH subsystem. As mentioned above, the specifications for the RICH analog dynamic range and timing resolution are less demanding than those for the beam-beam counters and the TOF subsystem.

As reported in the CDR, a prototype of the front-end electronics (FEE) for the RICH PMT readout was fabricated and tested. A study using this prototype revealed a potential problem. Although the prototype worked perfectly well when the external clock was 224 ns (the Day-1 RHIC beam crossing interval), the capacitor discharge and the digitization could not keep up with a 112 ns clock rate (the higher luminosity running condition).

The digitization problem was due to a misconception about how the flash ADC operates. The analog input signal is sampled when the flash ADC is gated by the beam pulse. However the digital output of the flash ADC appears only after two additional gating pulses are

applied. Our first solution was to derive an extra two pulses from each beam pulse and apply all three successively to the flash ADC to obtain the digital output before the next beam pulse arrived. However, since the flash ADC has only 20 MHz bandwidth, it takes 200 ns to gate it three times. Therefore this scheme works for an external clock of 224 ns, but not for an external clock of 112 ns.

The problem was solved after we realized that the flash ADC's internal stack is really a FIFO, so we can use a pipeline readout scheme in which we get the digital output after the flash ADC has been gated by the next two beam pulses. The problem of capacitor discharge was investigated further, and it does not appear to be a critical problem. However, reevaluation of the RICH front-end conceptual design, which includes the interface to the DAQ system, is in progress.

7.4.2 High Voltage Power Supply

A prototype high voltage power supply was fabricated that can supply up to 1.8 kV with maximum current of 3 mA. Ten phototubes can be operated with one supply. The supply uses a DC-to-DC converter to generate a high voltage source from a 6-V DC power source, with conversion efficiency of approximately 50%. Output voltage is controlled by another low voltage input. A test is in progress to study: (1) the noise generated by the power supply, (2) the long term voltage stability, and (3) the reproducibility of the output voltage set by the control voltage.

7.5 Performance

In order to evaluate the performance of the proposed RICH detector in a realistic environment, a beam test was performed at KEK in October 1992, as reported in CDR, and at BNL in June 1993. For the PMT option, the test results at BNL paralleled and verified the results obtained at KEK. Near the end of the beam test at BNL, a 20-cm diameter prototype thin mirror was tested successfully with the PMT array. The thin mirror, a 0.050-in. thick carbon-fiber/epoxy substrate with a replicated reflecting surface, gave results that were indistinguishable from those obtained with the thick Pyrex mirror used in all previous tests.

7.6 Cost and Schedule

The cost estimate and schedule for the RICH is presented in Table 7.1. In this table the RICH mechanical (WBS 5.2.7) and front-end electronics (WBS 5.3.7) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

WBS	Description	Total Cost	1993				1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	
5.2.7	RICH	\$4,678,648.80																								
5.2.7.1	Design	\$185,832.00																								
5.2.7.2	Prototype	\$270,000.00																								
5.2.7.3	Fabrication	\$4,118,016.80																								
5.2.7.4	Assembly	\$42,800.00																								
5.2.7.5	Test	\$82,000.00																								

WBS	Description	Total Cost	No	1993				1994				1995				1996				1997				1998			
				Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4		
5.3.7	RICH FEE	\$1,144,197.52																									
5.3.7.1	Design	\$82,000.00																									
5.3.7.2	Prototype	\$65,003.20																									
5.3.7.3	Fabrication	\$982,994.32																									
5.3.7.4	Assembly	\$16,000.00																									
5.3.7.5	Test	\$38,200.00																									

Chapter 8

Time-of-Flight System

8.1 Overview

The PHENIX time-of-flight (TOF) system identifies hadrons. The TOF system consists of 1024 elements of scintillator slats with photomultiplier tube readout. The radial distance is approximately 5 m from the vertex, and the angular acceptance is 40° in θ and 30° in ϕ . The time resolution of the TOF system is about 80 ps, which provides for a π/K separation of up to 2.4 GeV/c in 4σ .

The TOF system has an excellent trigger capability, which is especially important for $p + p$ collisions at RHIC. In addition, the good timing resolution of the TOF system, in conjunction with tracking information, provides an excellent particle identification trigger.

8.2 Changes

1. The slat sizes and the overall design of the TOF system are changed somewhat since the CDR, because the geometrical segmentation of the outer detectors in the two central arms of PHENIX is now an octant (four sectors per arm), instead of the previous quadrant (two sectors per arm).
2. The chosen TOF configuration is CDR design A, in which a scintillator rod is attached to two prism-shaped light guides and the phototubes. There are two azimuthal sections: one with the acceptance of a full octant (22.5° in ϕ) and the other covering one-third of the adjacent octant (7.5° in ϕ).
3. Testbeam results of TOF slats in Design A were obtained that indicate it is possible to achieve time resolutions of better than 90 ps.
4. There are no significant changes in the TOF front-end electronics, but there is progress on the design and prototype tests of new discriminators, TVC/AMU, and integrator/AMU chips.

Figure 8.1: A typical TOF slat showing a scintillator bar of cross section ($1.5 \times 1.5 \text{ cm}^2$) attached to two prism shaped light guides and the PMTs.

8.3 Mechanical Design

The TOF system is located between the PC2 pad chamber and the lead-glass calorimeter in the allocated radial space between 5.030 and 5.180 m in r . The coverage is the full η range ($70^\circ \leq \theta \leq 110^\circ$) of the central arm detectors over 30° in ϕ . The TOF slats are oriented along the r - ϕ direction.

The changes in the configuration of the outer detectors (TOF, TEC, and EMCal) are described in Chapters 2, 6, and 9. Instead of the previous quadrant design (two sectors per arm), an octant structure (four sectors per arm) is now adopted. In this new design each detector octant covers 22.5° in ϕ ; therefore, the TOF system covering 30° in ϕ with 1024 slats occupies one full octant (22.5° in ϕ and 768 slats) and a third of the adjacent octant (7.5° in ϕ and 256 slats). In the CDR two options for the mechanical design were described and labelled as designs A and B. After performing beam tests at KEK and BNL and evaluating both options, we chose design A for the TOF system. Certain changes to this design were necessary to accommodate the octant structure.

Figure 8.1 shows the design of a typical scintillator slat. The cross section of each scintillator slat is square with a dimension of $1.5 \times 1.5 \text{ cm}^2$. Each end of the scintillator is attached to a prism-shaped light guide and to a photomultiplier tube. The arrangement uses BC404 plastic scintillators produced by the Bicron corporation and R3478S photomultiplier tubes produced by Hamamatsu. The characteristics of the scintillators and photomultiplier tubes are described in the CDR. The performance of the scintillator slats is described in Section 8.5.

Figure 8.2 illustrates the arrangement of the TOF slats (left), and the location of the TOF system in relation to the position of the TEC and the EMCal (right). The slats are oriented along the r - ϕ direction. In the fully covered octant, where the TOF system covers the full 22.5° in ϕ , three different lengths of scintillators are used in order to avoid geometrical conflict among neighboring photomultiplier tubes. Part of the arrangement, indicating the arrangement of two columns of slats, is shown at the left side of Fig. 8.2. The first column comprises three scintillators with lengths of 55, 65, and 75 cm, and for the second column the lengths are 75, 65, and 55 cm. This arrangement is repeated over the full section in the octant. Notice that the 75-cm long slat has a different light guide arrangement on one side; instead of a prism shape light guide, there is a 30° bent light pipe. This is necessary to avoid overlapping the photomultiplier tubes at the ends (on top and bottom) of the TOF system.

Figure 8.2: The arrangement of the three lengths of scintillator slats into two different column types is shown in the left side of the figure. The location of the TOF system in relation to the position of the TEC and the EMCAL is shown on the right.

For the partially covered octant, where the TOF system covers only 7.5° in ϕ , the requirements are less stringent, and we chose to use 65-cm long scintillator slats. For all the slats in this octant, one end of the scintillator slat is attached to a prism type light pipe, whereas the other end has a 30° bent light pipe. The slat positions for the light guides will alternate in order to properly accommodate the photomultiplier tubes.

The overall design for the mechanical support structure is essentially the same as the CDR version. Details of the mechanical support structure are being developed.

8.4 Electronics

The design of the TOF front-end electronics was described in CDR Chapters 8 and 11. Three cables are attached to each photomultiplier tube: a fiber optic cable attached to the face of the photomultiplier tube, a high voltage cable on the tube base, and a signal cable. These cables exit the counters along the z direction. The high voltage cables connect to a high voltage distribution box, and the fiber optic cables go to a pulsed laser system.

Crates containing the TOF front-end electronics are mounted close to the TOF system. Each crate has slots for 16 boards and each board has 16 channels for eight TOF slats. The signal from each of the tubes is split into two, one of which is used to provide the charge and the other is used for timing information. The charge signal goes first to an integrator, then to an analog memory pipeline, and finally to an ADC. The timing signal goes first to a discriminator, and then to the LVL-1 trigger. A second timing output from the discriminator is used as a start pulse for the TVC. Progress continues on the design and testing of the new discriminator boards, TVC/AMU chips and integrator/AMU chips. Final designs are expected soon.

8.5 Performance

In the CDR two options were proposed for the configuration of the TOF slats and were labelled as designs A and B. In design A a scintillator rod was attached to two prism shaped light guides and the photomultiplier tubes, whereas in design B a scintillator rod was attached to either straight light guide or to the light guides, which are bent twice. The scintillator rod for design A was 85-cm long with a square cross section of $1.2 \times 1.2 \text{ cm}^2$, while for design B the rod was $130 \times 0.8 \times 0.8 \text{ cm}^3$.

The relative timing performance of the two options was studied systematically using the test beam facilities at KEK and BNL. The timing resolutions for designs A and B were studied at various positions along the scintillator slat. Whereas both designs achieved less than 90 ps at the center of the slat, the timing resolutions at the edge were found to be considerably poorer for design B than for design A. This results from the fact that the light attenuation of design B is larger than that of Design A. This conclusion is consistent with our previous R&D studies with long plastic scintillators: the timing resolution of a long slat correlates strongly with the light attenuation length, which improves for shorter and wider (thicker) scintillator rods.

Figure 8.3: Time-of-flight spectrum (left) and a scatter plot of time-of-flight vs. ADC.

Because the overall timing resolution is better and the mechanical installation is simpler, design A is chosen for the PHENIX TOF detector. Here below are the details of the timing performance tests.

We tested the scintillator rod, which was made of Bicron BC404 and had dimensions of $100 \times 1.2 \times 1.2 \text{ cm}^2$. Prism-shaped Lucite light guides are attached on both ends of the rod to reflect the scintillation light to the photo cathode of the tube (see Fig. 8.2). The same photomultiplier tube chosen for the PHENIX TOF, Hamamatsu R3478S, was used in the tests. The timing resolution of the start counter used for the tests was less than 50 ps.

Comparison of the light yields, with and without the prism-shaped light guides, indicated that the light loss in the light guide is about 40%, which is consistent with our earlier estimates. Figure 8.3 shows the typical timing performance; the left panel shows the TOF distribution between the test counter and the start counter. The distribution is consistent with a Gaussian distribution ($\sigma = 88 \text{ ps}$) over more than two orders of magnitude. The right panel of Figure 8.3 gives a scatter plot showing a clean correlation between the TOF and the charge of the pulse after the slewing correction.

Figure 8.4 shows the position dependence of the TOF resolution, which is observed to be approximately 90 ps for all locations measured. The overall timing resolution obtained with the design A provides adequate timing performance for the PHENIX TOF.

As described in Section 8.3 above, the design of the PHENIX TOF was modified to accommodate the recent design change for the tracking chambers and the EM calorimeter. The typical size of the scintillator rods is now $65 \times 1.5 \times 1.5 \text{ cm}^3$; therefore, the length of the

Figure 8.4: The position dependence of the TOF resolution.

scintillator is shorter by 35 cm, and the width and the thickness is wider by 0.3 cm than the counter tested. Our previous studies of scintillator rods demonstrate that the timing resolution correlates strongly with the light attenuation along the scintillator, and that the light attenuation length scales roughly with the inverse of the width (thickness) of the rod. Thus, the new design with the shorter and the wider (thicker) scintillator should provide even better timing resolution.

In the current design, the photomultiplier tube is located behind the scintillator; therefore, there it is possible for a particle to travel through both the scintillator and the photocathode at the same time. Such a particle may produce Cherenkov light in the photocathode in addition to the scintillation light. The Cherenkov light may affect the timing of the leading edge of the scintillation light pulse, because the emission mechanisms for scintillation and Cherenkov light are different. The effect of the Cherenkov light emission on the timing was fully investigated. The timing differences between the photomultiplier tubes were measured at various locations. Figure 8.5 shows a plot of the measured minus expected timing differences versus the hit position of the particle at the scintillator rod.

All the data points, except those at the edge of the scintillator, show the same timing difference. This suggests that the velocity of light traversing along the scintillator is constant. However, at the edge region the time jumps by about 150 ps, which corresponds to the difference of emission times between the Cherenkov and scintillation light. The timing jump is seen in the region of 2–3 mm from the surface of the photomultiplier tube, which is the approximate thickness of photocathodes.

The studies described above demonstrate that the effect of Cherenkov light emission is detectable by looking at the difference of timing. In other words, when the hit positions of the

Figure 8.5: The difference between the measured and expected times for different hit positions.

particle along the rod are determined from the timing difference between two photomultiplier tubes, the earlier timing of the affected Cherenkov light produces the observed hit position outside of the rod. Fiducial cuts made on the scintillator rod can easily eliminate such tracks. Because the area of the photocathode seen from the side is $0.4 \times 0.4 \text{ cm}^2$, the probability for an effect from the Cherenkov emission is less than 1% of the total events recorded. Furthermore, as noted above, such Cherenkov-produced events are easily detectable and separable from the others.

Chapter 9

Electromagnetic Calorimetry

9.1 Overview

The primary role of the electromagnetic calorimeter (EMCal) system in PHENIX is to provide a measurement of the energies and trajectories of photons and electrons produced in heavy ion collisions. The sizes of the calorimeter cells are $55.2 \times 55.2 \text{ mm}^2$ and $40.0 \times 40.0 \text{ mm}^2$ in the lead-scintillator (PbSc) and lead-glass (PbGl) sections, respectively. These cell sizes, coupled with the substantial radial distance from the interaction region, result in low enough occupancy to carry out the desired measurements, even in the highest multiplicity Au + Au events at RHIC.

We will use the calorimeter to trigger on rare events with high p_T photons and electrons. We will also incorporate the EMCal signals in LVL-1 triggers for high multiplicity or large total transverse energy.

The calorimeter is read out with phototubes, which preserve its good intrinsic timing characteristics and allow some rejection of photon and electron backgrounds by time-of-flight measurements. EMCal energy determinations combined with momentum measurements from the tracking system are expected to provide an e/π suppression factor of about 10^{-2} at 1 GeV.

The EMCal design is based on planar sectors of individual modules that are first assembled into supermodules. Each of the two central arms of PHENIX comprises four sectors (octants) of EMCal modules. Already existing are a total of 10,000 PbGl blocks assembled into 24 block supermodules, which are currently used in CERN experiment WA98. The lead-glass system, which provides one-quarter of the EMCal coverage (two sectors), will be moved to BNL in time for installation and commissioning. The remainder of the coverage is lead-scintillator modules, which are based on Shish-kebab geometry waveshifter fiber read-out. Because PbSc modules have superior timing performance, the PHENIX high-resolution TOF system will be placed in front of the PbGl walls, which means that in these two sectors the calorimeter is 20 cm farther from the beam.

The PbSc design is nearly complete. A wide range of tests have been performed with individual modules both in the laboratory, at the TRIUMF M13 beamline (March 1993) and in the B2 beamline of the AGS (July 1993). The purpose of these tests was to explore

different designs and technologies mentioned in the CDR and to test their performance.

Meanwhile, components of a 256 channel supermodule were fabricated in Russia and assembled at BNL during July and August. This supermodule will be tested at the AGS next spring and give us the chance to test some of the design choices in a large system.

9.2 Changes

There is significant progress both toward final design choices (for example, the components for the PbSc optical system) and in the development of the calibration and monitoring system, which is an important part of the system design. Significant changes and developments since the CDR include:

1. **Wall geometry change:** the wall structure is changed from four to eight sectors in ϕ , which alters most mechanical parameters of the PbSc section, but is not expected to affect overall performance.
2. **PbSc design choices:** there are new results and progress toward the design in the areas of scintillator-fiber material, calibration components, PMT's and High Voltage distribution.
3. **Progress on PbGl monitoring system:** in parallel with supermodule construction for CERN experiment WA98, a monitoring system is being developed for the PbGl calorimeters to be used in PHENIX.

9.3 Lead-scintillator Calorimeter

9.3.1 Supermodule and Wall Design

Recently, the PHENIX tracking group proposed a design change, which was later adopted, for the Time Expansion Chamber (TEC). The geometry is now more nearly cylindrical with four, instead of two, sectors in each of the two arms (see Chapter 6). The TEC chamber frames position at least one- X_0 of material in front of the EMCal right at the middle of a wall, which compromises both the electron and photon measurements, particularly at the trigger level. Therefore, it was decided to split the EMCal walls to follow the new TEC geometry by building octants instead of quadrants (four walls per arm). The new geometry is shown in Fig. 9.1. Each octant covers 22.5° in azimuth, amounting to a total of 180° . However, the fiducial volume for electrons and photons decreases by about 5% due to the newly introduced edges. The impact on jet physics will be investigated. The z -extent of the calorimeter walls is chosen to match the tracking coverage of $\pm 10.6^\circ$ in ϕ ($\delta\eta = \pm 0.35$). For the PbSc the projection of the tracking coverage to the back of the calorimeter is more than a tower inside the edges. The azimuthal symmetry is slightly broken because the radial distance of the lead-glass octants must be 20 cm larger to accommodate the TOF detector.

The entire support structure (wall) is being redesigned. Also, the supermodule, module and tile dimensions are changed to suit the new wall geometry. The critical parameters of the

Figure 9.1: New radial envelope of the electromagnetic calorimeter.

Table 9.1: Critical Parameters of the Lead-scintillator and Lead-glass Mechanical Design

Item	PbSc	PbGl
Dimension (active area) of single wall	2005.2×4011.4 mm	2132.0×3936.0 mm
Number of modules in a wall	$18 \times 36 = 648$ modules	$52 \times 96 = 4992$ modules
Dimensions of supermodule	6×6 modules	4×6 modules
Wall dimension	3×6 supermodules	13×16 supermodules
Module cross section	110.94×110.94 mm	40.0×40.0 mm
Scintillator dimensions	$55.2 \times 55.2 \times 4.0$ mm	(not applicable)
Spacing between holes for fibers	9.27 mm	(not applicable)

new design are summarized in Tables 9.1. A full mechanical design of the PbSc calorimeter is underway at IHEP, Protvino, and should be complete by the end of 1993.

9.3.2 Module Design

Fiber Choice

The dependence of light yield on the choice of wavelength shifting (WLS) fiber was studied. The starting point was single-clad BICRON fibers with fast green fluor (G-2) used in the original tests. When the fluor concentrations were increased up to a factor four times that of the standard product (BCF-92), the light output increased by 20%. Increasing the fiber diameter from 1.0 to 1.2 mm resulted in a 27% increase of the light yield. Our collaborators at the Institute of Nuclear Research, Moscow have made polystyrene waveshifter fibers with double cladding of progressively lower refractive index, which give rise to larger light yield due to the higher numerical aperture in K-27 doped fibers. Some new fibers made with the G-2 fluor and using the multiclad technology may also be tested.

It is known that the higher concentration of fluor used in the tests decreases the light attenuation length in the fiber from approximately 3.10 to 2.17 m [2], so that the longitudinal uniformity of the calorimeter response is degraded. Simulations show that an "effective attenuation length" of 2 m or more will produce a constant term in the energy resolution of less than 0.5%.

We measured the longitudinal nonuniformity of the towers due to fiber absorption by injecting light from the N₂ laser onto a scintillator tile edge. The tile was part of a "pseudo-module" consisting of fibers and a few scintillator plates built to study uniformity issues. The tile was moved to different depths in the module, and the total charge delivered by our photomultipliers was plotted. Figure 9.2 shows the result of the mapping of longitudinal response with minimum and maximum concentrations of G-2 fluor.

Even with the highest fluor concentrations tested, the module longitudinal uniformity was found to be acceptable. The current choice for WLS fibers is the 1.2-mm diameter fiber

Shish-Kebab EMCaI - Longitudinal scan

Figure 9.2: Longitudinal scan of a calorimeter module, showing the nonuniformity due to waveshifter fiber attenuation. The fiber type "BCF-99-29b" is identical to "BCF-92" but with a factor of four greater fluor concentration.

with the highest concentration of G-2 fluor we have tested. This will result in a signal of approximately 2000 photoelectrons/GeV.

Scintillator-edge Preparation

The same "pseudo module" apparatus was used to study the optical isolation between tiles in a module and transverse uniformity within a tile. The choice has not yet been made among different options for edge preparation of the tiles. Of interest are the overall light yield, uniformity of response within a tile, and the optical isolation between towers within the same module. Monte Carlo simulations have shown that without optical isolation, the calorimeter occupancy would increase by about 50% in a single HIJET environment [4]. Assuming the same single HIJET multiplicities, the performance is degraded relative to that of an isolated shower; for example, the position resolution will deteriorate by about 20% and the energy resolution will degrade by at least 1%. Therefore, it was decided that towers within a module should be isolated optically.

For mechanical reasons it is preferable to prepare the tile edges to accomplish optical isolation, and to achieve the required uniformity by applying a permanent coating. The coating procedure can leave one corner of each tile exposed so that UV light can be introduced from the monitoring system. Among the coatings tested are white reflective paints and vacuum aluminum deposition.

Figure 9.3 shows the result of a scan accross one face of a scintillator, along a line midway between fiber rows. The tile edges were vacuum aluminized and the measured non-uniformity is less than 2% except at the edges where measurements are still incomplete. Currently the edge aluminization looks most favorable, and ways to further improve the edge response are being investigated. Among the possibilities are improvements to the mirror quality, variations on the fiber pattern near the tile edge, and suppressing response at the center of the tile with a mask pattern on the reflecting paper.

Longitudinal-segmentation Schemes

We decided not to pursue the idea of longitudinal segmentation of the calorimeter modules. Both simulations and the 1992 AGS beam test results showed that longitudinal segmentation of our lead-scintillator modules could be achieved using different decay time constants for the first six X_0 of the calorimeter modules [6]. This "fast-slow" segmentation scheme improved e/π separation by about a factor of two. However, it would require more complicated and costly electronics to take into account the spread in arrival times for particles of interest in PHENIX. Although this technique may be suitable for a high energy application, in which arrival times are equal for particles of interest, it is believed that the modest benefits for RHIC would not justify the additional cost.

9.3.3 Calibration System

Extensive use was made of the CDR calibration scheme (illustrated in Fig. 9.6 and described in Section 9.2.3. of the CDR), but the basic design continues to be improved. A stable

Figure 9.3: Transverse scan of a scintillator tile within a calorimeter module. The edge of the scintillator tile was prepared by vacuum-aluminum deposition.

fraction of UV light (337 nm) from a pulsed laser is split through an integrating sphere, going simultaneously to a biplanar phototube and to the distribution system on the calorimeter. When normalized by the reference (biplanar tube) signal, the calibration signal in a typical channel has a Gaussian distribution of $\sigma \sim 1\%$, depending on laser intensity. A 1:37 beam splitter was constructed using 200 μm fibers. The light output of the individual fibers spreads by only a factor of two, instead of by a factor of 4.2 for a commercial splitter (C Technologies).

Finally, the production techniques of leaky fibers were evaluated for light yield and reproducibility. These leaky fibers will distribute the laser light to the four towers throughout the entire depth of the module. With current technology the amount of light that leaks out of the fibers decreases exponentially with distance (depth). This approximation to the signal from electromagnetic showers is already superior to uniform light deposition.

The measured light transmission and conversion efficiencies of individual system components are listed in Table 9.2. This monitoring system has been tested in the laboratory [3], then transferred to the B2 beamline at the AGS. Except for some problems with the mechanical stability of the elements, the monitoring system performed well [1]; therefore, it is likely to be the final choice.

The goal is to be able to predict the calibration (GeV/channel) of each tower to within 10% prior to data taking. Calibration of all modules using a testbeam setup would be a major undertaking, so the possibility of using cosmic ray muons was investigated. Since the

Table 9.2: Efficiency of the Laser Light Distribution System

Item	Efficiency
Integrating sphere	$2.5 \times 10^{-6}/\text{port}$
Splitter (total)	0.28
Splitter/module connection	≈ 0.7
Leaky fiber	≈ 0.7
Light to photoelectrons	$1.7 - 2.4 \times 10^{-4}$

cosmic-ray muons have a peak width of about $\sigma \leq 13\%$, it is likely that for trigger purposes they could be used to provide a 10% initial calibration of all channels [1].

Of direct interest is the typical variation between towers after correction for the measured gain and relative quantum efficiency of the phototubes. Since each tower consists of very many statistically independent optical joints (36 joints/tile and 66 tiles/tower), it is likely that actual differences between scintillators will be the largest contributor to variations between towers. The variation of the light yield at the end of the tower was measured from tile to tile by scanning the edge of a module with a UV fiber. A variation of $\sigma = 3.7\%$ was found and attributed to the combined effects of UV light injection variation into the tile and light collection variation in the module.

9.3.4 Photomultipliers and High Voltage Supply

Several twelve-stage, 3-cm diameter FEU115M phototubes were carefully tested at BNL and by a group at CEBAF [5]. All 100 tubes delivered in the original shipment were tested at IHEP. All tubes were still functional but one, which seems to have been mishandled during shipping. Their overall performance is satisfactory [1]. Noise is low, gains are several times 10^6 , linearity is better than 2% up to 40 mA peak current. We specified a quantum efficiency of greater than 12.5% at 500 nm to the manufacturer and have measured the average to be 16%. Figure 9.4 shows the pulse height distribution of a typical tube. The peak at five channels above pedestal corresponds to a single photoelectron. Cathode response uniformity within a 1.1 cm square area is better than 10% as required by our specification to the manufacturer. This area is still larger than that of the 7-mm diameter of the WLS fiber bundle coupled to the phototube.

As part of the FY94 R&D effort, we decided to design and build at least two test boxes, one of them deployed at IHEP, the other at BNL to facilitate quality control and preselection of all 16,000 phototubes. The results of these measurements will be the primary input to a database, which will be used to set the gains and monitor the long term performance of individual channels.

Figure 9.4: Pulse height spectrum from an FEU115M phototube developed for PHENIX. The single photoelectron peak is clearly visible.

High Voltage Distribution

We are studying two options for high voltage distribution and control. In the first scheme the same high voltage would be deployed to a group of channels, and the HV in each channel would be adjusted with motor driven potentiometers. The range of adjustment is 300–400 V under computer control. Similar systems have been used by the GAMS experiment at IHEP, by CDF at Fermilab, and are now made commercially. The second option would use a 5 kHz generator and a flyback transformer; high voltage in the individual channels could be adjusted in the entire range (including setting them to 0) via a reference voltage. All channels could be addressed and set individually by computer. Members of the IHEP group have experience with the operation of an 800 channel system of similar design used in a neutrino experiment. We are planning to build a prototype of both systems and test them February 1994 at the AGS.

Figure 9.5: Compilation of available data on timing resolution of the FEU 115M phototube and of a full calorimeter tower equipped with it (solid squares). Note that the vertical scale is full width half maximum.

Timing discriminators

Detailed simulations have shown that the best strategy for timing is to derive the timing signal from the channel with maximum energy in a cluster (rather than forming some weighted average of all channels in a cluster). This strategy has also been confirmed by laboratory tests, and it simplifies the use of timing information at the trigger level.

We have compared the performance of leading edge and constant fraction (CFD) discriminators. CFDs offer some advantages at very low pulse heights because there is automatic correction for slewing; on the other hand, CFDs are more expensive and more prone to misadjustment. We found that for pulses with small amplitude A the slewing curve is steep ($1 \text{ ns}/50 \text{ MeV}$), but more or less proportional to $1/A$ and in principle it can be corrected if necessary with the flash ADC signal. We think that a leading edge discriminator with a threshold corresponding to 50 MeV or less will provide adequate timing in the trigger even without an online slewing correction: with 1000 photoelectrons the measured full width half maximum of the timing is 400 ps [1]. Figure 9.5 shows the measured timing resolution for

Table 9.3: Cost Estimate and Schedule for the PHENIX PbSc Calorimeter

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
5.2.9	PbSc (Shish-kebab) EMCal	\$7,755,114.16																							
5.2.9.1	PbSc EMCal Mechanical Design	\$585,002.16																							
5.2.9.2	PbSc EMCal Fabrication and Com	\$6,918,112.00																							
5.2.9.3	PbSc EMCal Assembly	\$144,000.00																							
5.2.9.4	PbSc EMCal Commissioning	\$108,000.00																							

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
5.3.9	Cal(PbSc) FEE	\$613,995.20																							
5.3.9.1	Design Cal(PbSc) FEE	\$91,056.00																							
5.3.9.2	Prototypes of Cal(PbSc) FEE	\$76,184.00																							
5.3.9.3	Fabrication of Cal(PbSc) FEE	\$419,431.20																							
5.3.9.4	Assembly of Cal(PbSc) FEE	\$0.00																							
5.3.9.5	Testing of Cal(PbSc) FEE	\$27,324.00																							

cosmic-ray muons in one of the calorimeter modules read out by an FEU115M PMT. Also shown for comparison are all available data from different laboratories on the resolution obtained with laser pulses and with fast LED.

9.3.5 Cost and Schedule

The cost estimate and schedule for the lead-scintillator calorimeter is presented in Table 9.3. In this table the PbSc mechanical (WBS 5.2.9) and front-end electronics (WBS 5.3.9) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

9.4 Lead-glass Calorimeter

9.4.1 Monitoring System Design Criteria

The monitoring system for the PbGl photon spectrometer is currently being implemented for use in the WA98 experiment at CERN. The design criteria for the monitoring system are:

1. It must provide gain control for the 10,000 phototubes with better than 1% accuracy.
2. It should allow a timing calibration of better than 300 ps.
3. It should be of the same modular structure as the mechanical design of the spectrometer (supermodules of 24 blocks) to allow independent transportation of the supermodules with the calibration system kept intact.

4. It must be reliable.
5. It should need minimal maintenance effort.
6. It should allow linearity measurements.

To meet these requirements the PbGl monitoring system will be based on pulsed LED. The LED light output will be monitored by photodiodes incorporated with each supermodule.

9.4.2 LED for Monitoring System

Ideally, the monitoring light pulses should be as similar as possible to Cherenkov light, that is, fast blue pulses. However, this is problematic because the blue LED available have both low light output and slow response. On the other hand, the existing yellow LED have very high light output and fast response. Therefore, the PbGl monitoring system will use both the blue and yellow LED in order to provide pulses of alternatively the appropriate wavelength or pulse shape.

The dominant effect concerning changes in the LED light output is the temperature dependence. It is known that LED output intensity has an exponential dependence on the temperature: $I(T) = I(25^{\circ}\text{C}) \cdot \exp k(T - 25^{\circ}\text{C})$. Typical temperature coefficients are approximately 1%/C°.

9.4.3 Pulse Generators

The effective response time of the blue LED has been studied with an avalanche pulser with an intrinsic response time of a few ns. LED pulse rise times of the order of 25-30 ns have been achieved. However, the total pulse width is still larger than 400 ns and the intensity of the LED is very low, which makes this mode of operation unusable. For the yellow LED, rise times of 8-10 ns and total pulse widths of 30-40 ns have been obtained with satisfactory intensity using a similar avalanche pulser.

For the above mentioned reasons the PbGl monitoring system will consist of a system of three LEDs each with a different pulse generator:

1. An avalanche pulser (AP) for use with a yellow LED, which will provide fast pulses with the proper time response.
2. A slow pulser (SP) for use with a blue LED, which will provide light pulses with the most appropriate wavelength.
3. A fast pulser (FP) with variable intensity for use with a yellow LED, which will allow linearity measurements.

Table 9.4: Cost Estimate and Schedule for the PHENIX PbGl Calorimeter

WBS	Description	Total Cost	Funct	1994					1995					1996					1997					1998				
				Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4				
5.2.10	Lead-Glass EMCaI	\$5,662,243.20																										
5.2.10.1	PbGl EMCaI Mechanical Design	\$20,240.00																										
5.2.10.2	Prototyping	\$19,200.00																										
5.2.10.3	PbGl EMCaI Fabrication and Com	\$5,612,803.20																										

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
5.3.10	Cal(PbGl) FEE	\$677,618.40																							
5.3.10.1	Design Cal(PbGl) FEE	\$150,576.00																							
5.3.10.2	Prototypes of Cal(PbGl) FEE	\$118,233.60																							
5.3.10.3	Fabrication of Cal(PbGl) FEE	\$385,264.80																							
5.3.10.4	Assembly of Cal(PbGl) FEE	\$0.00																							
5.3.10.5	Testing of Cal(PbGl) FEE	\$23,544.00																							

9.4.4 Light Distribution System

The light distribution to the phototubes is achieved very simply. The LEDs are mounted inside a plastic cover sitting at the front side of the supermodules. The LED light shines onto the inner surface of this cover, which is coated with a highly reflective, difusing paint. The light then enters through holes in the plastic front face of the modules into the lead glass and is transmitted to the phototube. The intensity variations at the phototubes, due to the different positions in the supermodule, are reduced by varying the hole size with position.

The three LEDs will sit in thin metal tubes for shielding to avoid annoying radiation from the pulses. One separate cable (RG174) for each LED will exit the tubes and the plastic cover.

9.4.5 Reference Detector

A PIN photodiode will be mounted on the plastic cover. The diode will look through a hole in the plastic and directly view the LED light. This will allow us to monitor the long term variations of the LED intensity.

9.4.6 Cost and Schedule

The cost estimate and schedule for the lead-glass calorimeter is presented in Table 9.4. In this table the PbSc mechanical (WBS 5.2.10) and front-end electornics (WBS 5.3.10) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

Bibliography

- [1] PHENIX DC Meeting and Workshop, Aug. 29 - Sep. 2 1993, PHENIX Note 94
- [2] C. Hurlbut, Bicron Corp., private communication.
- [3] E. Kistenev, C. L. Woody: Laser Light Distribution for PbSc EMCal, PHENIX Note 86, June 1, 1993
- [4] G. David: Effects of Missing Optical Isolation Between EMCal Towers, PHENIX Note 96, Sep. 3, 1993
- [5] S. Majewski: Tests of two FEU115M phototubes, CEBAF Internal Note, June 12, 1993
- [6] Proceedings of the III. International Conference on Calorimetry for High Energy Physics, Sep. 29 - Oct. 2, 1992; presented by S. White (to be published)

Chapter 10

Muon Arm

10.1 Overview

The purpose of the PHENIX muon arm is to facilitate the study of vector mesons decaying into dimuon pairs, to allow the study of the Drell-Yan process and charm production via $D\bar{D} \rightarrow \mu^+ + \mu^- + X$, and to provide the muon detection in $D\bar{D} \rightarrow \mu + e + X$. To meet these objectives, the muon arm must both track and identify muons, as well as provide good rejection of pions and kaons. As discussed in Chapter 3, the mass resolution for the PHENIX muon arm is sufficient to resolve the ϕ from the $\rho + \omega$, the ψ' from the J/ψ , and the $\Upsilon(2S+3S)$ from the $\Upsilon(1S)$.

Figure 10.1 shows an updated elevation view of the PHENIX muon arm, including the current layout of the muon identification system, and Fig. 10.2 illustrates details of the absorber region. The muon identifier comprises six walls of steel absorber interleaved with layers of limited streamer tubes of the Iarocci type. The streamer tubes are instrumented with sets of pickup pads to give low resolution tracking capability and provide signals corresponding to a useful geometrical segmentation to the first and second level muon triggers. The first absorber wall is the steel end-plate flux return of the muon magnet. It is followed first by two 10-cm thick steel layers, then by three more, which are 20-cm thick. More details of the design are given below.

10.2 Changes

The major changes in the muon arm since the CDR include:

1. An improved design of the tracking chambers.
2. An improved plan for the muon identification segmentation and number of identification layers.
3. A better understanding of the effects of neutrons on the pattern recognition in the tracking arm.

Figure 10.1: Elevation view of PHENIX muon arm in the major facility hall.

4. A small modification to the layout of the specialized muon arm absorbers.
5. An improved understanding of the triggering capabilities of the muon arm using the muon ID layers.

As shown in Table 10.1, the design presented in this update results in an increased number of channels in the muon arm. There is also additional progress. For example, the results from a testbeam run at BNL provided a better understanding of muon identification performance, along with the opportunity to compare measured versus simulated performance. The resulting R&D activity is described below in Section 10.4.4. Construction of a prototype drift chamber, whose design is appropriate for the tracking chambers, is now in progress.

Finally, since the time of the CDR, a major technical review of the muon subsystem was held at BNL (6/93). Summarized here is any technical review information, which differs from the CDR. Further details on the muon arm technical review are given in Ref. [1].

Table 10.1: Summary of Channel Counts for Muon Arm

Detector	Layer No.	No. of Channels
μ T station	1	4416
	2	5520
	3	9744
	Total	19,680
μ ID plane	1	7632
	2	3816
	3-6 (each)	1728
	Total	18360

10.3 Acceptance and Coverage

The CDR design of the tracking chambers represented just the beginning of the full technical design process. For example, "frameless" chambers were assumed in the CDR analysis of the acceptance and coverage of the muon arm. Described below (in Section 10.4.3) is a more realistic, but still preliminary, chamber design. The acceptance was studied using chamber frame widths (as measured in the r - ϕ view) of 5, 7.5, and 10 cm for stations #1, #2 and #3, respectively. The analysis required that the muons from the vector mesons, generated with the same kinematic distributions as described in the CDR, penetrated to the third layer of the current design of the muon identification system and did not hit any frames in the tracking system. The fraction of events where both muons survive the cuts are 16% for the ϕ , 47% for the J/ψ , and 66% for the Υ . Drell-Yan and $D\bar{D} \rightarrow \mu^+ + \mu^- + X$ rates are reduced by about the same fraction as for a resonance of the same mass. The acceptance losses have a rapidity dependence. Increased losses occur at higher rapidity, where the width of the chamber frames is a larger fraction of the total area inscribed by the solid angle. For now the above fractions should be applied to all event rates and acceptances quoted in the CDR versions of this section and Chapter 3. However, as discussed below in Section 10.4.3, much more work will be done to optimize the chamber design. It is likely that a significant fraction of the above losses are recoverable.

10.4 Layout and Mechanical

10.4.1 Interface to Magnet

The design for the specialized absorbers, which are part of the central magnet, was modified slightly since the CDR. The changes are described below in Section 10.4.2. The design changes result partly from the relocation of the beam-beam counters, but also from a precautionary measure in response to higher estimates of hadron shower leakage into the muon arm.

The design of the muon tracking chambers was also studied in more detail since the CDR. It is now realized that the chambers must be designed to mount directly to the muon magnet lampshade. The new design is described below in Section 10.4.3.

10.4.2 Specialized Absorber Sections

One change in the specialized absorbers is the specific inclusion of a reentrant tip on the muon magnet piston. In the CDR design, a copper absorber was located between the main magnet pole-tip and the station #1 chambers. The absorber held the beam-beam counters. Since the beam-beam counters have been relocated, a similarly shaped reentrant tip has been added to the front of the muon magnet piston. This arrangement serves to protect the station #1 chambers from most of the electrons arising from gamma-ray conversion on the front of the piston. The layout is shown in Fig 10.2.

Figure 10.2: View of muon arm showing the specialized absorber layout.

A more significant change is the addition of a neutron absorber. As described in Section 10.6.1 below, new calculations of shower leakage into the muon arm indicate that earlier estimates were low. The increase in the number of expected neutrons is nearly a factor of three greater than the CDR estimate. The design is now changed to allow the possibility, as shown in Fig. 10.2, of incorporating a borated polyethylene wall as a neutron shield, which would be positioned between the central magnet pole tip and the lead curtain. The neutron absorber may not be included because (as is discussed in more detail below) the higher neutron fluxes now being estimated may not present difficulties for the tracking chamber pattern recognition and track finding algorithms. However, because the muon energy threshold would be increased by the incorporation of either the neutron shield or the lead curtain, the acceptance for ϕ events would be diminished significantly. Therefore, the plan is to incorporate the required mounting capabilities (mostly bolt holes) into the system, but not necessarily to install the absorbers themselves, unless it is demonstrated that they are actually needed for background reduction in the tracking chambers.

10.4.3 Muon Arm Tracking Chambers

After the CDR, a detailed study of the tracking chamber design was begun. A mounting scheme is envisioned, which uses kinematic mounts attached directly to the muon magnet lampshade. That would isolate the chambers from the motion of the magnet and reduce acceptance losses due to an external support structure.

An active surveying system for monitoring chamber positions in real time is also planned. It is presently conceived to consist of a series of light emitting diodes mounted on station #1, together with corresponding lenses on station #2, which focus the light on a series of quadrant diodes on station #3. The locations of these devices will be chosen to measure the motion of individual stations as well as the motion of stations relative to one another. A UV-laser may be used to illuminate the active volume of the chambers, in order to provide reference trajectories.

The chambers are constructed of stacked frames of wires. The frames are made of G-10 laminated to aluminum frames. Wires will be wound onto each individual frame, then the frames will be stacked to complete the chamber. A cross-sectional view of a stack is shown in Fig. 10.3. Cathode wire planes alternate with anode-field planes, resulting in nine cathode planes and six anode field planes per station octant. A shield of 25- μ m thick aluminized mylar is used to electrostatically separate the $U - U'$ from the $Y - Y'$ planes and to separate the $Y - Y'$ from the $V - V'$ planes. Aluminized mylar and acylar sheets provide both the outer ground planes and the gas barriers. In the present design, the chamber frame widths are 5.0, 7.5, and 10 cm for stations #1, #2, and #3, respectively. Further finite-element analysis and materials studies will be done to reduce the widths of the frames in order to minimize acceptance losses. For example, it may be possible to taper the chamber frame thickness near the piston, which would greatly improve the ϕ acceptance. The chamber frames include the chamber mounting fixtures, survey reference points, and elements of the active surveying system. Measurement accuracies of 250 μ m per plane were achieved previously with large chambers of this type [2].

Since the CDR, there were also some further studies of the wire spacing and geometry,

Figure 10.3: Cross-sectional view of the wire planes in one octant of a muon arm tracking station.

which included modeling the electrostatics with GARFIELD [3]. Figure 10.4 shows the resulting geometry. In the chambers of station #1, the cells will be 0.5-cm square, and in stations #2 and #3, the cells will be 1-cm square. That cell geometry was chosen to have the smallest channel count consistent with good pattern recognition. As part of those studies, the impact of neutrons on the chamber occupancy was examined and it was found that the occupancy would be increased by less than 2%. That increase in occupancy is small in comparison to the 10% occupancy due to charged particle tracks in station #1. Therefore, neutrons should not cause problems for pattern recognition in the muon tracking system.

Since the time of the CDR, a $30 \times 30 \text{ cm}^2$ prototype drift chamber was designed using the proposed wire geometry, and it is currently being assembled. The prototype will be used to study resolution issues and gas mixtures. The first such results should be available by late fall. At the same time, an ongoing R&D activity is investigating the use of very thin cathode strip chambers as an alternative to drift chambers. The advantages would be better tracking resolution and a faster readout which could, in turn, be useful in the LVL-2 trigger.

Figure 10.4: The wire geometry of the drift chambers.

10.4.4 Muon Identifier

As stated in Section 10.1, the muon identifier now consists of six walls of steel absorber interleaved with layers of limited streamer tubes of the Iarocci type. The layout is shown in Fig. 10.1. The first absorber wall is the 30-cm thick steel end-plate flux return of the muon magnet. It is followed by two 10-cm thick steel layers, then by three 20-cm thick steel layers. Fabrication and installation of the muon identifier steel is discussed in Chapter 4, since it is integrated with the magnet components for design, structural, and scheduling reasons. A segmentation of that type has an approximately 2.2 GeV low energy muon threshold and a normal incident muon would lose approximately 3.3 GeV passing through the entire stack. In the CDR the absorber slabs were 1-m thick concrete walls. That provided a longitudinal segmentation of approximately 2.5 radiation lengths per layer. Because of safety reasons, a finer granularity was not possible using concrete. The finer granularity is needed in order to reach our required μ/π sensitivity.

The streamer tubes have a resistive graphite coating on the inner surface that serves as the cathode. They are instrumented with sets of external pickup pads to give low resolution tracking capability and provide signals corresponding to a useful geometrical segmentation for the first and second level muon triggers. A cross-sectional view of these chambers showing the locations of the pickup pads is given in Fig 10.5, and the proposed pad segmentation

for the LVL-1 and LVL-2 trigger is shown in Fig. 10.6. In the first gap, there are two sets of LVL-1 pads and two sets of LVL-2 pads. How that can be realized is shown in Fig. 10.7. The first gap is to insure that this road starting gap operates at nearly 100% efficiency. The second gap has the same pad arrangement as the first but those pads are offset by one-half pad in θ and in ϕ with respect to the ones in the first gap. The second gap also contains one layer of LVL-2 pads. The remaining gaps (gaps 3 through 6) contain one set of LVL-1 pads each. The sizes and numbers of the pads have been chosen to give a good match to the road widths of muons passing through the muon identifier. With the proposed pad segmentation, it will be possible to trigger on tracks that have the trajectory of muons originating from the interaction region (where $dr/dz > 0$), while discriminating against beam-gas (where $dr/dz \approx 0$).

Present signal-to-noise studies of the muon arm indicate that a fraction of the charged pions are misidentified as muons. A μ/π ratio of 3×10^{-2} at 1 GeV/c and of less than 1×10^{-2} at 2 GeV/c must be attained in order to insure that pairs of pions misidentified as muons constitute less than 10% of the combinatorial pairs for the muon arm. After the CDR, R&D work took place on the muon identifier with the goal of ensuring that the required performance will be attained. A prototype muon identifier was prepared for the 1993 BNL-AGS testbeam run on the B2 line. The purpose of that prototype was to build a section of the muon identifier that was deliberately over instrumented and overly segmented (both in depth and in the transverse direction) compared to what should be required for PHENIX. Analysis of test beam results could then be made, using less and less of the information available, until the limits of needed performance for PHENIX are reached. Margins for safe operation can be established in this way.

The prototype muon identifier consisted of limited streamer tubes, read out by x -strips on one side of the tube and y -strips on the other, interleaved with layers of steel absorber. Each limited streamer tube had eight $9 \times 9 \times 1500$ mm³ long active volumes with an anode wire at the center of each. Those sections were separated one from another by 1-mm thick plastic septa. The cathode was coated with carbon, permitting the use of external pickup strips for signal detection. Some 150 streamer tubes were used. The streamer tubes were packaged in frames that held the readout strips in position. The readout strips were four centimeters wide in both directions and extended the length of the streamer tube array. The streamer tube units were mounted between steel plates on a large support stand. Ten 0.5-in. thick steel plates were placed before each layer of streamer tubes. The entire array, which included more than 125 cm of steel, corresponds to 7.5 hadronic absorption lengths. The data from that testbeam run is in the process of being analyzed and will be compared to a GEANT simulation of the prototype. The information from that study will be incorporated in the final design of the PHENIX muon identifier.

A further R&D study is under way to investigate the applicability of resistive plate chambers run in a proportional mode. Possible advantages include a shorter integration time (as short as about 20 ns as compared to about 110 ns for limited streamer tubes). This narrow time window would reject slow neutrons that interact in the muon identifier.

Figure 10.5: A cross section of a muon identifier limited streamer tube with its pickup readout pads.

10.5 Electronics

10.5.1 Front-end

Since the time of the CDR, additional investigation of the front-end readout has identified an attractive option for the drift time measurement TDC. As stated in the CDR, the University of Pennsylvania ASD chip will serve as the preamplifier-shaper. Instead of the LeCroy MTD132A ASIC TDC, a TDC developed by Y. Arai of KEK for SDC tracking is being considered. This multihit TDC ASIC, referred to as TMC 4004, uses a shift register approach that can be adapted readily to accommodate the RHIC 112-ns timing. The TMC 4004 chip yields a defined output that can be synchronized with the RHIC RF, which naturally fits into the pipelined architecture of the readout without further "framing" of the data. Use of the LeCroy TDC would require determining which RF time bucket a hit belongs to and using that information to choose the appropriate DMU location for storage. That rapidly becomes awkward with high hit rates. A second advantage of the TMC 4004 is that it helps to minimize deadtime because it is not read out in common stop mode.

The front-end readout for the muon identifier remains essentially as described in the CDR. The only changes since the CDR is that a timing circuit has been added, we know that pulse height information is not needed, and the electronics will have 16 channels per chip set. In order to minimize R&D efforts, the timing circuit will be based as much as possible on the corresponding circuit for the EMCal.

10.5.2 Trigger

After the CDR, it was realized that beam-gas interactions at RHIC occur at a higher rate than the collision rate at the interaction vertex. This led to new simulations of a possible implementation of a LVL-1 trigger processor that uses AND's and OR's to search for roads in the r - ϕ space of the muon identifier. The output of this LVL-1 trigger would be the number

Figure 10.6: The proposed segmentation of the pickup pads on the limited streamer tubes in the muon identifier. (a) The arrangement of the pads used by the road-finding algorithm of the LVL-1 trigger. (b) The 1° ϕ -measuring pads used by the LVL-2 trigger to project from the muon identifier into the tracking chambers. In both cases, the inside radius is at $\theta = 10^\circ$ and the outside radius corresponds to $\theta = 35^\circ$.

Figure 10.7: The arrangement of the pick-up pads and the muon identification streamer tube panels in the first muon identifier gap.

and “coordinates” of each muon track candidate along with the depth of penetration of the candidate track. It would work by using the hits from the LVL-1 pads and building roads through the muon identifier. The segmentation of the LVL-1 pads is designed to match the growth of road widths as a muon traverses the identifier. When finished, the LVL-1 would have a list of tracks, their coordinates, and their penetration depth. A cut requiring two or more tracks which penetrated to at least a depth of the third identification readout plane (a “pairs” cut) could be applied. For Au + Au events, such a “pairs” cut has a rejection factor of five.

Much as described in the CDR, the LVL-2 trigger would use the output of the LVL-1 trigger in conjunction with hits in the LVL-2 pads to provide masks for the digitized outputs of the $Y - Y'$ planes of the station #3 chambers. The process would be repeated working towards the station #2 and #1 chambers. If hits are found in the appropriate places in these stations, then the LVL-2 trigger would count it as a muon candidate. The complete output of the LVL-2 trigger would be the number and coordinates of muon track candidates. A code to simulate the trigger and which runs on a digital signal processor has been written and tested with good results. Based upon PISA simulations of background and J/ψ events, excellent muon track finding efficiency ($> 95\%$) and good rejection of background (80%) are obtained.

10.6 Performance and Backgrounds

More simulations of the performance of the muon detector have been carried out, primarily at ORNL and Vanderbilt University, since the completion of the CDR. These simulations were carried out with the integration of the complete muon arm geometry and analysis software incorporated into PISA. As part of these simulations, the FLUKA high energy particle interaction code was incorporated into PISA. The main result of those simulations have been a better understanding of the neutron, electron, and photon particle leakage into the tracking system. In addition, detailed studies of track reconstruction efficiency and resolution have been performed.

10.6.1 Absorber

As stated above, simulations of the shower leakage into the muon arm have been conducted since the CDR. The new simulations give better estimates of the numbers of neutrons, electrons, and photons that enter the muon arm. The results are summarized in Table 10.2, which is an update of similar results presented in the CDR. The simulations used the GEANT 3.15 and FLUKA packages, with tracking cuts of 1 MeV for all particles. Events from a UA1 event generator, with a hard p_T component, were used [π^+ , π^0 , π^- (1 : 1 : 1) and 6% charged kaons, 2 events, 10,240 particles/event]. The primary track acceptance cuts were $p > 0.0$ GeV/c and $0.4^\circ < \theta < 70.0^\circ$.

The results in Table 10.2 are different from the corresponding CDR numbers primarily because the neutron flux into the muon arm is now estimated to be nearly a factor of three higher than before. Consequently, the neutron detection efficiency of the proposed muon tracking chambers was investigated. Calculations indicate that an upper limit to the neutron efficiency should be about 0.5%. There is, in fact, evidence to indicate that it could be much lower. This investigation is continuing. However, even if the upper limit were the actual efficiency, the total occupancy of the chambers would only increase by about 2%, compared to the nearly 10% occupancy due to charged particles. Therefore, the track finding and pattern recognition algorithms should not be adversely affected. Furthermore, the effect on the muon LVL-1 trigger should be minimal; however, that is currently under further investigation.

10.6.2 Tracking and Spectrometer

There has been a significant improvement in the characterization of the tracking and mass reconstruction capabilities of the spectrometer. Results from double density central Au + Au collisions simulated with PISA give tracking efficiencies of 91% for the ϕ and 93% for the J/ψ and Υ . That demonstrates the robustness of the tracking algorithms even for very high occupancy events.

10.7 Cost and Schedule

The cost estimate and schedule for the muon tracking chambers is presented in Table 10.3. In this table the μT mechanical (WBS 5.2.12) and front-end electronics (WBS 5.3.12) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

The cost estimate and schedule for the muon identifier is presented in Table 10.4. In this table the μID mechanical (WBS 5.2.13) and front-end electronics (WBS 5.3.13) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

Table 10.2: Comparison of Shower Leakage into the Muon Arm With and Without Shields

	Station #1			Station #2			Station #3		
Leakage Particle Type	<0.1 GeV	0.1-1.0 GeV	>1.0 GeV	<0.1 GeV	0.1-1.0 GeV	>1.0 GeV	<0.1 GeV	0.1-1.0 GeV	>1.0 GeV
Without Neutron and Lead shields									
γ	1266.0	12.0	0.0	1227.0	14.5	0.0	1172.5	11.5	0.0
e^+e^-	72.0	2.5	0.0	6.5	2.5	0.0	19.5	1.0	0.0
Charge Hadrons	1.5	54.5	4.5	0.5	51.0	5.5	0.0	40.0	4.5
$\mu^+\mu^-$	2.0	4.0	0.5	0.0	11.5	1.5	0.0	10.0	2.0
Neutral Hadrons	2995.0	773.5	20.0	4137.5	968.0	29.0	4760.0	918.5	27.5
Neutrons only	2995.0	769.0	19.0	4137.5	963.5	28.5	4760.0	917.0	27.0
With Neutron shield but without Lead shield									
γ	1689.0	12.0	0.0	1399.0	13.0	0.0	1239.5	10.5	0.0
e^+e^-	72.0	2.0	0.0	5.5	5.0	0.0	19.5	3.0	0.0
Charge Hadrons	2.5	42.5	5.5	1.5	48.0	9.5	1.5	34.5	8.0
$\mu^+\mu^-$	1.0	3.5	0.5	0.0	7.0	0.5	0.0	7.0	0.5
Neutral Hadrons	1559.5	556.0	15.0	3011.0	756.5	24.5	3353.5	695.0	22.5
Neutrons Only	1559.0	553.5	13.5	3010.5	753.0	23.0	3353.0	692.5	21.0
Without Neutron shield but with Lead shield									
γ	873.5	9.5	0.0	1046.5	10.0	0.0	964.5	8.0	0.0
e^+e^-	69.0	1.5	0.0	4.0	1.5	0.0	12.5	1.5	0.0
Charge Hadrons	1.5	30.5	4.5	0.5	36.0	4.0	0.0	26.0	4.0
$\mu^+\mu^-$	0.5	4.5	0.0	0.5	4.5	0.0	0.5	6.0	0.0
Neutral Hadrons	2482.0	654.0	12.5	3846.5	839.0	19.0	4337.0	781.0	16.5
Neutrons only	2482.0	652.0	12.5	3846.5	837.5	19.0	4337.0	779.0	16.5
With Neutron and Lead shields									
γ	871.5	6.0	0.0	1019.0	9.0	0.0	900.5	7.5	0.0
e^+e^-	70.0	0.5	0.0	2.0	2.0	0.0	9.5	1.5	0.0
Charge Hadrons	0.5	32.0	3.0	0.0	34.5	5.5	0.0	24.5	4.5
$\mu^+\mu^-$	1.0	4.0	1.0	0.5	5.5	1.0	0.5	6.5	1.0
Neutral Hadrons	1586.5	552.0	7.0	3147.0	790.0	13.5	3438.0	735.5	13.0
Neutrons only	1586.5	549.0	6.5	3147.0	784.5	12.5	3438.0	730.5	13.0

Table 10.3: Cost Estimate and Schedule for the PHENIX Muon Tracking Chambers

WBS	Description	Total Cost	1994				1995				1996				1997				1998			
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
5.2.12	Muon Tracker	\$3,269,374.16																				
5.2.12.1	Engr. Design Inspection & Administration	\$533,440.00																				
5.2.12.2	Prototyping	\$140,000.00																				
5.2.12.3	Fabrication	\$2,392,654.16																				
5.2.12.4	Chamber Testing	\$69,584.00																				
5.2.12.5	Shipping	\$20,000.00																				
5.2.12.6	Installation	\$113,696.00																				

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998	
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2
5.3.12	Muon Tracker FEE	\$1,048,742.00																					
5.3.12.1	Design Muon Tracker FEE	\$99,320.00																					
5.3.12.2	Prototypes of Muon Tracker FEE	\$71,878.00																					
5.3.12.3	Fabrication of Muon Tracker FEE	\$855,240.00																					
5.3.12.4	Assembly of Muon Tracker FEE	\$0.00																					
5.3.12.5	Testing of Muon Tracker FEE	\$22,304.00																					

Table 10.4: Cost Estimate and Schedule for the PHENIX Muon Identifier

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998	
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2
5.2.13	Muon ID Start	\$1,469,737.84																					
5.2.13.1	Design Muon Identifier	\$82,415.84																					
5.2.13.2	Fabricate MuID	\$1,156,982.00																					
5.2.13.3	Test Muon ID	\$13,980.00																					
5.2.13.4	Assemble Muon ID	\$216,380.00																					

[illegible]

Bibliography

- [1] PHENIX Note PN-84, PHENIX Collaboration, PN-84, June, 1993.
- [2] C. Conta et al., UA2 Collaboration, Nucl. Instr. Meth. **224**, 65-74 (1984).
- [3] R. Veenhof, GARFIELD, a drift chamber simulation program, CERN Program Library entry W5050 (1992).

Chapter 11

On-line System

11.1 Overview

The principle design goal of the PHENIX On-line System is to select and archive events of potential physics interest at the rate given by the available RHIC luminosity. To achieve this goal the system must

1. Have enough bandwidth to permit digitization of all Au + Au interactions at the RHIC design luminosity (approximately 14 kHz peak interaction rate, 7 kHz average),
2. Implement a set of parallel triggers that select events of potential physics interest, and reduce the accepted interaction rate to about 200 Hz, commensurate with transfer to mass storage, and
3. Be extensible in both bandwidth and trigger selectivity, so that upgrades to handle higher interaction rates encountered with collisions between lighter beam species and/or future accelerator luminosity upgrades will be straightforward.

The particle multiplicity per event at RHIC can exceed 5000 for central Au + Au events, which implies that, even in a small aperture, physics-sensitive triggers must examine a large number of tracks. This high particle multiplicity leads to many more than one e^+e^- pair, or $\mu^+\mu^-$ pair, or hadron pair useful for HBT-type analyses, per event. Most of the leptons detected result from Dalitz decay, photon conversion, or weak decays of hadrons, so that most "lepton pairs" are in fact combinatorial pairs. The backgrounds to the physics observables are so large that in many cases the background event rate exceeds the true event rate by two or three orders of magnitude. Nowhere is this problem more evident than for the dielectron measurements. The backgrounds from π^0 Dalitz decay and photon conversions are such that in central Au + Au collisions, roughly 75% of the events have an electron pair, while only one in every 3000 events has a dielectron pair resulting from a vector resonance decay. The muon spectrometer has a lower background and thus finds interesting events at a much lower rate (one in every 5 central Au + Au interactions), but again only a few percent of these are true pairs. The PHENIX on-line system thus must be able to balance the raw triggering and data processing rates from such disparate parallel streams as the dielectron and dimuon ones,

while also passing gracefully from high event rate $p + p$ collisions to high event multiplicity $Au + Au$ collisions.

The desire to measure low mass pairs and low p_T particles in PHENIX makes it difficult to exclude these combinatorial pairs using any straightforward low-level trigger. The need for good statistics and control of final uncertainties requires that PHENIX examine nearly all potential lepton pairs and record a large fraction of them. The triggering algorithms must permit tagging and some examination of all events with lepton pairs for all colliding species and bombarding energies. A flexible, low-deadtime and multilevel triggering system is therefore essential.

The on-line system is fully pipelined, has three levels of triggering (hereafter denoted as LVL-1, LVL-2 and LVL-3), is free of deadtime through LVL-1, and has minimal dead time at LVL-2 and LVL-3, all in order to preserve the high interaction-rate capability of the experiment. Buffering is provided to handle fluctuations in the event rate and in trigger processing time. Pipelining allows the on-line system simultaneously to record data from the detectors, process LVL-1 trigger information from recent bunch crossings, perform LVL-2 and LVL-3 analysis on events from much earlier bunch crossings, and build and archive accepted events. The lower levels of the readout and the LVL-1 triggering system will run in parallel via distributed processing to handle on-line data correction, zero suppression, and calculation of trigger "primitives" to be used by higher levels of analysis. The first level of the data readout and triggering systems is clock-driven, with the bunch-crossing signals from the accelerator providing the pipeline clocking signal. The higher levels of the readout and triggering systems are data-driven, with the results of triggering and processing in one level propagating up to the next level only when the processing on the current event is finished. The only information propagating from higher to lower levels in the readout is for the purpose of suspending processing or inhibiting triggers when buffers at higher levels approach saturation due to extreme fluctuations in event size or trigger processing times.

Significant reduction in the raw event rate for observables such as electron pairs can only be obtained by reconstructing the electron momenta and calculating pair invariant masses. This requires assembling data from many of the detectors in the experiment, for example DC, PC, EMCal, RICH, and TOF. This can only be done in the second or third stage of the trigger. The muon spectrometer can provide a factor of about 5 rejection at LVL-1 simply by counting deep hits in the muon identifier system, and the EMCal can select photons of greater than 2 GeV (or 3 GeV), which occur in 20% (or 2%) of central $Au + Au$ events, respectively, via a simple LVL-1 cluster trigger.

PHENIX has approximately 200K detection elements, many providing both time and pulse-height and some storing multiple hits or requiring numerous samples for detailed waveform analysis. Several of the PHENIX detection subsystems have response functions that imply a memory time spanning as many as twenty beam crossings, enough that probabilities for pileup of events from different crossings in the detector volume are non-negligible at the highest interaction rates. Most of the detectors will be sampled within 2-3 beam crossings (BC) of a given crossing. However, the DC and the TEC have long memories relative to the bunch crossing rate. The readout system for these detectors must have sufficiently deep buffers to store the entire relevant memory of the detectors, together with the associated

Figure 11.1: Data flow architecture of the PHENIX on-line system.

control logic to be able to select those samples associated with earlier crossings.

11.1.1 General On-line Scheme

Figure 11.1 shows the architecture planned for the PHENIX on-line system from a dataflow perspective. This general on-line structure includes:

- **Front-end Electronics (FEE)** — PHENIX uses front-end electronics with minimal intrinsic deadtime, intelligent buffering, and as much signal processing as possible on or

near the physical volume of the detector. Samples of signals from all detector elements will be taken every beam crossing and stored locally while LVL-1 is applied to reduce the volume of data that must be digitized and read out from the detector. Raw signals from the various detector subsystems must first undergo analog processing [AP] in the FEE, with amplification and shaping to extract amplitude and/or time information and to develop inputs for LVL-1, where included. The FEE provides this analog processing and internal buffering of signal samples synchronously with the RHIC beam crossing. It also correlates these samples with particular beam crossings, produces signals for the LVL-1 trigger, accepts events in accordance with LVL-1, digitizes the events thus selected, and outputs event fragments under the control of the data acquisition system.

Figure 11.1 recognizes the need for different methods for handling the data of three types of detector subsystems. Implementation of Type-A subsystems, wherein all data is digitized in real-time, on or near the detector, is straightforward using available fast analog and digitization electronics as long as extremes of high-speed, extended precision, or low power density are not required. Several PHENIX detectors will use this method: BB, DC, RICH, TEC, and μ T. However, considerations of channel density versus power/heat loads make this approach problematic for MVD. Cost and availability concerns arising from precision and/or channel count requirements preclude using this method for others: TOF, PC, EMCal, and μ ID. These cases employ the Type-B strategy of holding the raw data in front-end analog buffers [AB] based on switched-capacitor-type Analog Memory Units (AMUs). This approach decreases digitization bandwidth requirements for these subsystems directly in proportion to the rejection provided by LVL-1. Finally, for detection systems with a long intrinsic analog processing times or which are not involved in triggering we could employ the Type-C approach, with digitization performed only after an event has been accepted by LVL-2. With the detection technologies currently envisioned we do not presently foresee such a need.

The FEE for the various subsystems is designed to appear as uniform as possible to the rest of the on-line system and uses a common approach to time-frame buffering. Details of the front-end electronics for various detector subsystems are discussed in Section 11.3.

- Level-1 Trigger (LVL-1) — This system provides fast filters for discarding empty beam crossings and uninteresting events before their data are fully digitized. Every 112 ns it receives trigger data corresponding to a new beam crossing in parallel from all requisite front ends over a relatively small number of dedicated high-speed paths. It operates in a synchronous pipelined mode with a fixed latency $\Delta t_{LVL-1} \simeq 4 \mu s$. It thus generates decisions every 112 ns, for beam crossings which occurred earlier by Δt_{LVL-1} . The latency has been chosen to permit reasonably sophisticated LVL-1 decisions without requiring prohibitively deep front-end buffers. The LVL-1 decisions are broadcast to the front-end electronics to initiate final digitization and readout. For accepted crossings the LVL-1 results, including tagging, masking and routing bits plus various sums, are also read out as data for input to more extensive filters. LVL-1 is discussed further in Section 11.4.1.

- **Data Acquisition System and Higher-Level Triggers** — For events accepted by LVL-1 the data acquisition system receives the data fragments from the front-ends and all data produced by LVL-1 itself. The DACQ system first buffers them while additional filters are applied in LVL-2 to reduce the accepted event rate further. Subsequent processing includes data compression to conserve readout bandwidth (whether by simple zero-suppression or more elaborate corrections plus parameterizations) by using digital signal processors [DSPs]. Preparations are made for subsequent routing operations by addition of data-block headers. Buffers [B] with smart I/O ports are used to accommodate fluctuations in input rate, hold data pending decisions on its disposition, and synchronize and merge results from different parallel sources.

The Level-2 trigger is designed to provide additional rate reduction through a flexible, programmable environment for applying both subsystem-local processing and global correlations between detector subsystems on a time scale of about 100–1000 μ s. Within each detector subsystem Local Level-2 trigger [LL2] processors can be used to apply detector-specific triggers and/or prepare additional primitives for global tests. Multiple sectors within a given subsystem as well as separate subsystems can be treated in parallel, thereby providing a means for load balancing the system. Because each leg is pipelined with respect to event-ID, different events can be processed concurrently at different points in the system, with synchronization postponed until required by the algorithms running in the Global Level-2 Trigger [GL2]. These algorithms will typically embody filters based on correlations between different detector subsystems. LVL-2 will also be a parallel, pipelined processor so that different sections of it can handle different events concurrently, with final overall synchronization postponed until required by the Event Builder [EB].

Events which pass the LVL-2 cuts are then assembled into complete event records by the Event Builder. The Event Builder includes a set of high-speed DSPs in its output legs. These processors perform final formatting and are the first processors which can examine the entire dataset for a given event. We will utilize this to apply yet another level of triggering, denoted LVL-2.5. The design permits subsequent event transmittal to a third level trigger, LVL-3, which can perform *full* reconstruction for a limited set of events, or subsets of events, if needed. This LVL-3 Trigger runs in high-speed general-purpose processors, as opposed to DSPs, and may also take advantage of the entire event information. The demand for archival storage must be reduced by this point to approximately 20 MB/s. This is described further in Section 11.5. Data for events passing this final level of filtering will be archived on permanent storage media by a set of data loggers in the On-line Computing System (ONCS).

- **On-line Computing System** — The On-line Computing System is the focus for overall control and monitoring of the front-ends, trigger system, data acquisition system, and all ancillary systems. It is responsible for archiving data, monitoring and controlling the data flow thru the system, downloading the more than 2000 embedded processors in the data acquisition system, and configuring the entire on-line system. It must handle calibration and monitoring tasks, providing needed database services for maintaining

histories of calibrations, surveys, run records, detector histories and operating parameters. It must present a uniform interface to users wishing to operate the system or analyse results and must provide needed software tools and code management facilities to enable software development. It has responsibility for monitoring and controlling all "slow" aspects of PHENIX such as high-voltage systems, gas systems, AC and DC power systems, and temperature monitors. Finally it provides general connectivity between the on-line system and the simulation and off-line analysis facilities as well as other external computing resources. An approach based on a standardized realtime workstation, standardized realtime interface cards, and a distributed database and distributed monitoring processors is being developed. This system will enable users to access all parts of PHENIX amenable to computer monitoring and control via a common set of protocols. See Section 11.6 for details.

The baseline data acquisition system will be able to archive data at 20MB/sec. At RHIC design luminosities this can keep up with the 0.5% most central Au + Au collisions and with p + p collisions, although not with p + A and some lighter A + A minimum-bias collisions. Thus in early running, or until RHIC reaches design luminosities, the data acquisition system will be able to collect data with this taping capability and without the LVL-2 trigger. The bandwidth of the entire system, including the implementation of the LVL-2 trigger, will increase to keep up with expected increases in rates.

11.1.2 Event Rates at RHIC

The expected raw interaction rates and data volumes in PHENIX are summarized in Table 11.1. The table is updated from the PHENIX CDR to reflect the present PHENIX Baseline design as well as to incorporate a better estimate of expected event sizes and data derived from each subsystem. It lists projected luminosities, interaction rates, charged particle multiplicities and data rates for many different types of collisions, ranging from p + p to central Au + Au. The table indicates a factor of 300 variation in the raw interaction rate, but a factor of less than 3 variation in the product of (interaction rate) $\times$ (multiplicity) (excluding asymmetric $A_1 + A_2$ ion collisions and the somewhat artificial p + p "central"). Also included in the table is an estimate of the zero-suppressed data rates that would be necessary to read out every interaction for the given species combinations. (The number of bytes per track is developed below as shown in Table 11.5 in section 11.3.1.)

All of the minimum bias data rates exceed the planned Day-1 data archiving capacity of approximately 20 MB/s by factors of 20 to 50. A reduction of the raw minimum bias data rate via selective triggering is thus needed. The simplest triggers for the A + A combinations are centrality triggers, which are also quite important for physics reasons. However, the centrality requirements assumed for Table 11.1, namely the most-central 10% of events, still do not reduce the data rate enough to record all events. The light-ion and p + A collisions have extremely high interaction rates, requiring event selection already at LVL-1 in order to comply with the 40 μ s digitization time chosen for ADCs. The mean time between interactions expected for collisions of heavy ions does not require LVL-1 rejection,

Table 11.1: Interaction Rates, Time Between Events and Data Volumes in PHENIX

Beams	Type	$\mathcal{L}$ ($\text{cm}^{-2} \text{s}^{-1}$)	Rate (kHz)	T_{avg} (μs)	$\frac{dN_{ch}}{dy}$	$\langle N_{ch} \rangle$	Tracks/s	$\frac{\text{kB}}{\text{event}}$	MB/s
p + p	MB	$1.1 \cdot 10^{31}$	454.0	2.2	2.6	0.92	$4.2 \cdot 10^5$	1.9	863.
O + O	MB	$1.6 \cdot 10^{29}$	179.3	5.6	13.3	4.66	$8.4 \cdot 10^5$	5.6	1006.
Si + Si	MB	$7.1 \cdot 10^{28}$	118.5	8.4	24.7	8.63	$1.0 \cdot 10^6$	9.6	1133.
Cu + Cu	MB	$1.7 \cdot 10^{28}$	47.3	21.2	60.2	21.06	$1.0 \cdot 10^6$	21.9	1034.
I + I	MB	$5.1 \cdot 10^{27}$	23.3	43.0	130.1	45.53	$1.1 \cdot 10^6$	46.1	1074.
Au + Au	MB	$2.2 \cdot 10^{27}$	13.9	72.2	210.9	73.80	$1.0 \cdot 10^6$	74.1	1030.
p + p	Cn	$1.1 \cdot 10^{31}$	4.5	220.3	26.3	9.20	$4.2 \cdot 10^4$	10.1	46.
O + O	Cn	$1.6 \cdot 10^{29}$	17.9	55.8	55.5	19.43	$3.5 \cdot 10^5$	20.2	362.
Si + Si	Cn	$7.1 \cdot 10^{28}$	11.9	84.4	102.7	35.96	$4.3 \cdot 10^5$	36.6	436.
Cu + Cu	Cn	$1.7 \cdot 10^{28}$	4.7	211.5	250.7	87.74	$4.2 \cdot 10^5$	87.9	413.
I + I	Cn	$5.1 \cdot 10^{27}$	2.3	429.6	542.1	189.72	$4.4 \cdot 10^5$	188.9	435.
Au + Au	Cn	$2.2 \cdot 10^{27}$	1.4	721.6	878.6	307.49	$4.3 \cdot 10^5$	305.6	428.
p + O	MB	$1.3 \cdot 10^{30}$	406.1	2.5	3.5	1.23	$5.0 \cdot 10^5$	2.2	899.
p + Si	MB	$9.0 \cdot 10^{29}$	403.1	2.5	4.0	1.39	$5.6 \cdot 10^5$	2.4	962.
p + Cu	MB	$4.3 \cdot 10^{29}$	340.1	2.9	4.8	1.68	$5.7 \cdot 10^5$	2.7	906.
p + I	MB	$2.4 \cdot 10^{29}$	306.5	3.3	5.7	2.00	$6.1 \cdot 10^5$	3.0	912.
p + Au	MB	$1.6 \cdot 10^{29}$	276.6	3.6	6.4	2.24	$6.2 \cdot 10^5$	3.2	890.
Si + Cu	MB	$3.4 \cdot 10^{28}$	76.2	13.1	67.4	23.58	$1.8 \cdot 10^6$	24.4	1857.
Si + I	MB	$1.9 \cdot 10^{28}$	55.9	17.9	79.5	27.83	$1.6 \cdot 10^6$	28.6	1597.
Si + Au	MB	$1.3 \cdot 10^{28}$	45.0	22.2	87.1	30.47	$1.4 \cdot 10^6$	31.2	1404.

but the overall triggering system would benefit from some selectivity at LVL-1 in order to reduce the processing burden required in LVL-2 and LVL-3 for compliance with the data-archiving bandwidth. The overall rejection factors required from the various trigger levels are summarized in Table 11.2, which again updates the similar table in the PHENIX CDR to reflect the revised estimate of event size and the planned mean time of 40 μs between LVL-1 accepts. The resulting numbers of events per second which may be archived are listed in the rightmost column of this table.

11.2 Changes

11.2.1 Front-end Electronics Changes

1. We have decided upon 64-cell-deep AMUs.

Table 11.2: Required Trigger Rejection Factors and Accepted Events/Second

Beams	Type	Reduction required from LVL-1	Rate (MB/sec) after LVL-1 reduction	Reduction required from LVL-2*LVL-3	Accepted Events per sec.
p + p	MB	18.2	48.	2.4	10,400
O + O	MB	7.1	141.	7.0	3550
Si + Si	MB	4.8	238.	11.9	2100
Cu + Cu	MB	1.9	548.	27.4	910
I + I	MB	1.0	1074.	53.7	430
Au + Au	MB	1.0	1030.	51.5	270
p + p	Cn	1.0	46.	2.3	1960
O + O	Cn	1.0	362.	18.1	990
Si + Si	Cn	1.0	436.	21.8	550
Cu + Cu	Cn	1.0	413.	20.6	230
I + I	Cn	1.0	435.	21.7	105
Au + Au	Cn	1.0	428.	21.4	65
p + O	MB	16.0	56.	2.8	8900
p + Si	MB	16.0	60.	3.0	8320
p + Cu	MB	13.8	66.	3.3	7610
p + I	MB	12.1	75.	3.8	6650
p + Au	MB	11.1	80.	4.0	6240
Si + Cu	MB	3.1	608.	30.4	820
Si + I	MB	2.2	715.	35.7	700
Si + Au	MB	1.8	779.	39.0	640

2. The RICH will now use flash conversion instead of AMUs for analog storage.
3. An improved estimate of data volume has been developed, taking into account necessary tagging and labelling information.
4. Development of monolithic electronics has continued as part of the RHIC R&D program, with concentration on AMU development, silicon detector front-end development, ADC development, and timing channel (discriminator and TAC) development.
5. An agreement has been reached with KEK to utilize their shift-register pipeline-TDC for digitization of timing signals from DC and μ T. A development project to tailor this chip to our architecture has started.

6. Control and bus lines have been enumerated.

11.2.2 Trigger Changes

1. Level-1 work has centered on refining algorithms to improve background rejection and/or improve efficiency and on defining the GL1 trigger.
2. The Beam-Beam LVL-1 trigger has been redesigned to build histograms correlating hits in time to improve noise, beam-gas and QED-pair rejection.
3. The MVD trigger has been changed to use a sum of fired discriminators instead of an analog sum of pulse heights to improve noise rejection and enable operation for low-mass colliding ions. A lookup table using the vertex from Beam-Beam LVL-1 has been added to reduce effects arising from the long RHIC interaction diamond.
4. The RICH LVL-1 trigger has had an "adjacency check" added to avoid double-counting of rings and preserve good efficiency at realistic number of photoelectrons thresholds.
5. The μ ID trigger has been completely redesigned to follow roads in θ and ϕ . This improves background rejection for beam-gas events and shower-secondary hits arising from leakage from magnet surfaces, particularly the piston magnet. The μ ID itself has been redesigned, as described at the June 1993 Muon TAC Review; the μ ID LVL-1 trigger incorporates these changes.
6. A concept for the GL1 trigger involving cascaded lookup tables and a programmable routing layer has been developed.
7. A concept for LVL-1 to LVL-2 data transfer has been developed.
8. A general architecture for LVL-2 has been developed. This has led to the addition of a Level-2.5 trigger included in the output legs of the Event Builder.

11.2.3 DACQ Changes

1. Formats for data at the various stages through the data acquisition system have been developed.
2. We have decided to read out all data, including empty channels, from the Front-end Cards and transmit it to Data Collection Modules. These modules are responsible for zero-suppression, formatting, calibration and preparation of initial LL2 trigger primitives. They incorporate programmable devices (FPGAs and DSPs) to do this. Output ports on DCM crates provide parallel links to the EB, further stages of LL2, and the Monitoring network.
3. A design for the Event Builder which can operate at 1 GB/s has been developed. A concept for the buffers collecting the DCM outputs for transferral to the EB has been developed.

4. A decision has been made to use the hi-speed processors in the Event Builder output legs for another level of triggering, denoted Level-2.5. This trigger level is the first at which *all* parts of an event are brought to one point.

11.2.4 On-Line Computing System Changes

1. An entirely new section entitled "On-line Computing Systems" has been written (UCDR 11.6). It incorporates the original CDR sections "Event Archive" (CDR 11.4.3), "Ancillary Systems Control" (CDR 11.6) and "On-line Computer System" (CDR 11.7). Several topics in CDR section "Data Acquisition" (CDR 11.4) are also addressed.
2. A first global on-line computing model (for organizing the analysis of high-level requirements) is presented. A framework for system development, a kernel environment for system integration and operation, and classes of generic services are also described. Emphasis is on a common approach to development.
3. The architecture has evolved according to a network-transparent client-server paradigm, influenced by the EPICS model for control and monitoring.
4. Ancillary Systems Controls has been subdivided into an Ancillary Systems Interface System and subsystem-specific ancillary systems equipment.
5. The Master Timing, Event-data Flow Control, Event Builder/Level-3 Trigger, and Event-data Archive subsystems are now treated on an equal footing with detector and trigger data acquisition subsystems.
6. The "routing layers" of the CDR have been incorporated into the Monitor Network.
7. An overall development strategy is now described.

11.3 Front-end Electronics

11.3.1 Subsystem FEE and Data Volumes

The purpose of the Front-end Electronics (FEE) is to convert sets of signals from detector electrodes into digital event fragments in accordance with selections made by the LVL-1 trigger, order and label the event fragments for transmission to the DCMs, and provide for calibration and monitoring cycles. This involves analog signal processing with amplification and shaping to extract the optimum time and/or amplitude information from the signal, development of trigger input data and its transmission to dedicated trigger paths, buffering to allow for the time needed by LVL-1 to process selected data and decide whether to keep the event, plus digitization and a reasonably uniform interface to the downstream readout system. The designs and timing relations listed in the PHENIX CDR still describe the FEE planned for PHENIX, with the exception that the decision was made to use FADCs to convert the pulse height and timing information from the RICH counter PMTs.

Table 11.3: Channel counts, digitizer channels and LVL-1 trigger bits/sums for PHENIX.

Subsystem Symbol	Detector Channels	ADC Channels	TDC Channels	LVL-1 Inputs	
BB	132	132	132	132	FADC
MVD	34,816	34,816		96	sums
DC	10,240		10,240		
PC	46,400	46,400			
TEC	28,032	28,032			
RICH	6,400	6,400	6,400	6,400	FADC
TOF	2,048	2,048	2,048	2,048	bits
EMCal	25,536	51,072	25,536	6,384	FADC
μ T	19,680		19,680		
μ ID	18,360	18,360		17,280	bits

Table 11.4: Front-end Electronics Parameters for PHENIX Subsystems

Sub-System	preA mult.	Shaping Time(ns)	Resoln (bits)	Range (bits)	AMU/	ADC	TDC	Special Needs
BB	1	(PMT)	10	10	DMU	yes	yes	FADC, History
MVD	32	60	8	8	AMU	yes	no	History
DC	8	≤ 10	9	16	DMU	no	yes	pTDC
PC	16	400	11	11	AMU	yes	no	
TEC	8	70	4	"9"	DMU	yes	no	FADC
RICH	8	(PMT)	8	8	DMU	yes	yes	FADC
TOF	8	(PMT)	12	12	AMU	yes	yes	
EMCal	8	(PMT)	11	15	AMU	yes	yes	Dual ADC
μ T	8	30	9	16	DMU	no	yes	pTDC
μ ID	16	70	8	8	AMU	yes	no	

Summary of Channel Counts

A updated summary of the number of channels for the various detector subsystems in PHENIX Baseline is given in Table 11.3. This includes the number of detector channels, the number of ADC and TDC channels, and the number of signals provided to the LVL-1 system.

An updated summary of some of the parameters of the different PHENIX FEE systems

for minimum-bias events of various types are given above in Table 11.1.

11.3.2 Development of Functional Blocks

The R&D work carried out on monolithic electronics since the CDR has concentrated on developing circuits useable by the MVD, the EMCal and PAD detectors, the TEC, and the TOF. Circuits developed for the MVD include a preamplifier, analog memory and ADC. Those for the EMCal and PAD chambers include analog memory and ADC. FADC development was continued for the TEC, and discriminator and TAC development was carried out for the TOF and other detectors needing timing information, such as EMCal.

The silicon-strip preamplifier has an 85μ channel-to-channel pitch, sensitivity of 170mV/MIP (i.e. 41mV/fC), a risetime of 60ns and a noise of 900 e_{rms} for 4pF input capacitance. The integral linearity is 2% up to 11 MIP-equivalent signal. It requires a single 5V supply and dissipates $\sim 1\text{mW/channel}$. It is 1mm long for all three stages of the preamp. It requires a reset pulse periodically to discharge the feedback capacitor. This chip has been fabricated and tested thru ORBIT.

The remainder of the front-end electronics chain for the MVD includes an AMU and a multichannel Wilkinson ADC. Versions of both circuits are now designed and submitted for fabrication. The AMU is an 8-channel, 16-cell test version with at present only linear address decoding. Simulations indicate its power consumption will be dominated by the $250\mu\text{W/channel}$ for the readout amplifier. The ADC is an 8-channel, 8-bit Wilkinson designed to run at clock speeds up to 200 MHz , although actual operation will only be at 20 MHz , resulting in $12.8\mu\text{s}$ conversion time, to minimize power consumption.

A pad preamplifier, originally designed for the GEM IPC, has been submitted for fabrication. It is designed for 30ns peaking time, much faster than PHENIX PAD requires. Power dissipation from four stages is simulated to be 25 mW (four stages of shaping are used to obtain $1/t$ compensation for gas detector use) and the noise performance is expected to be $\sim 1800\text{e}$ for 35 pF input capacitance. The output stage has a baseline restorer for optimal drive of the AMU; it should drive a 25pF output load with no slewing. A PHENIX-specific design will employ 400ns shaping, resulting in better noise performance.

Various AMUs have been fabricated and tested. One under development for use with the EMCal has been tested at 11 MHz writing rate using an FPGA to drive the addressing lines in order to simulate real-world noise sources. It exhibits a memory cell "gain" of 0.956 on average. The pedestal variation measured is 3.8mV_{rms} (i.e. 0.8 LSB at 10 bits) while the integral nonlinearity is measured to be $\pm 0.05\%$ (i.e. 1 LSB at 11 bits .) The linear operating range extends up to 4.5V input for a 5V supply voltage.

Two types of discriminator have been fabricated and been given initial tests. One is a leading edge type, which would be useful for TOF, RICH and PbSc, while the other is a constant-fraction type, being developed for the PbGl, whose existing PMTs have a modest rise time of 15ns . Both types are auto-zeroed to remove offset and drift. They provide a timing output within $5\text{-}10\text{ns}$ for a 1V input and 15 mV threshold setting. The minimum settable threshold is 15mV in this design. The layout pitch is 170μ , meaning 16 easily fit onto a standard $4.8\times 4.8\text{ mm}^2$ monolithic die of the type envisioned for many of the monolithic FEE systems in PHENIX.

Two separate TAC development efforts have been pursued. Both have yielded tested devices this year. One has completed testing with a high-resolution setup and exhibits a noise floor of 25ps or less at full range, with the noise increasing as the square root of the time interval, as expected if TAC error is dominated by thermal noise on the chip. The other has only been tested with a system providing 10 bits measurement resolution at 100ns range, i.e. with a 100ps test channel width. Well over 100K counts still stay within one ADC channel again meaning that at 4-sigma level, a noise figure of ~ 25 ps is reached. Both results are encouraging for use in PHENIX, as systems with 40-60ps intrinsic resolution are planned. Again, these TACs are laid out with 170μ pitch, meaning 16 easily fit on a 4.8×4.8 mm² die. In order to provide sufficient time for settling, both TACs are designed to acquire on one bunch crossing, write on the next crossing, and reset on the third. This means a common-stop approach will be used to avoid deadtime in the TAC channels. This also means the 25ps resolution for the TAC must be maintained for full scale (as the test results have demonstrated), as that is where the particles with shortest flight times are recorded in a common-stop system. Both efforts have now moved on to laying out the TAC on the same substrate as the AMU in order to study integration and cross-talk issues.

The FADC development for the TEC has progressed thru testing of the first monolithic chip, which was fabricated thru MOSIS. The design allocates 28 levels for dE/dx measurement, ranging over 2fC-120fC of signal, and 3 levels for TRD photon measurement, located at 400,600 and 800 fC. The components on the test chip all work, and the FADC functions at the needed clock rate. Further development includes moving to a 1.2μ process and determining the level of clock feedthru on the substrate.

The ADC development concentrated on a 10-bit multichannel Wilkinson design, building on a design originally prepared for the leadglass detector in CERN experiment WA98. Test results show that 0.15% integral nonlinearity and 1/3 LSB differential nonlinearity are achieved over full range. The channel profile is 1.13 bins. The clock rate has been tested as high as 167MHz with single-edge triggering. This suggests 334MHz operation with double-edge triggering is possible, but this has not been tried. The power consumption is 3.3 mW/ch with a 100MHz clock and 31.8 kHz conversion rate (higher than the mean LVL-1 accept rate planned for PHENIX). Again, the chip has been laid out with 170μ pitch to facilitate multichannel chip layout.

11.3.3 Control and Bus Lines

The control and bus lines needed by the Front-end electronics cards are listed below. This particular example is for the case of the MVD. Other systems will have the same control lines, or a superset, whereas the bus lines listed necessarily depend upon the subsystem. The intent is to use the RHIC RF as the master timing reference in order to insure the PHENIX electronics operate synchronously with the beam. This insures a phase lock and also allows the FEE to track changes in beam revolution frequency in RHIC as a function of beam momentum. (This amounts to $\sim 0.5\%$ over the full beam momentum range at RHIC.) Other control lines are then viewed as arming signals whose action occurs on the next beam crossing pulse. A serial interface is provided to download, e.g., calibration and threshold information to the FEE. A latch line is included in the serial set to allow series-parallel

broadcast of this information. The bussed lines typically are output by the state machine responsible for overall operation of the front-end electronics on a card. This state machine was called the heap manager in the CDR.

- **Control Lines**

- Beam crossing clock (8.9 MHz)
- Initialize/reset
- LVL-1 trigger accept
- Integrator reset
- Test qualifier (inject calibration pulse on channels identified via information loaded from serial lines.)

- **Serial Lines**

- Serial data in
- Serial data clocking line
- Serial data latch (copy data from shift register to latch register when this line pulses true.)

- **“Bussed” Lines**

- AMU
 - * 6-bit write address
 - * 6-bit read address
 - * 1-bit write enable
 - * 1-bit read enable
 - * 1-bit reset address
- ADC
 - * 8 lines for links to ADC internal state machine
- Serial
 - * 5 lines to preamp front-ends
 - * Data, clock, latch, calibration trigger and calibration voltage
 - * Note the calibration voltage can be an analog DC level
- Buffer
 - * 5-bit address line to ADC output buffer
 - * Selects one of 32 on-chip buffers for output to FIFO

the Muon ID LVL-1 Trigger to have the four outputs shown, two of which will be input to the global LVL-1 logic of the GL1 Box. Note that there are more data bits in the LVL-1 Muon ID than there were in the CDR. The other change is to the input to the GL1-Box that was shown in CDR Fig. 11.30 as a 1-bit "calibration" word. This 1-bit word is now replaced with 3 bits of "mode/monitor" information. We anticipate that a 3-bit word called MODE will suffice to indicate the "mode" or type of data being taken, which could be data for monitoring the LVL-1 trigger, running test patterns, running diagnostics, or any special runs other than normal data taking.

The updates to the LVL-1 algorithms which provide inputs to GL1 involve four of the six sub-detectors used for GL1. (1) The BB algorithm has had a major redesign to detect the space-time location of the largest particle-hit activity on each side. (This replaces the CDR algorithm for the first-hit space-time location on each side.) (2) The MVD will provide a sum of fired discriminators and this sum will be input to a vertex-correction lookup-table circuit that will correct for the η -dependent multiple-hit probability. (3) The RICH algorithm has had an adjacency-check step added, resulting in higher efficiency and avoiding multiple counting at low ring-count thresholds (which was a problem with the original algorithm in the CDR). (4) The μ -ID algorithm has been completely redone to require physically meaningful paths through the μ -ID layers and thereby exclude the shower secondaries and beam-gas interactions that the CDR algorithm picked up. These changes are discussed further in the following sections.

Beam-Beam

The purposes of the BB LVL-1 trigger as described in the CDR are unchanged, but the method used to achieve them has been improved. In the first three paragraphs of the BB writeup on pp. 11-44 and 11-45 of the CDR, the update only changes one sentence near the end of the second paragraph. The sentence to be replaced is "The time for each array is taken to be the earliest time (within certain limits) from all of the elements of each array with valid hits." (The BB has two arrays, each with 66 "flash-light" Cherenkov detectors.) Instead, this sentence should be replaced by the following sentence: The time for each array is taken to be the time bin (within certain time limits) in the (new, updated) histogram algorithm containing the greatest number of counts. That is, we no longer consider it acceptable to rely on the earliest time. As in the CDR, other criteria of time and count limitations must be met, and, if both arrays are hit, an average time and an estimate of vertex position are deduced.)

The reason for the change is the concern that the "earliest time" might come from a noise "hit" or a beam-gas, electromagnetic-dissociation type, or other background type of interaction rather than from the desired speed-of-light particle(s) from the beam-beam nuclear interaction. The latter signals will all be detected within a small time interval and should be numerous whereas background interactions should produce isolated hits. Beam gas and electromagnetic dissociation events will plikely occur at the correct time but will have too small multiplicity of hits. Noise hits will be random in time. The histogram algorithm will search out the time interval with the most hits in each array. To do this, we will use 6 bits from the 10-bit time-corrected FADC signals from the front end electronics of each of

Figure 11.3: Schematic diagram of LVL-1 trigger for BB (early stages).

the 66 Cherenkov detectors. The least significant of the 10 bits will correspond to 25 ps, thus we will use bits 2 (50 ps) through 7 (3.2 ns) to construct a 64 time-channel histogram, each channel 50 ps wide, for each of the 2X66 detectors. These histograms are summed channel by channel over the PMTs to form 64 histograms (per side) of the number of pMTs firing within each time channel. The time channel with the most hits for each side is then used to produce the various outputs of the BB LVL-1 trigger (which are the same as described in the CDR).

Schematic diagrams for different stages of the BB LL1 processing are shown in Fig.11.3 and Fig.11.4. These two figures replace Figs. 11.32 and 11.33 of the CDR. The steps in the BB LL1 processing pipeline are described in sequential order below. Also included are the number of bunch crossings (assuming 112 ns between bunch crossings) taken by each step in the pipeline.

Step 1 - Flash-digitization and time correction (1 crossing): This is done on the BB front-ends and provides time-corrected 10-bit words. (Only 1 bunch crossing

Figure 11.4: Schematic diagram of LVL-1 trigger for BB (later stages).

time is required.) Bits 2 (50 ps) through 7 (3.2 ns) form the 6-bit words $PMT(M)$, where $1 \leq M \leq 66$ for each side, that are the inputs to the BB LVL-1 circuit.

Step 2 - Time bin routing (2 crossings): For each array, the 6-bit $PMT(M)$ words produce the 1-bit time words "PMT-Time(M,N)", $1 \leq N \leq 64$. PMT-Time(M,N) is 1 for a hit $PMT \# M$ in time bin $\# N$ and 0 for a missed PMT, where hit or missed means within the 6.3 ns dynamic range in time. Each $PMT(M)$ can produce a hit in only one time bin N . The 66 inputs from each side are collected to sum the hits in time bin N in the next step.

Step 3 - Histogram adder stage (8 crossings): The PMT-Time(M,N) values are summed over M (the PMT index) by a pipeline digital summer, independently for each time bin N , to form the 13-bit word "Time(N)", which is the desired histogram in time of the number of hits in the array in that time bin. Of the 13 bits, 7 bits are needed to cover the 67 values of the sum (from 0 to 66) and 6 bits are needed as an index for the 64 sums of hits corresponding to the 64 values of the time index N .

Step 4 - Time bounds (1 crossing): The 64 values (per side) of Time(N) are then compared to upper and lower time limits. Any value of Time(N) outside these

limits is set to 0 in the "time range gate" circuit. This mechanism will limit event acceptance to times closely correlated to the RF clock and the region ± 30 cm in z around the center of the interaction diamond.

Step 5 - Finding maximum counts (7 crossings): For each side, a "count sorter logic" finds the Time(N) bin with the most hits. This is done with cascaded comparators, each of which finds the larger of two numbers. (Step 7 is an example of this logic.) If two (or more) time bins have the same number of hits, the earliest one is chosen by this circuit. If there are no hits, then the lowest time, Time(1) (which will be below the lower time limit) is chosen. The chosen Time(N) for the right side is output as the 6-bit word "T-Right" after having been checked to see if it exceeds the minimum-count threshold. If it does, the 1-bit word "Right Hit" is set to 1 for TRUE. The analogous words for the left side are denoted "T-Left" and "Left Hit".

Step 6 - Interaction time and vertex (1 crossing): The left and right Hit Times are added to get "T-Average", which is the time at which the event occurred, and subtracted to get "T-Vertex", which is the location of the event's vertex along the z -axis. (Note that the time sum needs to be renormalized to get a vertex in cm.) These two words are also made available to LVL-2. It must be recognized that the quantities "T-Average" and "T-Vertex" (which are produced for each beam crossing) are valid only for events for which the 1-bit "Valid Hit" from step 5 is true for both sides.

Step 7 - Vertex Validation (1 crossing): The vertex position is tested against preset maximum (V(max)) and minimum (V(Min)) values. This step provides two output signals, "Event Flag" and "Validity Flag". Event Flag is TRUE if either of the two Beam-Beam arrays have valid hits. Validity Flag is TRUE if both arrays were hit and the vertex is within the preset window. If the Validity Flag is FALSE, the rough vertex to be used by LVL-2 must be computed by other means.

The LL1 system for the BB counters produces 136 data words, of which 128 are 13-bit time-histogram words passed on to LVL-2. There are four 6-bit times: "T-Right", "T-Left", "T-Vertex", and "T-Average" (for which the least-significant bit is 50 ps). There are four 1-bit words: "Right Hit", "Left Hit", "Event Flag", and "Validity Flag". All 136 words are passed on to LVL-2. The two words "T-Vertex" and "Validity Flag" are used in the LVL-1 algorithm for the MVD to make vertex corrections to the Si strip multiplicity. Only the two words "Validity Flag" and "Event Flag" are used in the GL1 Logic Box. These two words, and perhaps other BB words, can also be made available for global distribution.

MVD

The update of the LVL-1 logic for the MVD consists of two changes. The first is that the inputs to the FADCs will be the 256-channel sum of fired discriminators instead of

Figure 11.5: Schematic diagram of LVL-1 trigger for MVD.

the 256-channel analog sums that were described in the CDR. The change to a sum of fired discriminators was made to eliminate the summing electronic noise from 256 channels. (Although the signal-to-noise ratio is expected to be at least 10, the vast majority of the 256 channels will have only noise.) The second change occurs in the LVL-1 algorithm for the MVD. It consists of the addition of a lookup table step, with the input address to the lookup tables being the 7-bit FADC digitization of the 256-channel sums of fired discriminators. The lookup tables have dual purposes. For the Si strip detectors, they provide a vertex correction that uses the 7-bit word "T-vertex" (time difference signal) and the 1-bit word "Valid Vertex" (which is 1 if the T-vertex is meaningful) from the BB detector. The vertex-correction lookup-table will correct for the η -dependent multiple-hit probability. The second purpose is to allow, by adjusting the lookup table settings, the 15-bit word called "Total Sum" to represent either the Total Sum, Right Sum or Left Sum, where left and right are either side of the $z=0$ center of the interaction diamond. For asymmetric beams, it may be desirable to use in the GL1 Box the sum of one side instead of the sum of both sides. This

will be accomplished by adjusting the lookup tables so that the single word "Total Sum" to provides the desired sum.

The new lookup table step becomes step 3:

Step 3 - Adjust via lookup tables (1 crossing): The inputs to each lookup table is the 7-bit FADC from step 2, the 7-bit word "T-vertex", and the 1-bit word "Valid Vertex". (The two vertex words come from the BB detector.) The lookup table "adjusts" the multiplicity inputs and provides a 7-bit adjusted output to step 4.

Steps 2, 3 and 4 of the CDR are thus renumbered to be Steps 3, 4, and 5, respectively. Figure 11.34 of the CDR is now replaced by Fig.11.5, which updates Fig. 11.34 by including the lookup table step and T-vertex input.

The number of beam crossings which must elapse before "T-Vertex" can be provided to the lookup tables is somewhat uncertain, thus the first beam crossing shown in the figure is labeled with an N. Nine additional beam crossings are required to finish the MVD algorithm (most of which are taken by the pipeline digital summers). The value of N is expected to be about 20 (as shown in the update for the BB LVL-1 trigger) thus the "Total Sum" is expected to be available in about 29 beam crossings. The uncertainty in N is expected to be 1 beam crossing and is caused by the unknown time to do the 10-bit flash digitization and time correction to the input to the BB LVL-1 trigger. It was assumed that this would be done within 1 beam crossing in arriving at 20 beam crossings for the BB "Valid Vertex" word. The 29 beam crossings would become 30 if an additional beam crossing were needed for the BB front-end electronics to complete flash digitization. However, if the pipelined digital summing were made faster than the one-pair-sum-per-beam-crossing rate assumed, then the total time would be several beam crossings less than 29.

The CDR also states that the three sums, "Total Sum", "Left Sum", and "Right Sum" are all input to the GL1 Box. For the reasons stated above, this is no longer correct, and only "Total Sum" is used at GL1. All three sums remain part of the data stream, however.

RICH

Simulation studies of the algorithm given in the CDR for the RICH have shown that two (or more) adjacent unit cells (refer to CDR Fig. 11.35 on p. 11-51 for the definition of the RICH unit cell) could each exceed the ring-count threshold for a single incident electron, leading to double (or multiple) counting of the Cerenkov ring produced by that electron. To eliminate this, we modified the RICH algorithm of the CDR by adding an adjacency-check step as step 9 of the RICH algorithm. The adjacency check involves a 3x3 array of unit cells, with the center unit cell being checked against 4 of the other 8 unit cells:

Step 9 - Adjacency check (2 crossings): If any of the 4 adjacent unit cells are over threshold, then the unit cell at the center of the adjacency group will not be counted in the ring sum. Otherwise the central unit cell will be counted (if it were a 1).

Only 4 of the 8 adjacent unit cells in the 3x3 array are checked, because the adjacency check is done simultaneously for all unit cells (128 per arm) and we want *one* of the unit

Figure 11.6: Schematic diagram of LVL-1 trigger for RICH (last 2 steps).

cells to be counted for an incident electron. Were all 8 adjacent unit cells included, the simultaneous checking would prevent *any* unit cell from being counted. (For example, if unit cells A and B were adjacent and all 8 adjacent cells were included, then A would prevent B from being counted when A was the center cell, but B would also prevent A from being counted when B was the center cell, and of course both are treated as center cells but in different parts of the algorithm.)

The "ring-sum step" (step 9 in the CDR) now becomes step 10 and the inputs to the ring sum are the "valid cell" outputs of the new step 9. This is shown in Fig. 11.6, which updates Fig. 11.37 of the CDR. Only steps 9 and 10 are shown in Fig. 11.6 because steps 1 through 8 of the RICH algorithm are unchanged from the CDR. (The modified algorithm now requires 22 beam crossings instead of 20.)

The performance of the RICH algorithm was checked with PISA simulations and the results are shown in Fig. 11.7. The efficiency of the RICH algorithm depends upon the value of the ring-count threshold which is given in terms of the number of photoelectrons per module. The figure shows the frequency f at which an electron ring is counted, where the subscript is the multiplicity of times a ring is counted. The left side of the figure gives the results for the original algorithm of the CDR, and the right side gives the results with the adjacency step. The right side shows no multiple counting, and the frequency of single counting is the desired 100 appreciably below 100 number of detected photoelectrons. (The

Figure 11.7: Efficiency of the RICH algorithm with and without the adjacency check.

mean number is 24, at which value the frequency of single counting is 50% labelled $1 - f_0$ shows the frequency of counting a ring at any multiplicity; it is identical to the frequency of single counting when the adjacency check is made. The left side of the figure shows why the adjacency check is needed. At low values of the threshold, a ring is more likely to be multiply counted than counted once. Indeed, for thresholds from 1 to 6, double or higher counting of a ring is more likely than counting it once. As the threshold increases, the multiple counting frequencies drop significantly, but so does the single counting frequency (since more rings are formed whose number of detected photoelectrons is on the low side of the Poisson distribution). Therefore, simply raising the threshold in the original algorithm to reduce multiple counting was not viable, making the development of this algorithm necessary.

Figure 11.7 clearly shows that adding the adjacency check as step 9 of the algorithm provides a dramatic improvement in the performance of the RICH algorithm. With this step, we should be able to operate the RICH algorithm at a threshold setting of about 10 photoelectrons (or 40 of detected photoelectrons) with high efficiency and with little dependence of the efficiency on the value of the threshold. The latter point is significant because the mean number of detected photoelectrons depends upon the path length of the particle in the RICH gas and thus upon the emission angle θ of the electron entering the RICH (and, to some extent, also upon the vertex position of the event), and we insist upon an

Figure 11.9: Schematic diagram of LVL-1 Trigger for Muon ID (later stages)

at LVL-1. The detection of a pair of muons is relatively straightforward in the acceptance of the muon system so the muon LVL-1 trigger should be easily implemented, and it should place less of a burden on the muon LVL-2 trigger. The muon LVL-1 trigger is designed to find tracks by following roads through all six layers of the Muon ID and to count identified tracks penetrating to at least layers two or three.

As discussed in Chapter 10, the Muon ID consists of a set of six absorber layers, after each of which is a limited-streamer-tube plane with pad readout. The Muon ID readout pad planes are segmented in θ and ϕ to assist in the trigger identification problem resulting from multiple scattering as muons pass through the absorber layers and to discriminate against particles that do not maintain a nearly constant value of the angle θ as they penetrate the muon ID. While both θ pads and ϕ pads are instrumentated, only the θ pads are used for the LVL-1 trigger. The ϕ pads are used to mask the μ T chambers at LVL-2 and later refine tracks found in those chambers. In the following the θ pads will be referred to as LVL-1 pads. See Chapter 10 for further discussion of the pad sizes and layout.

In Layer 1, the LVL-1 pads have the dimensions (θ by ϕ) of 1° by 2.813° , thus there are $27 \times 128 = 3456$ pads to span the 27° θ range and the 360° ϕ range. There are two sets of equal-size pads, viewing independent gas volumes, in Layer 1. Corresponding members of these two sets are ORed (as described below) in the LVL-1 trigger to assure a high efficiency for muon detection at Layer 1. In Layer 2, the LVL-1 pads have the same dimensions of 1° by 2.813° . There is only one set of LVL-1 pads in each of Layers 2 through 6. Layers 3 through 6 θ pads have the dimensions 1° by 5.625° , making $27 \times 64 = 1728$ pads in each of these layers. The total number of LVL-1 pads in the LVL-1 trigger is thus 17,280.

The LVL-1 Muon-ID trigger defines muon "roads" through the muon endcap. The roads

account for the increase in multiple scattering as the muons penetrate each absorber layer. Each "road" starts on a unique group of 4 pads in Layer 1, one pad wide in θ by two pads wide in ϕ , (2 adjacent pads on the first layer of LST detectors of Layer 1 and the two aligned adjacent pads on the second layer of LST detectors of Layer 1). These 4 pads are ORed together (to provide high efficiency) and define 1728 (27 in θ and 64 in ϕ) unique "roads" for the muon ID. While the roads are disjoint at Layer 1, the effects of multiple scattering in θ and ϕ plus track pitch in ϕ due to the preceeding magnetic field in μ T cause the road's shape (the envelope of roads for muons with energies from 1 GeV up) to flare into a larger cross sectional area as the muon progresses along its course. Thus adjacent "roads" share common LVL-1 pads in the layers beyond Layer 1. The size at each layer was deduced from simulations, and the results converted into the nearest larger integer to specify the number of adjacent pads to OR together. Layer 1 has 4 pads (2 front and 2 back), which define the 1728 roads. Layers 2 and 3 each use 6 adjacent pads (two in θ by three in ϕ for Layer 2 and three in θ and two in ϕ for Layer 3, which has the larger pads), Layer 4 uses 9 adjacent pads (three in θ by three in ϕ), and Layers 5 and 6 each use 12 adjacent pads (four in θ by three in ϕ).

The muon LVL-1 trigger finds and records the particle hits in the roads at Layers 1, 2, 3, 4, 5 and 6 independently. The only valid roads are those for which there is a hit at Layer 1. A valid road identified as having reached a depth N may have failed to fire one or more of the layers from 2 to N as long as it does not "skip" more than 1 consecutive layer. The layers at which each given road has a hit are passed on to LVL-2 as 1728 6-bit words called "Muon Road". The number of valid roads with hits at Layers 2 and 3 are separately summed and passed on as 11-bit words "Muon Sum-2" and "Muon Sum-3". A third sum, consisting of the number of valid roads with a hit at either Layer 2 or Layer 3 is called "Muon Sum-1". Muon Sum-3 and Muon Sum-1 are also provided to GL1 to be used in determining the number of candidates for $e\text{-}\mu$ pairs and $\mu\text{-}\mu$ pairs, respectively.

While massive in implementation, the LVL-1 trigger for muon-ID is simple in concept and involves only four logic steps. Schematic diagrams for these four steps are shown in Fig.11.8 and Fig.11.9. (These figures replace Fig. 11-41 of the CDR.) The steps in the Muon ID LL1 processing pipeline are described below:

Step 1 - Pad Hit Finding (1 crossing): The LVL-1 pad hits are determined by discriminators in the front-end electronics of the Muon ID system (This front-end step is not shown in Fig.11.8.)

Step 2 - μ Road ORs in Each Layer (1 crossing): All of the LVL-1 pads associated at each layer with a given muon road are ORed together and latched to form the inputs to the next step. As Fig.11.8 shows, the width of each road increases from layer to layer as a particle progresses into deeper layers.

Step 3 - Depth Recognition Logic (1 crossing): This step simultaneously receives all inputs for a given road "N" and records the depth of that road subject to two rules: (a) Each road recognized as "live" must be hit at Layer 1 and at Layer 2 or 3, and (b) A live road may contain missing or unhit layers, but two missing

layers in succession "stop" the road, and its depth is then recorded as the last layer that was hit. Also in this step the depth-2 and depth-3 signals are ORed together to form the word "Layer(1,N)" which is used to count all live roads. The 1728 6-bit words called Muon Road, which have the depth identification for each layer for each of the 1728 roads, are also passed on to LVL-2.

Step 4 - Form Depth Sums (11 crossings): Three sums of "live" roads are made using pipeline digital summers to provide a count of live paths at the associated depths. "Muon Sum-1", "Muon Sum-2" and "Muon Sum-3" are respectively the sum of live roads reaching at least depths 2-or-3, depth 2, and depth 3. Two of the three sums are sent to GL1 and all three are passed on to LVL-2.

The LVL-1 trigger for the Muon ID provides two outputs: a depth record called "Muon Road" for each of 1728 roads, and 3 "muon sums" for roads that reach (at least) Layers 2 and 3. "Muon Sum-1" is used at GL1 to make μ - μ pair triggers and "Muon Sum-3" is used at GL1 to make e- μ pair triggers. The depth record contained in Muon Road will be used by the muon track finding/fitting algorithm at LVL-2.

LVL-1 FIFO Architecture

The update to the CDR on the topic of the LVL-1 FIFO Architecture (pages 11-64 to 11-66 of the CDR) consists of the following additional comments concerning the LVL-1 to LVL-2 Data Transfer.

Upon a decision by GL1 to accept an event, all LVL-1 pipeline information for that event is moved into a temporary storage facility called the "Transfer FIFO". This facility, which was discussed in the CDR, is organized as a FIFO (First-In, First-Out) memory with a word length sufficient to accept all LVL-1 information from an event as a single word. Availability considerations suggest that this FIFO will be 64 words deep. The transfer will occur in a single 112ns beam-crossing cycle.

The information contained in a single FIFO word will be transferred in parallel to two places. One copy is moved into the experimental data stream, for inclusion in its entirety in the entire data block for that event. Because of the necessity for LVL-2 to gain access to this same data in a more selective manner, a second copy will be made available to LVL-2 in a multi-drop fashion, enabling that facility to capture only portions of the event required for specific trigger algorithms in LVL-2. Issues of data rate and event pipeline traffic statistics suggest that these transfers must be made in less than 32 μ s to prevent bottlenecks at this point. Input ports on the Data Collection Module (DCM) crates and DCMs themselves will be used to merge this "trigger" copy of the LVL-1 data selectively into the datastream.

The LVL-1 data stored in the Transfer FIFO may be characterized by three properties:

- a. *Location:* All geometrical information associated with LVL-1 data is implicit in the hardware location (i.e. specific address) at which the data resides within the long FIFO word. (Geometrical information refers to the indices of arrays such as the array of RICH modules, EMCal tiles, etc.)

b. Content: This refers to an actual data word stored at a specific hardware location.

c. Integrity: While this attribute of LVL-1 data is transparent, it is imperative that measures be taken to insure that the individual components constituting the Transfer FIFO maintain proper time alignment. This will be accomplished by embedding critical parity information into each Transfer FIFO word.

The integrity of LVL-1 data must be confirmed across the entire length of the FIFO word to insure that no time skew has occurred within the word. Should this test fail, the data transfer operation must be aborted and the LVL-1 hardware pipeline must be stopped and restarted.

To move the LVL-1 data into the rest of the experiment data stream, the long output of transfer-FIFO words must be changed into a 24-bit, word-serial format common to all other PHENIX data accepted by the DCMs. This requires that the location property listed above be made explicit and joined in appropriate fashion with the content property. This will be accomplished by formatting the data into 24-bit address/data units, adding header and trailer blocks recognized by PHENIX data transfer services, then transferring the resulting data structures in word-serial fashion over 24-bit wide data paths to LVL-2 and the DCMs (Data Collection Modules).

Global Level-1 Trigger

The Global Level-1 subsection on CDR page 11-61 lists the types of input data to the GL1 Box. Under **Status Information** the statement for *Calibration Status* should be replaced by:

Mode Status - A 3-bit word showing whether the data is "real" data (Mode=0) or one of seven other modes such as monitoring of LVL-1 data, reporting LVL-1 settings, reloading LVL-1 settings, taking standard or special (calibration) test patterns, or running diagnostics.

Under **Centrality Measurements**, the "multiplicity" quantities *Total Sum* and *Left and Right Sums* from the MVD are listed. (The MVD is divided in halves at $z=0$.) The choice of which of the three sums (left, right or total) from the MVD is passed to the GL1 Box will be made by setting bits in the lookup table for the Total Sum, and only the Total Sum will be input to the GL1 Box. Thus the CDR mention of *Left and Right Sums* here for central measurements should be ignored and the description for *Total Sum* changed to:

Total Sum (Multiplicity) - Total number of hits seen in MVD. (For asymmetric beams the Total Sum may be reprogrammed to equal either the Left Sum or the Right Sum of hits in the MVD.)

Following the listing for *Hit Sum* on CDR page 11-62 under **Hadrons**, the balance of the Global Level-1 Trigger section should be updated. Since the lookup tables described in

Figure 11.10: Schematic diagram of the GL1 Logic Box

the update accomplish the functions of the PPMD described in the CDR, and therefore the PPMD and Fig. 11.42 (the schematic diagram of a PPMD) are obsolete.

The GL1 Box will be designed to allow flexibility in programming the GL1 decision matrices. It will be possible to program it so that most of the GL1 trigger decisions other than the lowest-level interaction triggers are able to combine data from different sub-systems in PHENIX. Many LVL-1 trigger decisions only require an AND between a "valid interaction" decision based on the BB data and a threshold on an energy sum or a hit multiplicity. However, some of the observables of interest require further correlations. For example, we plan to use the Photon-2 sum, which counts photons above a relatively low energy threshold,

in low-multiplicity events where the photon showers are well separated, and plan to rely more on the Photon-1 sum, which uses a higher threshold, in the high-multiplicity events. We therefore envision using cascaded lookup tables in order to implement several such trigger decisions in parallel, to provide the necessary flexibility in the trigger, and to make the GL1 trigger decision fully (re)programmable.

Fig. 11.10 shows how we plan to use lookup tables in the GL1 "Logic Box". In the figure the subsystems contributing to GL1 are shown to the far left. The final accept/reject output of GL1 is shown at the bottom right. The flow in time is from left to right. The inputs to GL1 come from six of the PHENIX subsystems, namely BB, MVD, TOF, RICH, EMCal, and μ ID. Additional inputs include the RHIC beam status, the transfer FIFO status, and the 3-bit word called "Mode". (These inputs were shown earlier in Fig. 11.2.) Note that two options are shown for the EMCal information, depending on whether or not there is an indication of charge from PC-3. With PC-3 charge, the EMCal provides an electron trigger (called e^- in the figure). Without PC-3 charge, the same EMCal sum would provide us with γ -3, the sum of EMCal tower hits above a third energy threshold.

We first consider the GL1 logic of Fig. 11.10 for the case of normal data taking, for which the value of Mode is 0. (The bit counts listed at the exits of the first set of lookup tables give the numbers of distinct conditions possible for $\text{MODE} = 0$.) The first active step taken in the GL1 Logic Box is called the "Discriminator Lookup Step." This reduces the number of bits from the maximum channel count to a more manageable, and useful, number for LVL-1 decisions. Many of the inputs to the GL1 trigger are in the form of particle counts or energy sums. We will not be interested in the exact values of these inputs but rather in whether or not they fall within a given range of energy or multiplicity. Most of the "Discriminator Lookup Tables" take a 9-13 bit word as input and produce a 3-bit output word indicating which of seven non-zero programmed thresholds the input number exceeds. For the two multiplicity inputs (one is the EMCal's "Global Sum" and the other is the MVD's "Total Sum"), 4-bit words are used to allow setting of 15 non-zero programmed thresholds. It is quite probable that the number of output bits shown is one or two bits more than the minimum number needed. For example, most GL1 decisions on the RICH's "Ring Sum" need only two bits rather than the three shown because two bits provide us with the four discriminator values 0, ≥ 1 , ≥ 2 , and ≥ 3 , which should be enough to make any GL1 decisions that are based on a total ring count. The extra bit or two shown in the figure provides margin for additional choices that could be used to keep track of subdetector region, such as the RICH arm or the EMCal wall, in order to program regional information into the lookup tables.

The third column of Fig. 11.10 is symbolic. It shows how some of the LVL-1 signals are combined to provide the "physics" combinations of interest at LVL-1 before imposing later conditions based on event multiplicity or the location of the interaction vertex or status bits. The third column shows where LVL-1 signals from all six PHENIX subdetectors involved in LVL-1 are combined for the first time. (Earlier in the pipeline, before signals are input into GL1, there were local combinations of PC-3 sums and EMCal towers needed to generate the EMCal LVL-1 electron outputs, and there were inputs from BB to the MVD LVL-1 logic to provide the vertex corrections needed to obtain the MVD's multiplicity output "Total Sum".)

Table 11.6: Sample GL1 logic for the Mode=0 γ -1 trigger using Mult-1

Type	Word*	Trigger and Prescale Values									
		G1	G2	G3	G4	G5	G6	G7	G8	G9	G10
γ -1	γ -1 Sum	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 2	≥ 2
Mult-1 (MVD)	Total Sum	3	2	1	0	3	2	1	0	3	3
BB Flags	Validity Flag	1	1	1	1	-	-	-	-	-	-
	Event Flag	1	1	1	1	1	1	1	1	1	-
Status flags	L-beam Status	1	1	1	1	1	1	1	1	1	1
	R-beam Status	1	1	1	1	1	1	1	1	1	1
	"Mode=0?"	1	1	1	1	1	1	1	1	1	1
Prescale (One in 2^N)	Factor N	0	1	3	6	0	1	3	6	0	0

* γ -1 Sum and MVD's Total Sum are shown as 2-bit words in this sample table. The symbol "-" means "don't care".

This overall combination is accomplished by a (reprogrammable) routing and grouping of bits to inputs of a second stage of lookup tables.

The lookup tables called "Trigger" in Fig. 11.10 have as inputs the physics combinations, two types of multiplicity (or centrality) inputs, the two 1-bit time flags from BB, and the two 1-bit RHIC beam-status words. The largest number of bits input to a single Trigger lookup table is 18 for e - μ and for e' - μ . We anticipate using identical, standard-size lookup tables at all locations in the GL1 Box to conserve on engineering and fabrication costs. The number of outputs from these lookup tables depends upon the word size of the lookup table; we anticipate between 10-20 Trigger outputs of each type. An example of trigger and prescale logic is given below. Each of these 10-20 outputs for each trigger type then goes to a prescale logic step in which a programmed prescale factor is applied. The outputs of the prescale steps in Fig. 11.10 are labeled by a letter for the trigger type and a number for the number of triggers of that type (shown arbitrarily as 10 for each type in the figure). The grand total of 100-200 prescaled triggers are then fed to a cascade of logical ORs, so that any one of them can provide an accept condition for GL1. We plan to adjust prescale factors at run time to establish the desired fractions of events taken of each trigger type. However, as indicated by the final logical AND, an accept at GL1 also requires the transfer FIFO to be ready to accept transfer of the LVL-1 output, which means the experiment cannot be busy.

Table 11.6 gives some samples of the logic involved in the trigger lookup table and the associated prescale factors for the γ -1 set of triggers for Mode=0. The γ -1 trigger has a high-energy threshold (about 3 GeV) in an EMCAL tower. It is a useful trigger for a high-multiplicity collision. The basic idea is to accept at LVL-1 any event with a single high-energy photon in a central collision. There are two types of centrality or multiplicity input that can be used to do this. Table 11.6 illustrates how "Mult-1" (from the Si MVD) could be used. Table 11.7 illustrates how "Mult-2" (from the EMCAL) could be used. These two tables

Table 11.7: Sample GL1 logic for the Mode=0 γ -1 trigger using Mult-2

Type	Word*	Trigger and Prescale Values									
		G11	G12	G13	G14	G15	G16	G17	G18	G19	G20
γ -1	γ -1 Sum	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 1	≥ 2	≥ 2
Mult-2	Global Sum	3	2	1	0	3	2	1	0	3	3
BB Flags	Validity Flag	1	1	1	1	-	-	-	-	-	-
	Event Flag	1	1	1	1	1	1	1	1	1	-
Status flags	L-beam Status	1	1	1	1	1	1	1	1	1	1
	R-beam Status	1	1	1	1	1	1	1	1	1	1
	"Mode=0?"	1	1	1	1	1	1	1	1	1	1
Prescale (One in 2^N)	Factor N	0	1	3	6	0	1	3	6	0	0

* γ -1 Sum and EMCAL's Global Sum are shown as 2-bit words in this sample table. The symbol "-" means "don't care".

are nearly identical. In implementing the γ -1 triggers, these the logic of these two tables would be done in parallel, with an OR of all prescaled outputs, so each multiplicity type can be used independently of the other. Thus the lookup table labeled G1-G10 in Fig. 11.10 would actually consist of two parallel lookup tables to accomplish the trigger logic such as is illustrated by the two sample tables shown here. (There is no implied limit of 10 triggers per multiplicity type. Both this number, and the limited number of input bits shown in the table, are merely for brevity.)

In Table 11.6 the 2-bit word " γ -1 Sum" is shown with the four values 0, 1, 2, and 3 which respectively stand for thresholds of 0, ≥ 1 , ≥ 2 , and ≥ 3 for the number of γ -1 candidates in the event. The 2-bit word "Total Sum" is shown with the values 0, 1, 2, and 3 which respectively stand for four discriminator settings imposed by the discriminator lookup step. (The sample tables show only 2-bit words, but we will design for a 3-bit word for γ -1 Sum and a 4-bit word for Mult-1 to allow for more values for these variables.) The special 1-bit word "Mode=0?" is included to indicate that Mode must be 0 in order for any of the "real" data triggers to be satisfied. (The above comments apply also for Table 11.7 which replaces Mult-1 by Mult-2.)

In the ten sample trigger settings of Table 11.6, all ten require Mode 0 (for "real" data) and all ten require both left and right beam buckets to be full. Triggers G1 through G4 also require the Validity flag from BB (both sides fire and the vertex within limits), but for four different values of Mult-1. The prescale factor N is 0 for G1 since discriminator 3 is the highest and we want to keep every event that satisfies this condition. For discriminator 2 (lower than discriminator 3) the prescale factor N is set to 1 so that only 1 such event in 2 would be kept. For trigger G3 and discriminator 1, only 1 in 8 of such events would be kept, and for the low setting of discriminator 0 in trigger G4, only 1 in 64 would be kept. Triggers G5 through G8 are similar, but with the weaker BB requirement of the Event flag

only (i.e., only one side fires). Samples G9 and G10 require 2 or more EMCAL towers above the tower discriminator setting of the γ -1 trigger and $K > K_4$, with G9 requiring an Event flag and G10 requiring no BB flags. It is easy to imagine extending the possible conditions beyond the 20, each with its own prescale factor, shown in these two tables.

Triggers similar to those shown in these tables would be accomplished in a similar manner for the other trigger types, with a parallel lookup table for each of the two multiplicity choices. For the "e-e" triggers at the top of Fig. 11.10, the basic idea is to trigger on pairs of electrons, so the logic required would be similar to that shown in the tables above for the γ -1 triggers (except that a minimum of two electrons would be required for this e-e trigger). If two bits are sufficient for e-e, then the third bit shown in Fig. 11.10 could be used to indicate which RICH arm was involved. This would permit the lookup table to be programmed to require both electron candidates in the same arm, which would favor events with high transverse momentum for an e-e pair. It is worth repeating here that designing GL1 with such flexibility is imperative, in order to allow for programming of additional triggers.

The use of two parallel lookup tables is an efficient way to accomplish alternative choices such as the two multiplicities. In the example of table 11.6, the input to the lookup table for Mult-1 has five 1-bit words (BB, Status, Mode) and two 2-bit words ($\gamma - 1$ and Mult-1), for a total of nine input bits. The Mult-2 lookup table also has nine input bits. In order to accomplish the same logic in one table that has both Mult-1 and Mult-2 as 2-bit inputs would have 11 input bits, thus requiring 4 times as many storage entries (because each extra input bit to a table doubles the number of possible addresses), each twice as long as for the separate tables (since twice as many logic equations must have their outcomes stored at each address.) Significantly more memory addresses would be required for a single table with more input bits than for two parallel lookup tables, thus the latter is the (more efficient) way we propose to implement two multiplicity choices.

In the preceding discussion of the GL1 Logic Box, the 3-bit word Mode had the value Mode=0 as appropriate for "real" data. The other seven values of Mode would be used for other functions, such as monitoring, reporting settings, diagnostics, test patterns, etc. It is useful to think of Mode as a "routing" bit. This could be implemented by generating a 1-bit word for the question "Mode=0?" and requiring this word to be true in the real-data lookup tables, as was done in the above sample tables. For non-zero values of Mode, a different set of lookup tables would be in use. This includes the lookup tables of the "discriminator step" as well as the trigger lookup tables and their prescales. For the trigger lookup tables, one could envision three layers of the lookup-table logic shown in Fig. 11.10, with the top layer being Mult-1 tables, the second layer being Mult-2 tables, and the third layer being the non-zero Mode lookup tables.

In order to discuss the nonzero Mode case, it is useful to specify sample values for the seven non-zero values of Mode. (However, at the present time these are tentative, as is making Mode a 3-bit word.) Each of the input boxes in Fig. 11.10 has its own Mode. The following are sample values for the generic Mode for any subdetector:

1. Mode=0 for taking real data
2. Mode=1 for monitoring real data (i.e., data to online monitors)

3. Mode=2 for reporting settings (discriminator levels, etc.)
4. Mode=3 for reloading settings (changing levels, etc.)
5. Mode=4 for running an "all 0s" test pattern (tests summers)
6. Mode=5 for running an "all 1s" test pattern (tests summers)
7. Mode=6 for taking a test pattern using stored test pattern (tests all logical decisions)
8. Mode=7 for running diagnostics (when tests reveal a problem)

Each LVL-1 input has its own Mode and can be checked independently of the others. At present it is not yet clear how, or if, these tests will be combined with front-end electronics tests or calibration runs. However, in the design of the GL1 Logic Box, allowance must be made for a separate track of lookup tables to accomplish whichever type of mode functions are desired. This should be possible with a third layer of lookup tables specifically for the mode operations. The number of input bits to these tables would be about the same as those discussed above for real data. To illustrate this, consider again the γ -1 trigger type. To check out this LVL-1 input, it would not be necessary to input either Mult-1 nor Mult-2 in order to check the operation of the LVL-1 trigger electronics for this signal from the EMCal. However, the 3-bit word Mode would be input, so the number of input bits would be about the same. Our present planning is to use the same size lookup table for each layer and each trigger type, with this size determined by the largest table needed. At present a lookup table with 20 input bits is being considered, which has a few bits to spare.

The GL1 Logic Box will produce a mask of trigger bits, one for each trigger decision. This mask, as well as all inputs to and all outputs from each lookup table and each prescaler, will be part of the data stream. (The inclusion of these in the data stream is not shown explicitly in Fig. 11.10.)

11.5 Data Acquisition System

11.5.1 Overview of Data Acquisition System

The PHENIX experiment is designed to study nuclear collisions at RHIC. Each collision produces particles that we measure in the various PHENIX detectors. The set of data produced by all detectors in response to the particles created in such collisions will be called an event. The beams in RHIC are bunched so that collisions occur at regular intervals when a bunch from one beam crosses a bunch from the other beam. These bunch crossings occur at a frequency of 8.9 MHz with a corresponding time between crossings of 112 ns. A signal called the **BeamClock** will be generated to indicate when each bunch crossing occurs.

The probability that an interaction occurs when two bunches cross is determined by physics and the spatial size and intensity of the two beams. For our current purposes, all of the physics of the nuclear interactions can be summarized with two quantities, the cross-section for interactions and the particle multiplicity. The important beam characteristics are

contained within one quantity, the luminosity. The product of cross-section and luminosity gives the rate of nuclear collisions, while the amount of data that must be collected for each collision is roughly proportional to the multiplicity, with the proportionality constant determined by the amount of data produced by the various PHENIX detectors. The product of collision rate and event size gives the PHENIX DAQ data rate. Details of the calculation of RHIC luminosity, collision cross-section, collision multiplicities, and resultant data rates are given in the PHENIX CDR and references therein. The results of these calculations can be summarized by the following statements:

1. The collision frequency depends on the type of beam particle and varies from a maximum of 454 kHz in proton-proton interactions to a minimum of 14 kHz in Au-Au interactions.
2. The event size also depends on beam particle type and varies from a minimum of 1.9 kB/event in proton-proton collisions to a maximum of 305 kB/event in Au-Au interactions.
3. The data rate is roughly independent of beam particle type and varies from 0.9 GB/sec for proton-proton collisions to 1.0 GB/sec for Au-Au collisions.

The above raw data rates are clearly higher than can be written to tape. An estimate of reasonably achievable tape writing speeds (see below) suggests that the raw data rate is too high by a factor of 100. This means that the experiment cannot write every event to tape and must trigger on desired events. The trigger must be able to reduce the raw data rate to the tape writing rate and should enhance the number of "interesting" events written to tape relative to the non-interesting events, based on some physics considerations.

The PHENIX DAQ system can, then, be logically separated into two parts, the **trigger** and **readout** systems. The trigger system is responsible for processing data from the detectors with the goal of deciding whether or not an event should be written to tape, and the readout system is responsible for moving data from the detectors, combining this data into complete events, and writing the assembled event to tape. These two systems are interconnected at several points because in order to make a decision to keep or reject an event, the trigger system needs a subset of the data obtained by the readout system. This data that is provided for trigger processing is generally referred to as **trigger primitives**. Because the interactions that PHENIX wishes to study are very complicated, the calculations necessary to determine whether or not an event is accepted are involved and may require several stages of computation and decision.

The PHENIX detectors produce signals that must be processed and digitized before the data are written to tape. This processing may include discrimination of pulses, charge integration, and time measurement. An important step in this processing is the conversion of analog data to digital form. In principle, all analog signals could be digitized on each bunch crossing, but this solution requires expensive flash-ADC's and large amounts of memory to hold the resulting digital data. Alternatively, the analog signals could be stored in analog memory units (AMU's) until all trigger decisions have been made, at which point the analog data would be digitized. However, in order for the trigger to perform its calculations, some

Figure 11.11: Overview of the PHENIX data acquisition system. The figure is greatly simplified to illustrate only the major flow of data. Some of the items omitted for clarity are: distribution of the BeamClock, buffers for intermediate storage and time alignment, and monitoring and control paths.

fraction of the data would have to be digitized, requiring an alternate data path. PHENIX has chosen to take an intermediate path, which is illustrated in Figure 11.11. Most of the analog signals are stored in AMU's until a first level (LVL-1) of trigger decision is made. The signals are then digitized, processed, and buffered and also made available for use in further trigger calculations. If this second level (LVL-2) of trigger processing decides the event should be retained, then the digitized data from all detectors is assembled in the **Event Builder** (EB). The event builder provides a stream of complete events which may be written to tape. The event builder also contains enough intelligence to allow another round of trigger decision based on sophisticated algorithms that search for complex physics signatures in the events.

This structure for the DAQ system does require that some data be flash-digitized for use by the LVL-1 trigger. This flash digitization occurs on every bunch crossing, and in order for the trigger system to keep up with the flash-ADC's, the LVL-1 trigger must operate as a pipeline. This means that the trigger processing must be broken up into several short steps, each of which is finished within one bunch-crossing time. On each bunch crossing, the results from one step are passed to the next stage and calculation is restarted on data from the previous stage. This design allows the trigger processing to run continuously without

Figure 11.12: Logical flow of data in the LVL-1 trigger pipeline.

ever having to stop. While the LVL-1 calculation is being performed, the event data must be buffered in digital or analog memories that also function as pipelines since each bunch crossing provides new analog data that must be kept synchronized with the LVL-1 calculations and the LVL-1 trigger decision. The operation of both the trigger and memory pipelines are synchronized to the beam clock because each bunch crossing provides new data and thus requires an empty bucket in the pipeline.

The logical flow of the LVL-1 pipeline is shown in Figure 11.12. Each tick of the Beam-Clock (i.e., each bunch crossing) causes the signals from all channels in all detectors to be stored in either analog or digital memories. The digital data are obtained from flash-digitizers which may run at a frequency higher than the BeamClock. The analog or digital memories are located on either the detector itself (e.g., the MVD), or more commonly, off-detector on the Front-End electronics Cards (FEC's). The analog and digital memories used to store the event data do not actually transfer data on each bunch crossing, but mark the cells in which data for a given bunch crossing is stored as "occupied" until the LVL-1 decision has been made.

A subset of the detector outputs is provided to the LVL-1 trigger decision logic. The LVL-1 trigger uses a combination of flash-digitized signals and discriminator outputs to perform its calculations. Some of the detector outputs undergo analog processing (e.g. summing) before the digitization stage in order to reduce the number of computations that the LVL-1 trigger must perform. The LVL-1 trigger processing is broken up into two stages. In the

first stage, Local LVL-1 (LL1), data from a single detector are reduced and a limited set of output data is generated. In the second stage, Global LVL-1 (GL1), the data from the various detectors are combined and used to make a decision on whether or not the event is to be kept or dropped. The LVL-1 data propagates through the LL1 and GL1 logic in synchronism with the BeamClock, so that the LVL-1 decision for a given crossing appears at a fixed number of BeamClock ticks after that crossing occurred. This time is referred to as the LVL-1 latency. The LVL-1 decision is distributed to all FECs, synchronized with BeamClock. A **LVL-1 Reject** causes the FECs to release the memory space occupied by the data from that crossing. A **LVL-1 Accept**, on the other hand, causes the analog signals to be digitized and formatted for the next level of processing. The maximum time from a LVL-1 Accept until all data for a given event are digitized and formatted (loosely referred to as the **Digitization Time**), has been chosen to be $40\mu\text{s}$.

Digitized data from the FECs are transferred to **Data Collection Modules (DCMs)**. There, data are zero-suppressed, corrected, and reformatted. The data corrections may include adjustments for AMU pedestal or gain variations and calculation of "physics" quantities. The DCMs may also produce a reduced set primitives for use in the LVL-2 trigger. After the DCMs, the data from a given detector are buffered and merged onto a small number of cables which are the inputs for that detector to the Event Builder.

The Event Builder receives the data fragments from each of the detectors and assembles them into complete events. To achieve the required bandwidth, the processors which perform this assembly are deployed as a farm, i.e., several independent legs processing "successive" events in parallel. The word successive here is used loosely: the processing time per event in the EB may vary, so a given event is assigned to the next available leg, rather than via a deterministic algorithm. Because the Event Builder is the first place the entire event is assembled, it is a logical place to implement higher level triggers (Level-2.5 and higher) which rely on physics signals determined from the correlations between many sub-systems. Survivors of the combined processes (Event Building + Level-2.5 trigger) are written as complete events to mass storage. The design throughput for archiving is 20 Mb/s.

The flow of the data after digitization still proceeds in a pipeline fashion but it is no longer synchronized to the beam clock, but is instead event-driven. Each stage of processing fetches a data fragment, processes it, and stores the result in an output buffer where it will wait to be fetched by the next processing stage. At each stage, many operations are proceeding in parallel with no "lateral" synchronization. At each stage in the DAQ the alignment of data fragments from the various parallel processors is performed using FIFO buffers. These buffers also serve to absorb fluctuations in the arrival time of events.

The DAQ system described above is said to be data-driven and deadtimeless. Data-driven implies that the data contain control and format fields which provide all routing information necessary to move that data through the various stages of processing. In other words, the data "push" themselves through the system, instead of being actively routed by a control system.

The pipeline flow of the LVL-1 trigger is said to be deadtime-less in that acquisition of new data from new bunch crossings proceeds at the same time trigger decisions are being made and data is being digitized for preceeding bunch crossings. The extent to which this

property is guaranteed for later stages of the data acquisition system is determined by the depth of the readout buffers and the LVL-1 accept rate. LVL-1 trigger rates in excess of the mean throughput guarantee deadtime regardless of the buffer depth while for LVL-1 rates less than the mean throughput, a modest buffer depth (e.g. 5 events' worth) provides deadtimes of 10% or less. The remainder of the DAQ system is also deadtime-less within the constraint that no buffers in the system become full and unable to accept data from the output of a processing stage. Since a few such overflow's are inevitable due to event size and rate fluctuations, busy signals must be provided at every point in the DAQ system to halt processing further upstream in the pipeline until the blocked buffer has emptied sufficiently.

11.5.2 Data Formats

The data-driven architecture we have chosen depends in an essential way on the efficient coding of address, routing and control information within the detector data itself. The general scheme relies on separating all data "words" into two fields, a **Label** field and a **Data** field. The Label field describes the contents of the Data field. Here "Data" is used in a general sense - the Data field can contain information such as event number, packet identifiers, etc., in addition to items such as pulse height or time. In each case, the interpretation, routing, and processing of this information depends on the contents of the Label field. It is important that this encoding is done as compactly as possible, since the required width in bits directly affects the size of all data cables in the system.

The data is first configured into the above specified format by the FECs. The output of these cards is a string of 16-bit words, with 4 bits and 12 bits allocated for the Label and Data fields, respectively. We refer to this as the "Raw" data format. The raw data stream is received by the DCMs which in turn produce data in the "Corrected" format with 8 bits of Label field and 16 bits of Data. These two formats are described in further detail below, using the TOF detector as an example.

Raw Data

This data format describes the structure of the data immediately after digitization but before zero suppression, i.e. the format of the data transferred from the FECs to the DCMs.

```

/*****/
/*
/* Purpose   : Header file for TOF Data Formats
/*
/* Author    : WAZ, 03-May-93
/*
/* Variable
/* Names     : Standard prefixes used:
/*             b      Bit variable (0 or 1)
/*             bf     Bit Field (subset of int, longword, etc.)
/*             i      Integer, especially indices

```

```

/*          li          Long integer          */
/*          ui          Unsigned Int          */
/*
/*      TOF FORMAT, Raw Data
/*
/*      +---+---+ +---+
/*      |1|1|1| |0|0| bfEvent(1)
/*      |1|1|0| |1|0| bfEvent(2)
/*      |1|1|0| |0|0| bType
/*      |0|0|1| |0|0| bDisc
/*      +---+---+ +---+
/*      |Q|Q|T| |C|E|
/*      |1|0| | |E|V| bfData
/*      | | | | |L|E| (12 bits)
/*      | | | | |L|N|
/*      | | | | |T|
/*      | | | | | |
/*      +---+---+ +---+
/*      |
/*      |
/*      Repeated
/*      for all
/*      channels -----> Data Flow
/*      on card
/*****

typedef unsigned int ui      ;
typedef struct   rawTOF{
    ui   bfEvent : 2 ; /* 0 = Event, 1=CellNum, 2=TVC, 3=ADC */
    ui   bType   : 1 ; /* 0/1 for TVC/ADC data */
    ui   bDisc   : 1 ; /* 0/1 for discriminator not fired/fired*/
    ui   bfData  : 12 ; /* 12 bits reserved for raw data */
}      ;

```

Note that the Event Number is limited to 12-bits, and therefore obviously represents some sub-counter. The true event number for a given run will be formed elsewhere by combining this sub-counter with a "frame number" maintained by the master clock.

Corrected Data

This format describes the structure of the data produced by the Data Collection Modules.

```

/*****
/*
/* Purpose : Header file for TOF Data Formats
/*

```

```

/*                                                                 */
/* Author   : WAZ, 03-May-93                                     */
/*                                                                 */
/* Variable                                                                 */
/* Names    : Standard prefixes used:                           */
/*          b      Bit variable (0 or 1)                         */
/*          i      Integer, especially indices                  */
/*          li     Long integer                                  */
/*          ui     Unsigned Int                                  */
/*                                                                 */
/*                                                                 */
/*          TOF FORMAT, Corrected Data                           */
/*                                                                 */
/*          -----> Data Flow                                   */
/*                                                                 */
/*          +-+..+-----+-----+-----+-----+-----+-----+
/*          |0| |1|1|1|1|0|0 |1|1|1|1|1|0|0 1|1| bfHeader(1) */
/*          |1| |1|1|1|1|0|0 |1|1|1|1|1|0|0 1|1| bfHeader(2) */
/*          |0| |0|1|0|1|0|0 |0|1|0|1|0|0|0|0|0| bfType(1)  */
/*          |0| |0|1|1|0|0|0 |0|1|1|0|0|0|0|0|0| bfType(2)  */
/*          |0| |1|0|0|0|0|0 |1|0|0|0|0|0|0|0|0| bfType(3)  */
/*          |0| |0|0|0|0|0|0 |0|0|0|0|0|0|0|0|0| bfType(4)  */
/*          |0| |0|0|0|0|1|0 |0|0|0|0|1|0|0|0|0| bfAddr(1)  */
/*          |0| |0|0|0|0|0|0 |0|0|0|0|0|0|0|0|0| bfAddr(2)  */
/*          +-+ +-----+-----+-----+-----+-----+
/*          |W| |Q|Q|T|T|G|G |Q|Q|T|T|G|G|D|E|
/*          |o| | | | |e|e| | | | |e|e|a|V| bfData
/*          |r| |p|p|R|L|o|o |p|p|R|L|o|o|t|E| (16 bits)
/*          |d| |o|r|i|e|A|A| |o|r|i|e|A|A|a|N|
/*          |C| |s|e|g|f|d|d| |s|e|g|f|d|d|I|T|
/*          |o| |t| |h|t|d|d| |t| |h|t|d|d|D| |
/*          |u| | | |t| |r|r| | | |t| |r|r|W|N|
/*          |n| | | | | | | | | | | | |o|u|
/*          |t| | | | |#|#| | | | |#|#|r|m|
/*          | | | | |2|1| | | | |2|1|d| |
/*          +-+..+-----+-----+-----+-----+
/*                                                                 */
/*****

```

```
typedef unsigned int ui ;
```

```

typedef struct corTOF {
    ui    bfHeader: 2 ; /* 0 ==> "Data" is address info */

```



```

/* 1 ==> "Data" is EventNum */
/* 2 ==> "Data" is WordCount (EOD) */
/* 3 ==> Interpret "Data" as per below */
/* +-----+ */
ui  bfType : 4 ; /* | This field is detector specific | */
/* | as per value of DataIDWord; | */
/* | all other fields are valid | */
/* | PHENIX-wide | */
/* | | */
/* | 0 ==> "Data" is DataIDWord | */
/* | 1 ==> "Data" is T_left | */
/* | 2 ==> "Data" is T_right | */
/* | 3 ==> "Data" is Q_pre | */
/* | 4 ==> "Data" is Q_post | */
/* | 5 ==> "Data" is T_sum | */
/* | . ==> . . . | */
/* | . ==> . . . | */
/* | . ==> . . . | */
/* +-----+ */
ui  bfAddr : 2 ; /* 0 ==> "Data" is geo. add. (1st wd) */
/* 1 ==> "Data" is geo. add. (2nd wd) */
/* 2 ==> "Data" is logical add. (1st wd) */
/* 3 ==> "Data" is logical add. (2nd wd) */
ui  bfData : 16 ; /* 16 bits reserved for corrected data */
} ;

```

The following should be noted:

Type-Checking and Redundancy: The method outlined here has two useful features:

- Not all 2^8 Label patterns are valid. For instance, if `bfHeader` $\neq$ 0, then values in `bfAddr` have no meaning. This provides some weak data integrity checking, and allows for simple 1 or 2-bit logic to determine distribution/action of a given word (i.e., is it data or address?).
- There is a method available of near-infinite extensibility, via the `DataIDWord`. When `bfHeader` = 3, the corresponding Data field contains a descriptor which further labels the contents of the packet. There are 2^{16} possible values for this sub-label, so that one can use this general scheme for all sub-detectors; the value of `DataIDWord` is used to determine the meaning of the following words labeled with the `bfType` field.

Addressing: Explicit provisions have been made for two methods of encoding addresses: Both *Geographical* $\equiv$ (*Rack, Crate, Module, Channel*), and *Logical* $\equiv$ (*Arm, Wall, Row, Column*) modes are supported, depending on the value of `bfAddr(2)`. While in principle only 16 bits should be necessary to describe any sub-detector channel number in PHENIX, two words (32-bits) have been reserved to allow for upgrades.

DataIDWord: The `DataIDWord` is encoded using the sub-fields below:


```

typedef struct DataIDWord {
    ui    bfClass      : 4 ; /* 0 = L0 data          */
                                /* 1 = L1 data          */
                                /* 2 = L2 data          */
                                /* 3 = L3 data          */
                                /* . = . . .          */
                                /* . = . . .          */
                                /* . = . . .          */
                                /* 13 = Calibrations    */
                                /* 14 = Processed data  */
                                /* 15 = Raw data        */
    ui    bfSubDetector : 12 ; /* Sub-detectors bit-mapped according to: */
                                /* Bit 1: BeamBeam      */
                                /* Bit 2: MVD-Strip     */
                                /* Bit 3: MVD-Pad       */
                                /* Bit 4: DC            */
                                /* Bit 5: PAD           */
                                /* Bit 6: RICH          */
                                /* Bit 7: TRD           */
                                /* Bit 8: EMCal         */
                                /* Bit 9: MuTrack       */
                                /* Bit 10: MuID         */
                                /* Bit 11: Reserved    */
                                /* Bit 12: Reserved    */
} ;

```

Although this bit-mapping of the various sub-detectors does not make economical use of the potential 2^{16} values of the DataIDWord, it is thought that this is offset by its logic and flexibility for labeling arbitrary packets of information assembled in the system, e.g., a specific LVL-2 packet might provide electron candidate primitives with Bits 5, 6, 7, and 8 of the bfSubDetector field set. Note also that there is a finer level of description possible using detector-specific values of the bfType field in each data word. For example, one could further subdivide EMCal into Pb-Sc and PbGl.

11.5.3 DCM Cards

The basic tasks of the DCM are

- Remove empty channels from the data flow. (zero-suppression)
- Convert the hit information from channel-dependent values to physical values (wherever possible).
- Format the remaining data (hit channels) into the form required for downstream devices.

- Produce primitives for the LL2 trigger processors.

In addition, we may also want to use the DCMs to monitor (e.g. histogram) the data stream, to calibrate AMU cell-to-cell pedestal and gain variations and to determine ADC pedestals and timing offsets during calibration runs.

LVL-1 data is needed in the DCMs for zero-suppression or calculation of physical quantities. For example, time-of-flight data must be corrected for the vertex position and the interaction T_0 provided by the Beam-Beam LL1 processing. As another example, for events that satisfy an EMCal LVL-1 trigger, one might decide to always read out a cluster of EMCal channels regardless of their contents based on the data produced the EMCal LL1 processing. Since the processing algorithms may change and may require access to LVL-1 data that was not initially anticipated, the distribution of LVL-1 data to the DCMs should not be determined by cable connections, but instead by intelligent routing and broadcasting. This will be discussed further in section 11.5.3.

The DCMs are being designed to handle 64 channels of data for a typical system, assuming three pieces of data (e.g. 2 charge and 1 time sample for TOF) per channel per event. In our current design, these data are presented to the DCM at a rate of one sample per BeamClock tick, so the total time required to receive the data from the FECs at the DCM is $192 \times 112 \text{ ns} = 21.5 \mu\text{s}$, comfortably within the $40 \mu\text{sec}$ we have chosen as the digitization time. A count of the number of DCM cards required is shown in Table 11.8.

Figure 11.13 shows an implementation of a DCM with the desired functionality. The module consists contains three FIFO (First-In First-Out buffer) ports, a FPGA (Field Programmable Gate Array), RAM for a lookup table, a DSP (Digital Signal Processor), RAM for the DSP, and dual-ported RAM for an output buffer. The FIFO ports are used to buffer and time-align digitized data and LVL-1 information. The zero-suppression is performed by the FPGA in conjunction with the lookup tables. The DSP is used to download constants into the tables, to correct the data, to generate primitives for the higher-level triggers and to format the output data. The following sections will discuss each element and its function in more detail. We have used the requirements of the TOF readout to develop our description of the DCM and will use the TOF example extensively in describing the DCM operation. Its specifications are sufficiently general to be useful elsewhere in the system.

FIFO Ports

Each DCM has three FIFO ports, two data ports (A & B) and the LVL-1/Control port. The two data ports are fed by front-panel cables, while the LVL-1/Control port is connected to the backplane bus. The two data ports allow one to either combine data from physically separate portions of a detector (e.g. to form the coincidence between ends of TOF slats) or to simply double the density of channels addressed by one DCM. The backplane distribution of LVL-1/Control data satisfies the requirement that the LVL-1 data distribution not be determined by cable connections. The FIFO's serve to synchronize (align) the data received from the three different sources. They also buffer new data received from the FECs while an event is being processed in the DCM.

Our current design calls for a FIFO depth of 4K (16-bit) words, which provides space for $4096/192 \approx 21$ events, which is more than adequate to absorb fluctuations in the LVL-1

TRY2_1.XLS

Detector	channel	samples per channel	16 bits word per samples	data volume 1 in 16bits word	address tag	data volume 2 in 24 bits word	scale factor A	DCM A modules	DCM A crate	scale factor B	DCM B modules	DCM B crate
beam-beam	132	2	1	264	1	396	66	2	1			
MVD vertex	28,672	2	1	57,344	1	86,016	512	112	3			
MVD pad	6,114	2	1	12,228	1	18,342	512	24	1			
DC	10,240	20	1	204,800	1	215,040						
PC	46,400	3	1	139,200	1	185,600	384	363	10			
TEC	28,030	80	0.3625	812,870	1	840,900				64	438	11
EMCAL	25,536	3	1	76,608	1	102,144	384	200	5			
RICH	6,400	3	1	19,200	1	25,600	384	50	2			
TOF	2,048	3	1	6,144	1	8,192	384	16	1			
MUON T	19,680	12	1	236,160	1	255,840						
MUON ID	18,360	2	1	36,720	1	55,080	512	72	2			
Level 1												
Level 2												
TOTAL	191,612			1,601,538		1,793,150		839	25		438	11

Page 1

Table 11.8: Data volume load into the DCMs and number of DCMs needed as a function of subsystem.

accepted trigger rate as long as the average rate is less than the chosen maximum LVL-1 accepted trigger rate of 25 kHz. By deliberately choosing a very deep FIFO, we can use the the “partially full” signal from the FIFO to form the **System Busy** signal, thereby alleviating potential timing problems that would arise if System Busy only becomes TRUE when a buffer is completely full.

FPGA

The first step in the processing chain inside the DCM is zero-suppression. There are many possible algorithms for zero-suppressing data in the various PHENIX detectors, but the TOF-wall and other AMU-based detectors such as the MVD, PAD, and EMCAL will likely use either a threshold on the channel contents or LVL-1 data or a combination of the two to determine which channels contain useful information. Because the zero-suppression algorithm varies from detector to detector and may change in the future, the hardware performing this correction must be reprogrammable. However, we also feel that use of a DSP for this purpose is not currently feasible because to process 192 signals in the 40 μ sec digitization time with a DSP clocked at a typical 50 ns would allow only four operations per signal. This is only marginally adequate for our needs and would require an additional DSP on the DCM

Figure 11.13: Components of a PHENIX Data Collection Module (DCM).

board, else there would be no processing time available to perform the other tasks demanded of the DCM.

Since the data from each front end card in a given sub-system has a fixed format (see Section 11.5.2), the zero suppression is most efficiently performed using a programmable pipelined logic array, such as a FPGA. The FPGA will have the ability to access a lookup table that is also accessible from and downloaded by the DCM DSP. This lookup table may provide thresholds, correction constants, storage for LVL-1 data, and/or logical decisions to be used in the zero-suppression algorithm. For example, to implement a threshold-based zero-suppression algorithm, we must subtract an AMU cell dependent pedestal and compare to a threshold. This algorithm requires two data words to be fetched from the lookup table (one for pedestal and one for threshold), two "ADD" operations and a check for underflow. For detectors that have pre-samples, the above operation must be performed twice and then the corrected pre-sample must be subtracted from the corrected post-sample before the result is compared to a threshold. This algorithm would require 3 table lookups, 4 subtractions, and two checks for underflow for each channel. Regardless of the algorithm, when a given channel survives the zero-suppression step, the data associated with the channel (A, B, LVL-1/Control and Address) is passed on to the DSP for further processing.

The data field of the look-up table has been chosen to be 16-bits wide. The address field must be at least 15 bits wide to handle pedestal and AMU cell corrections (6 bits for AMU cell address and 9 for channel number). The selection of the FPGA's will be based on the

Figure 11.14: The interface between the FPGA and DSP in the PHENIX Data Collection Modules. Data arrives from the left and is zero-suppressed in the FPGA. Surviving data are passed on the DSP.

number of available gates and I/O pins. Estimates indicate that we need roughly 4000 gates and 100 pins, which is in the middle of the range of commercially available units.

The operation of the FPGA is controlled by information presented to the LVL-1/Control port. Zero-suppression may be enabled or disabled, and the precise mode of zero-suppression (based on channel contents, discriminator flags or LVL-1 signals) may be selected. The system is sufficiently flexible that zero-suppression can be controlled on a DCM-by-DCM basis or through the use of the lookup table, on a channel-by-channel basis within a single DCM.

DSP

The DCM DSP performs data correction, error checking and general maintenance and support operations. In normal operation the DSP receives the zero-suppressed data from the FPGA. We estimate that of order 10% of the hits in the TOF wall and other similar detectors will survive the zero-suppression stage. Since the empty channels have already been removed, the full power of the DSP can be used to correct and reformat the important (non-zero) data. For example, in the TOF readout, the magnitude of the PMT signal is recorded in PRE and POST samples from an integrator with a decay time constant of 112 ns. The DSP will correct both samples for AMU cell-dependent pedestal variations and TOF channel-dependent gain

variations, subtract the corrected PRE sample from the corrected POST sample, and then correct the resulting "charge" difference for the decay-time of the integrator. The times from the TOF slats are corrected using a similar algorithm except that no decay correction is necessary. The resulting time value is then also corrected for pulse-height slewing using the corrected charge value. The slewing correction is performed using a table look-up rather than a time-consuming analytic correction. After these corrections, the TOF data are suitable for use by the LVL-2 trigger. We have both estimated and simulated that these corrections should require 50-100 operations per measurement. Using typical DSP clock rates, this many operations can be completed within the available time per channel following suppression of empty channels.

The DSP is also responsible for reformatting the corrected data and writing them into the dual-ported memory for export to the EB. As discussed in Section 11.5.2, the data at this point will have an 8-bit Label space and 16-bit Data space. This export will be performed by the DMA co-processor contained within the DSP. The DSP is also responsible for generating and loading the lookup table used by it and the zero-suppression FPGA. Generation of the constants for the local lookup table will be performed by a separate calibration task run in the DSP during setup. By providing sufficiently large RAM for the DSP, it may also be possible to store histograms of calibration and/or monitoring data, in addition to the DSP program, correction constants and lookup tables.

Control and Monitoring

Each crate of DCM's will contain a controller to download code to all DSPs over the back-plane bus, and to program the FPGA's at system startup.

Buffering

Each crate will contain a card containing a minimum of three output buffers. These buffers have a simple structure of a FIFO input followed by a FPGA, which again forces time-alignment and checks for data integrity. The three buffers are connected to the Event Builder, the LVL-2 processors, and the monitoring network.

11.5.4 Event Builder

The PHENIX Event Builder will serve as the final event assembler and the Level-2.5 (LVL-2.5) trigger. As shown in Figure 11.15, the EB is configured as a crossbar. Data from each sub-system arrive at the event buffers associated with the leg(s) for that sub-system. The number of legs reserved for a sub-system is proportional to the anticipated bandwidth, and ranges from 1 to 10. Data from one event arrives at the various sub-system event buffers asynchronously, and is then transferred to the buffers associated with the next available EB horizontal leg. Those events passing LVL-2.5 will be further formatted and passed to a router box for recording on mass storage.

The EB level is the first point in the system that a trigger processor has access to the full data for one event. Note that this data consists of all LL1, GL1, LL2 and GL2 primitives,

Figure 11.15: Layout of the PHENIX Event Builder. Sub-system data arrives from the top. The buffers on each horizontal leg contain the sub-system data for one event. The processors pull the data from the buffers, assemble, format and (if desired) trigger on the event formation.

as well as the event data (raw and/or calibrated) itself. It is likely that LVL-2.5 algorithms use these pre-computed primitives to guide more refined calculations.

The bandwidth for the EB is determined by requiring that it must be able to absorb all the zero-suppressed data generated at the LVL-1 trigger rate for a given species, i.e., with the aid of a LVL-2 reduction. As our baseline system, we have taken a Au-Au minimum bias interaction rate of 14 kHz and an assumed event size for these interactions of 75 kB, leading to the requirement that the EB must accept 1.1 GB/sec.

Sub-Event Memories

Processors

The memory configuration of the EB processors is a direct reflection of the desired processing, formatting and transmission of the data. Its hardware configuration must be kept as simple as possible, while allowing the software to operate with minimal overhead in pointing to requested data segments.

A first pass at balancing these requirements leads to the configuration below:

- **Raw Data Bank:** This area contains an image of the received data, whose format

has already been designed for optimal compression and transfer.

- **General Pointer Bank:** A fixed-size bank devoted to pointers to the various sub-systems data. This bank's location is also fixed, so that the sub-system pointers have fixed values, permitting rapid access.

The general structure for this bank is

- Event Number
 - Run Condition(s)
 - Trigger Condition(s)
 - Error Code(s)
 - Word count
 - Starting address for raw data from detector N.
 - Starting address for decoded data from detector N.
 - Word count for data from detector N.
 - * Detector N, Region M starting address
 - * Word count for data from Detector N, Region M
- **Decoded Data Bank:** This bank contains data stripped of the routing labels used in data transmission, and stored in close to off-line format. This required size follows from the assumed event volumes found in Table 11.9.
 - **Geometric Pointer Bank:** Provides the map between a given detector channel and the corresponding hit numbers (if any) in the Decoded Data Bank. Typically the detector channels are configured in a (Row,Column) format corresponding to their layout in real space. The contents of a given (Row,Column) address is the offset address from the starting address of the decoded data for a given sub-system.

Two methods have been used to estimate the memory needed for each processor in the EB. The first assumes that each processor has sufficient memory to contain one *non-zero-suppressed* event (8 MB), plus the program memory, estimated to be 4 MB. The second approach requires that each processor's memory can store two Au-Au central zero-suppressed events, plus their associated decoded banks and program memory. Assuming that the decoded banks are twice the size of the raw data, that Au-Au central leads to 10% occupancy (a generous upper limit), and that the 192K of electronic channels requires 800 kB for the corresponding geometric pointer bank, one obtains $2 \times [0.1 \times 8 \text{ MB} + 2 \times 0.1 \times 8 \text{ MB} + 0.8 \text{ MB}] + 4 \text{ MB} \sim 10.4 \text{ MB}$, which is similar to the previously obtained 12 MB. Based on these considerations, we have chosen 16 MB as the design specification for the processor memories, i.e., 8 MB data, 4 MB program, and 4 MB reserved for absorbing fluctuations in event size.

Figure 11.16: Organization of the processors and memory in the PHENIX Event Builder.

Processor Units

As shown in Figure 11.16, the EB processor is logically divided into four blocks: the (previously discussed) memory, the main processor unit, the I/O unit and a co-processor.

The I/O unit controls all access to the memory and communications between the processor and co-processor, interfaces these elements to the backplane bus, and refreshes dynamic memory. Data from the buffers will be pushed into the memory through the I/O units once the processor signals it is ready to accept a new event. Upon completion of this process the I/O unit interrupts the processor, which then in turn notifies the co-processor to begin the decoding process.

Upon receipt of the signal from the processor, the co-processor begins the translation of the raw banks into the decoded data format. This process requires the starting address, word count and format of the raw data samples, which is also supplied by the processor. The processor/co-processor arrangement is designed to optimize the throughput and flexibility of the system: The processor (chosen to be a Texas Instruments DSP320C40) supplies the sub-system specific information and memory pointers to the co-processor, which is a gate array or FPGA configured to perform the required translations. This arrangement is made possible by the common data format adopted by all sub-detectors (see Section 11.5.2).

Buffer Card

TRY3_1.XLS

Detector	channel	No. of sample per channel	16 bits word per samples	address tag	data volume in 24 bits word	2 sample per hits	Average event size (24 bits) / Au-Au min. bias	occupancy for Au-Au min. bias	Total data bandwidth from Au-Au min. (24 bits) per sec	min. event builder legs	event builder legs	min. event builder size	Average event size (24 bits) for Au-Au central
beam-beam	132	2	1	1	396	1	396	100.00%	5,142,857	1	1	1,584	396
MVD vertex	26,672	2	1	1	86,016	2	4,220	4.91%	54,805,058	4	5	16,880	16,880
MVD pad	6,114	2	1	1	18,342	2	1,076	5.87%	13,975,325	1	2	17,216	4,304
DC	10,240	20	1	1	215,040	2	4,431	2.06%	57,545,455	4	5	17,724	17,724
PC	46,400	3	1	1	185,600	3	2,216	1.10%	28,772,727	2	3	17,724	8,862
TEC	26,030	80	0.3625	1	840,900	80	8,862	1.05%	115,090,909	7	10	20,256	35,448
EMCAL	25,536	3	1	1	102,144	3	2,068	2.02%	26,854,545	2	3	16,542	8,271
RICH	6,400	3	1	1	25,600	3	92	0.36%	1,198,173	1	1	1,476	369
TOF	2,048	3	1	1	8,192	3	98	1.20%	1,271,481	1	1	1,566	392
MUON T	19,680	12	1	1	255,840	2	479	0.19%	6,214,909	1	1	7,657	1,914
MUON ID	18,360	2	1	1	55,080	2	266	0.48%	3,452,727	1	1	4,254	1,063
Level 1							1,000		12,987,013	2	2		1,000
Level 2													
TOTAL	191,612				1,793,150	25,203			327,311,179	27	35		96,624

Aver. Au-Au central event size/no. of event builder legs
 (No. of data words * 56ns transfer clocks * 1.5) / 77 microsec
 (No. of data words * 56ns transfer clocks) / 77 microsec
 Average Au-Au min. bias event size * Au-Au min. bias interaction rate (77 micro-sec)
 Au-Au min. dndy*detector coverage*layers*hits per layer* (samples per hits*word per samples+Address tag)

Page 1

Table 11.9: Data volume load into the Event Builder as a function of subsystem.

The buffer cards located at the nodes of the EB act as FIFO's to synchronize the movement of data fragments from the sub-event buffer to a given EB leg. The buffer sizes are determined by the amount of data expected on a given data cable for Au-Au central events, plus a safety factor, as determined by Table 11.9. Each buffer card will consist of 64 KB of static memory plus a control unit to enforce FIFO operation. This small size is adequate for zero-suppressed readout, but will lead to significant bandwidth reduction when reading out non-suppressed data during diagnostics runs, since each buffer cannot absorb the entire data volume from a given sub-system in one transfer, which necessarily in turns blocks access to the "next" available output leg by a second event until the first one is absorbed by the output leg used by the first event.

Physical Implementation

The EB physical layout is shown in Figure 11.17. All elements reside in FutureBus+ 12SU crates, with 128-bit wide data backplanes. The buffer cards are fed by twisted-pair ribbon cable from two crates of sub-event buffers. Each buffer card has four front-panel connectors, two Data-In and two Data-Out. The two Data-in ports accommodate two independent vertical legs of data input from the sub-event buffers. The Data-Out ports are arranged to create a daisy-chain of buffers for each input leg spanning the several crates of Event Builder buffers. The buffer cards use the FutureBus+ backplane as the horizontal EB "output" leg

Figure 11.17: Physical layout of the PHENIX Event Builder. Digitized data from a set of channels arrives from the top and is zero-suppressed by the corresponding DCM. The results for each crate are stored in local buffers, then presented to the daisy-chained data cable.

within the crate, with crate-to-crate links provided via ribbon cable. The 128-bit backplane width is subdivided into 4 32-bit horizontal legs, with each leg ultimately leading to a corresponding set of processors. Thus, each Event Builder buffer card contains $4 \times 2 = 8$ buffers. We have chosen to implement four processors per processor card to match the above division of the FutureBus+ backplane datapath.

The vertical input legs will consist of 24-bit cable, corresponding to the (8-bit Label + 16-bit Data) format previously discussed. The horizontal legs will consist of 32-bit wide cables. All data transfers will be clocked at 20 MHz. In its optimal configuration, the number

of EB output legs should equal or exceed the number of input legs. Since the number of input legs (40) is determined by the amount of data generated by the individual sub-systems for Au-Au minimum bias rates (when transferred at 20 MHz), the nominal configuration of the EB should consist of 40×40 legs, with 25 processors per output leg.

L2.5 Trigger and Control

11.6 On-line Computing System

11.6.1 Overview

Requirements

The on-line computing systems (ONCS) will provide a standard framework for developing and operating the integrated PHENIX facility from the standpoint of computing. This includes software and hardware for a kernel operating environment, public protocols and interfaces, and tools and templates for common efforts. These features will be embodied in a set of component systems for development, run control (including detector and readout configuration), event-data flow control, event-data archiving, monitoring (including event analysis), calibration, diagnostics and maintenance, and interfacing to ancillary systems. These component systems will be layered on a set of generic services providing a common approach to user interfaces, database access, and message passing.

Validating and monitoring the physics performance of the PHENIX detector, an essential ingredient in guaranteeing high-quality physics data, requires high-performance on-line computing facilities, given the complexity, size, high data rates and large data volumes of the detector, and the computationally intensive analyses expected for RHIC events. The long running periods at RHIC make this function even more critical, so quality and reliability in the on-line systems are especially important.

The breadth of the PHENIX physics program, new physics that may be uncovered, possibilities for future upgrades, likely advances in data acquisition technologies, and the long lifetime of the detector all imply *evolving* requirements for on-line computing. It is therefore crucial that on-line computing systems be designed from the outset to be flexible and readily extendable, by stressing an open architecture and implementation strategies based on widely supported standards.

System Model

A model of the on-line computing systems is useful for mapping requirements into classes of relatively independent functions and services, and to help define boundaries with other PHENIX and RHIC systems. A simplified data flow diagram for a first pass at such a model appears in Fig. 11.18. This approach divides on-line computing into two overlapping domains of functions: (1) an "on-line framework" that provides an interconnected software and hardware environment that will be the basis for implementing and integrating on-line operations, (2) a generic set of PHENIX services that can be used to both build the framework and implement specific PHENIX subsystem applications. The framework includes:

Figure 11.18: Simplified Functional Model of On-line Computing.

run control, ancillary systems (see CDR section 11.6), monitoring (including on-line event analysis), diagnostics, event-data flow control, event-data archiving and calibration. The services provide standard interfaces, protocols, tools, and templates for creating operator/user interfaces, building and accessing various PHENIX databases, and communicating between distributed processes and between software components and users.

Scope and Assumptions

The ONCS software effort will be primarily limited to provision of a kernel operating environment plus standard tools and public protocols and interfaces, as opposed to subsystem-specific "applications." A general framework plus templates for subsystem-specific efforts will also be provided by the ONCS project as a part of the public software. This will include software components required to provide a reasonably uniform systems environment for re-

programmable devices embedded in the Data Acquisition System, to an extent agreeable to the ONCS and DAQ teams. Wherever reasonable and cost effective, standard industry-supported network hardware and software will be specified for ONCS interconnects. Responsibility for developing a PHENIX-standard operator/user interface, databases, and message services will be shared between the PHENIX ONCS and Offline Computing Systems (OFCS) projects.

ONCS hardware will be predicated as much as possible on off-the-shelf products. ONCS tasks dealing with custom components will be limited to those outside the domains of detailed design and fabrication. Components in the latter class include devices in the Data Acquisition (DAQ, which includes event-data readout, triggering, and event building) chain at and upstream of the final subsystem event-fragment buffers, and local slaves to the timing and trigger systems. The hardware implementation will be based on an evolution of the architecture shown in Fig. 11.45 of the CDR. This implies a distributed client-server architecture with standard microprocessors in industry standard frontend crates (for local read-out control, monitor and calibration coordination, and trigger and timing supervision) networked to each other and to general purpose workstations for operator interfaces, file and database services, and general-purpose computing resources.

Because of the long lifespan of the PHENIX project and the large fraction of uptime needed for the detector, quality control will be important during design and implementation. Establishing common standards, tools, templates, and testing schemes to facilitate integration of subsystem-specific hardware and software will be important activities for the ONCS group. The development process we envision is discussed in more detail in section 11.6.5.

11.6.2 Changes

1. The updated CDR Section "On-line Computing Systems" (UCDR 11.6) now incorporates the original CDR sections "Event Archive" (CDR 11.4.3), "Ancillary Systems Control" (CDR 11.6) and "On-line Computer System" (CDR 11.7). Several topics in CDR section "Data Acquisition" (CDR 11.4) relating to the on-line computing systems are now addressed in UCDR Section 11.6.
2. Discussions of software and software development strategies are essentially all new.
3. We present a first crude attempt at a global on-line computing model, to start organizing analysis of high-level requirements and mapping them onto classes of relatively independent functions and services that can facilitate a common approach to development for all on-line computing projects (UCDR 11.6.1).
4. We present a first delineation of a framework for system development, a kernel environment for system integration and operation (UCDR 11.6.3), and classes of generic services (UCDR 11.6.4).
5. The architecture has evolved according to a network-transparent client-server paradigm, influenced by the EPICS model for control and monitoring (UCDR 11.6.3).

6. Ancillary Systems Controls has been subdivided into an Ancillary Systems Interface System and subsystem-specific ancillary systems equipment (UCDR 11.6.3).
7. From the controls and monitoring standpoint the Master Timing, Event-data Flow Control, Event Builder/Level-3 Trigger, and Event-data Archive subsystems are treated on an equal footing with detector and trigger data acquisition subsystems (UCDR 11.6.3).
8. The "routing layers" referenced in the Event Archive and Online Computer System sections of the CDR have been incorporated into the Monitor Network, by considering the latter as a general-purpose high-speed fabric for connecting clients on the workstations to intelligent servers embedded in the Subsystem Control and Monitoring Nodes. (UCDR 11.6.3).
9. We present a first attempt at defining a strategy and plan for system development, keyed to needs of PHENIX subsystem evolution (UCDR 11.6.5).
10. Separate sections for costs and schedules for on-line computing systems have been added.

11.6.3 On-line Framework

Run Control System

The Run Control System starts and stops data acquisition runs in predetermined sequences. It also includes facilities for defining those sequences and for detector initialization, configuration, and general coordination of on-line operation.

Configuration and initialization can be a relatively complex process, as it extends to all subsystems and many levels of the on-line system. Detector partitioning, running modes, and initialization parameters and sequences must be specified either by an operator or via retrieval of archived descriptions. The ability to "partition" the detector is especially important during testing and commissioning, when it will often be desirable to run different detector subsystems independently. Initialization can include retrieval of frontend calibration data and software from a database and loading it into the proper hardware, loading trigger parameters and software, verifying that software and hardware are in known and valid states, etc.

The ONCS team will provide an integrated framework for run control. This framework will include a PHENIX operator interface, configuration generation tools, a configuration database and access mechanisms, tools for specifying initialization and operation sequences, and standard run control protocols and interfaces to all subsystems. Specific detector subsystem application software will be the responsibility of the subsystems groups. Integration of subsystem-specific applications into the on-line computing systems will be supported by the standard tools and interfaces, and will be done jointly by the ONCS and subsystems groups.

Ancillary Systems Interface System

See the definitions and overview of ancillary systems in the original CDR section 11.6.

The ONCS group will provide a framework for implementing ancillary systems and integrating them into the PHENIX on-line system. The framework will include an operator interface consistent with the PHENIX on-line operator interface, recommended hardware and software standards, generic control and monitoring and interface software, and templates which can be used as starting points for building subsystem-specific ancillary control and monitoring applications.

We consider three major types of ancillary systems: facilities for control and monitoring of detector, trigger, etc., equipment not in the primary data path (e.g., voltage, current, temperature, etc., instrumentation); interfaces to RHIC accelerator systems; and interfaces to safety systems. The subsystem groups will supply their specific hardware and software. The ONCS group will provide PHENIX-standard interfaces to RHIC accelerator and safety systems in collaboration with RHIC operations and safety engineers. Critical safety applications will, of course, have to be specially treated, and often implemented in external hardware not directly controllable by PHENIX. The ONCS group will, however, be responsible for interlocking PHENIX internal processes based on monitoring of these external systems.

EPICS. The Experimental Physics and Industrial Control System (EPICS), originally developed in the LANL Accelerator Technology Division, is widely used for accelerator and process control and is now being evaluated for control of experiments (CEBAF, STAR at RHIC, SDC and GEM at SSC). EPICS is based on a control system model wherein a local area network provides communication between frontend microcomputers (for intelligent local control over attached devices) and workstations (providing operator interface and general-purpose computational facilities). Widely supported standard protocols, processor families, and operating systems are used throughout its implementation. EPICS further provides tools for: describing the control and monitoring functions; triggering procedures to be executed in response to requests or changes of conditions; alarm definition, management, and presentation; and archiving of system events and monitoring results. Databases associated with these activities are explicitly distributed over the frontend microcomputers, giving maximal subsystem independence and enhanced system fault tolerance. Because EPICS appears to provide all the facilities necessary for "slow control" and monitoring of PHENIX ancillary systems, and for interfacing to RHIC accelerator and safety systems, it figures heavily in our plans for ONCS.

EPICS is currently oriented towards control and monitoring of hardware and software "channels" residing in VME and VXI crates and owned by server tasks running in microprocessors under the VxWorks real-time operating system. Its most widespread application has thusfar been towards "industrial process" control and monitoring of hardware associated with power, temperature, pressure, motor control, etc., systems. However, in a physics experiment much information which needs to be monitored is calculated as part of event reconstruction and analysis. We will thus be investigating an extension of EPICS to control and monitor "software devices" residing in UNIX workstations. If this is possible, EPICS

might provide the basis for global control and monitoring across the entire PHENIX experiment. This could significantly reduce ONCS startup effort, while also providing a very widely supported uniform operator/user interface for all control and monitoring functions.

Monitoring System

Functions. The on-line computing systems must provide capabilities and resources for continual monitoring and validation of detector operation and of the quality of the data being archived. This is a computationally intensive task, due to the large number of front-end channels and subsystem control and monitoring points. The ONCS Monitoring System will provide support for not only collecting information and statistics on aspects of detector subsystem operations at many levels but also effective ways for presenting the resulting information (alarms, messages, graphs, archived logs) to both operators and software applications. The Monitoring System will notably also provide a common framework for reporting and responding to exception conditions throughout the on-line system.

As in the case of run control, the ONCS team will provide an integrated framework for monitoring functions, including standard user interfaces and templates for subsystem and operator interface integration. Specific subsystem monitoring application software will be the responsibility of the relevant subsystem groups.

On-line event-data analysis. To help ensure maximal quality in the archived data, the on-line computing systems must also provide a framework and resources for on-line event analyses, including full event reconstruction for some fraction of accepted events. This requirement places demands on both bandwidth and computing resources. Its detailed implications require further specification of the mix of analyses, but for purposes of estimating minimum resources we have considered full reconstruction of minimum-bias Au+Au events.

The Offline and Simulations groups estimate that reconstructions of Au+Au minimum bias events require about 250 MFs/event (MFs = MegaFlops-second), and the DAQ group estimates the corresponding event size to be about 75 KB. On a 50-MF processor, it will thus take $250 \text{ MFs} / 50 \text{ MF} = 5$ seconds to fully reconstruct a Au+Au min-bias event. If we write data at 20 MB/sec with a livetime of 50% for 1/2 year of running, we will record 2.1×10^9 min-bias events. Reconstructing these events on one 50 MF processor would require $(2.1 \times 10^9 \text{ events} * 5 \text{ sec/event}) / (3.1 \times 10^7 \text{ sec/yr}) = 340$ years, or 1 year on 340 such processors.

In our current plan the on-line computing systems will include 16 high-performance workstations. With four of them dedicated to central on-line operations, 12 are left for other tasks. If we assume 50% of the CPU power of these 12 can be used for on-line event analysis, we will have $12 * 0.5 * 50 \text{ MF} = 300 \text{ MF}$, or the power to reconstruct $300 \text{ MF} / 250 \text{ MFs/event} = 1.2$ min-bias events/sec. This implies a data rate to the analysis nodes of about $1.2 \text{ events/sec} * 75 \text{ KB/event} = 90 \text{ KB/sec}$. This can be easily accommodated by an ethernet-based monitoring network. The fraction of min-bias events fully analyzed on-line according to these estimates is about 0.45% of all events logged.

However, most monitoring purposes will not require full event reconstruction. If the dominant event processing is 10 times faster than full reconstruction, then the required

bandwidth is about 1 MB/sec. Furthermore, it is not unreasonable to expect an eight-fold increase in workstation performance by beam time (to about 400 MF per processor). Then the bandwidth requirements for monitoring can reach 8 MB/sec (and we could fully reconstruct in real-time about 10 events/sec, or 4% of all recorded min-bias events).

Ethernet is clearly not appropriate here — a higher bandwidth implementation, like FDDI or FibreChannel, must be considered. We have thus included an FDDI-type (10 MB/sec bandwidth) network for estimating the cost of the ONCS "Monitor Network". The EPICS system notably already accommodates this possibility, as well as other high speed sub-netting strategies. However, we reiterate that for EPICS to enable on-line event analysis it must be extended to incorporate UNIX-resident "software devices."

Diagnostics and Maintenance System

Understanding and repairing system failures requires an integrated set of diagnostic tools. Most diagnostics will be very subsystem specific, relying on knowledge of detailed hardware features, so low-level diagnostic aspects must be addressed by the associated subsystem groups. But the ONCS team will provide a standard framework plus supporting tools, protocols, interfaces, and templates for integrating subsystem diagnostics in a unified PHENIX environment. These capabilities must, in fact, be incorporated early in the ONCS development, to support subsystem prototyping and testing.

The detailed requirements for diagnostics remain to be determined. Recent experience at HERA suggests that expert systems may aid operators in system-level diagnostic procedures, especially during production running. Further developments along these lines and related applications in other large-system contexts will be evaluated. At the very least the Diagnostic System framework should accommodate integration of such tools.

Event-data Flow Control System

A guiding principle underlying the data-driven architecture presented in the CDR (CDR Fig. 11.2) is that communication between processing elements is through memory. Operations proceed when all required inputs are available, and statistical fluctuations in the time between events and/or in the data volume per event are accommodated by multi-event buffers controlled via local "holds" derived from *current* considerations of *locally* available free space. This scheme generally optimizes usage of available bandwidth, automatically provides local load-leveling, and obviates any need for complex "busy chains."

The need for robustness and fault tolerance nevertheless requires high-level facilities for imposing "slow" flow control. This is especially true for subsystems with numerous inputs from disparate sources, where it is desirable to guard against loss of synchronization and to provide for dynamic fault recovery. A few levels of selective inhibit are also important for certain calibration and diagnostic operations, as it may be necessary to allow selected parts of the system to cycle without the overhead of operating the entire on-line system. And, in a completely mundane sense, normal runtime operations like initialization, startup, and shutdown require some central means for selectively inhibiting all data flow.

The event-data flow control system is intended as a flexible, programmable environment for providing these capabilities. Its detailed requirements will be developed in collaboration with the DAQ group. This system must at least interface to both the Control and the Monitor and Calibration Networks, to provide paths required for initialization, validation, and exception reporting/handling.

Event-data Archive System

See CDR Section 11.4.3.

Calibration System

In collaboration with the offline computing systems (OFCS) group, the ONCS team will provide a unified framework for acquiring, archiving, and organizing calibration data, and for monitoring the validity of the stored constants during operation. This will naturally include interfaces for down-loading calibration data to the appropriate destinations during on-line system initialization.

Because calibration is heavily dependent on subsystem details, applications addressing subsystem specifics for acquisition and handling of calibration data are necessarily the responsibility of the corresponding subsystem groups.

The potentially large quantity of calibration data that will have to be generated and accessed on-line may have important ramifications for the on-line computing resources and bandwidth, if rapid on-line processing is crucial to their validation and maintenance. The requirements for calibration remain to be explored in detail.

11.6.4 Generic Services

User Interface Services

The on-line computing systems group will provide a set of PHENIX-standard tools and software for creating and integrating interfaces for all PHENIX applications. Our goal is a common look and feel to all PHENIX user interfaces, with a standard way of referencing components and operations across all subsystems. This uniformity is especially important for the on-line operator interface, in order to simplify operator interactions with an inherently complex system. These user interface components and tools will support:

- a single standard framework for the on-line operator interface;
- a set of standard interface components ("widgets") for selection, data entry, data presentation, detector and event displays, alarm presentation, multimedia message presentation, etc.;
- standard access control mechanisms (for protection of critical functions and data);
- standard graphics software facilities;
- a standard graphical user interface (GUI) builder.

Database Services

The on-line system will include numerous databases. These include descriptive databases that document the detector and the on-line environment, and archival databases containing historical records of significant occurrences and changing conditions, with indexes and catalogs for aiding information retrieval. The event-data archive is a special case and is discussed in a separate section (See section 11.6.3.).

Identified database types include:

- Configuration database and archive - description and record of the detector and on-line system run-time environment and configuration, including partitioning, running modes, trigger settings, available resources, initialization and calibration data sets, software and hardware versions, selected options, etc.
- Detector Assembly & Survey database and archive - description and record of all detector components and their arrangement, including component locations, names, addresses, geometry, etc.
- Data Structures database and archive - a standard "dictionary" defining and recording all data objects to be handled by ONCS processes.
- Calibration database and archive - organized collection and record of all calibration constants required on-line, including data to be down-loaded into on-line hardware and software modules.
- Monitoring database and archive - current collection and record of detector and on-line status information, performance data, environmental and external-system conditions, exceptions, etc., relating to the operational state of the detector.
- Catalog of event data archive - an index, with appropriate cross-referencing and descriptive information, for identifying and locating archived event-data.
- Experiment Logbook - a database that assembles and records all information needed to define and later reconstruct the conditions, environment, objectives, problems, etc., occurring during PHENIX running (and testing). Our goal is an entirely "on-line" implementation, with hypertext and multimedia capabilities.
- Documentation Database and archive - a means for collecting, organizing, and retrieving all ONCS documentation, also "on-line" with hypertext and (perhaps) multimedia capabilities.

Because much of the information contained in these databases is also important to off-line and simulation computing, the ONCS team will collaborate with those groups in the definition and in the establishment of standard protocols, interfaces, and tools for creating, accessing and managing general PHENIX databases.

Interprocess Message Services

Because of its complexity and distributed nature, ONCS will require extensive multiprocessing for maximal performance, efficiency, and maintainability. Unified interprocess communication services will therefore be critical to a quality implementation for ONCS. The ONCS group will define PHENIX-standard protocols and interfaces for passing messages between on-line software components, and between software processes and users and programmers. These standards and services will provide a key ingredient for integrating both the overall ONCS framework and subsystem-specific applications.

11.6.5 Development plan

Organization

The ONCS Group will be a core team of system-level experts, providing central organization and support for global on-line computing efforts. Subsystem-specific applications must be provided by subsystem personnel, to exploit knowledge of specific hardware.

During the analysis/specification/design phase, the ONCS and OFCS teams will collaborate with the PHENIX Simulations team and representatives from the RHIC Computing Center to coordinate activities in areas involving major technology overlap. This will require a PHENIX Computing Systems Central Design Group of approximately 15 experts drawn from PHENIX DAQ (4 persons), ONCS (6 persons), and OFCS (4 persons), plus a PHENIX liaison to the RHIC Computing Center (1 person).

This group will be aided by a panel of PHENIX Detection Subsystems Computing Technical Leaders, formed by drawing at least one person from each of the PHENIX detector subsystem projects (10 total). They will have responsibility for guaranteeing success in all computing related aspects of their subsystems involvement with PHENIX. Initially this panel will serve (1) to provide computing requirements input to the Computing Systems Central Design Team, and (2) to evaluate the latter's proposals, specifications, designs, prototypes, etc., to help assure that their efforts meet all relevant PHENIX requirements. This panel's responsibilities will evolve to providing resources to the PHENIX computing systems projects from the subsystems teams for development of subsystem-specific applications. The ONCS project will, in particular, explicitly tax subsystem teams for service in overall ONCS integration and validation.

We assume that the aforementioned 25 positions are staffed entirely through contributions from PHENIX-collaborating institutions. These groups' depths are currently inadequate for efficient execution of all aspects required for development of PHENIX Computing Systems. Obviating overhead and delays associated with internally developing the necessary degrees of expertise in several key areas will require access to a resource pool of 30 FTE-years, for contracting work by established experts and toolsmiths in areas such as: operating systems and utilities, graphics, user interfaces, databases, CASE, and real-time embedded systems. Also 2 technician-FTE's will be needed during installation and validation of ONCS hardware at the PHENIX site, and 2 software and administration support staff are required for routine maintenance of ONCS facilities throughout system development.

Development Process

We are in the initial stages of creating a detailed development plan. Because of the relatively large and diverse on-line computing systems effort, and the importance of quality and reliable on-line operation, we intend to create a somewhat formal development process that defines explicit procedures and tracks progress through regular reviews and careful documentation. At all times there will be close coordination with the various subsystem developers.

The initial task of the ONCS group will be the definition and implementation of a development environment. This will include defining standards and acquiring tools for supporting: analysis and design methods, quality assurance, code configuration and management, coding style, operating systems, languages, graphics, user interfaces, databases, project management, standard processor families, interconnects, and networks. To make efficient use of these tools a training program will be necessary (starting, for example, in system development methods).

The standards used by the ONCS group will in most cases become de facto and recommended standards for all developers of on-line applications. To facilitate ONCS integration a set of PHENIX-specific software standards, tools, and quality assurance mechanisms will be defined and made available to application developers.

The on-line computing systems will be developed according to an iterative process of requirements analysis, design, and implementation. It will be iterative in the sense that the analysis, design, and implementation activities will not be carried through for the final system in large separate phases ("waterfall" method). Rather prototype systems, with some functions left out or not fully implemented, will be completed at the end of each of several analysis, design and implementation cycles ("spiral" method). Subsequent cycles of analysis, design, and implementation will augment and revise the system based on feedback from users' experiences using the earlier prototypes and increased understanding of requirements. An important component of each release will be complete documentation.

This evolutionary approach will allow us to independently develop subsystems of the ONCS, while always remaining closely coupled to evolving needs and schedules of the other PHENIX subsystem groups. It is aimed at providing required functionalities as early as practicable, and at the same time ensuring the quality of the final system by emphasizing continual product evaluation and feedback from the users.

We currently envision three major versions of ONCS:

- An "Early-prototype System" implementing essential functions that are needed by detector subsystem groups for development and testing. This prototype is intended to provide a basis for refining our understanding of ONCS user requirements via actual user feed back, and for testing early design ideas.
- A "Prototype System" based on the design resulting from the previous phase. This prototype is intended to test the overall system design and implementation and elicit further refinement of user requirements.
- A "Production System" implementing and integrating all functions needed for actual PHENIX operation at official RHIC startup, based on evaluation of previous versions.

Each major version will be broken into minor releases that will provide frequent stages of increasing functionality:

1. A minor release implementing functions that provide support for run control, ancillary systems, monitoring, diagnostics.
2. A minor release adding increased functionality to support common services (operator interface, databases, interprocess messages).
3. A third minor release implementing additional functions that provide support for remaining framework functions (event-data flow control, event-data archiving, calibration).

Two key releases will provide:

- A first prototype (Early-prototype System first stage) at the end of 1995, with functions supporting detector subsystem development and testing.
- A fully integrated prototype based on the overall design (Prototype System third stage) at the end of 1997, with all functions required for full detector cosmic ray tests.

11.7 Cost and Schedule

The cost estimate and schedule for on-line system is presented in Tables 11.10 and 11.11. The cost of the front-end electronics common systems (WBS 5.3.13), the data collection modules (WBS 5.4.1), the subevent transfer buffers (WBS 5.4.2), the event builder (WBS 5.4.3), the level-1 trigger (WBS 5.4.4), and the level-2 trigger (WBS 5.4.5) are shown in Table 11.10 at PHENIX WBS level 4. The cost of the master timing and control system (WBS 5.4.8), on-line computing system (WBS 5.4.9), the simulation computing system (WBS 5.4.10), the on-line computers, archiving and networking systems (WBS 5.4.12), and the common DAQ system work (WBS 5.4.13) are shown in Table 11.11 at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

[illegible]

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998				1999			
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4				
5.4.1	Data Collection Modules	\$1,676,387.33																										
5.4.1.1	Design	\$191,627.81																										
5.4.1.2	Prototype	\$204,911.52																										
5.4.1.3	Fabrication	\$1,335,848.00																										
5.4.1.4	Assembly	\$67,900.00																										
5.4.1.5	Test	\$86,400.00																										

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998				1999			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	
5.4.2	Subevent Transfer Buffers	\$349,928.96																											
5.4.2.1	Design	\$145,185.28																											
5.4.2.2	Prototype	\$120,712.48																											
5.4.2.3	Fabrication	\$26,431.20																											
5.4.2.4	Assembly	\$14,400.00																											
5.4.2.5	Test	\$43,200.00																											

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998				1999				
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	
5.4.3	Event Builder	\$2,531,164.48																												
5.4.3.1	Design	\$202,228.80																												
5.4.3.2	Prototype	\$350,863.68																												
5.4.3.3	Fabrication	\$1,891,862.00																												
5.4.3.4	Assembly	\$9,600.00																												
5.4.3.5	Test	\$76,800.00																												

[illegible][illegible]

Table 11.11: Cost Estimate and Schedule for PHENIX On-line System, Continued

[illegible]

WBS	Description	Total Cost	1994				1995				1996				1997				1998				1999		
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	
5.4.9	On-line Computing System	\$3,917,955.52																							
5.4.9.1	Design On-line Computing System	\$613,791.66																							
5.4.9.2	Prototype On-line Computing System	\$0.00																							
5.4.9.3	Fabricate On-line Computing System	\$3,194,819.84																							
5.4.9.4	Assemble On-line Computing System	\$109,344.00																							
5.4.9.5	Test On-line Computing System	\$0.00																							
5.4.9.6	On-line Computing System Completed	\$0.00																							

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998				1999			
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	
5.4.10	Simulation	\$192,320.00																											
5.4.10.1	Simulation Design	\$14,400.00																											
5.4.10.2	Simulation Prototyping	\$69,200.00																											
5.4.10.3	Simulation Fabrication	\$6,400.00																											
5.4.10.4	Simulation Assembly	\$11,520.00																											
5.4.10.5	Simulation Test	\$100,800.00																											

[illegible][illegible]

Chapter 12

Off-line System

12.1 Overview

The purpose of the PHENIX off-line computing system will be to provide all aspects of data and information handling that are not directly connected to the collection of data and the monitoring of the experiment. Examples of functionalities provided by the off-line system are: event and detector simulation, event reconstruction and analysis, and data and general information management.

The off-line system can be viewed as another subdetector, which will have to undergo the usual engineering phases: design, prototyping, fabrication, implementation and maintenance. The off-line system does, however, function in a somewhat different way, because it also can be viewed as a service organization for the rest of the collaboration. It will have to provide all the necessary tools for the physicists to transform the data, originally written to a storage medium by the data acquisition system in some digital format, into interesting physics papers and exciting conference talks. Whereas only a small subset of physicists will be directly interacting with each of the subdetectors, practically all members of the collaboration will interact with the off-line system on a daily basis (e.g., analyzing data, reading online manuals, producing figures, or accessing information on the state of the detector during a run two years).

The PHENIX off-line system will provide the same functions as most other off-line systems do in current large scale experiments. However, the challenge for PHENIX is the magnitude of the data flow coming from the detector. The requirements for data storage are two to three orders of magnitude larger than current experiments, and the computational needs exceed even the largest supercomputers currently available. In many respects PHENIX's off-line requirements are comparable to or even larger than those of the SSC detectors.

Table 12.1: Data Storage and CPU requirements for PHENIX and selected RHIC and SSC detectors.

	PHENIX	STAR	PHOBOS	SDC
Event recording rate (Hz)	100	1	400	100
Event size (kBytes)	300	20,000	12	1000
Recording rate (MBytes/s)	20	20	5	100
Events / year (10^6)	650	10	4000	1000
Raw data / year (10^{12} Bytes)	200	200	50	1000
Gigaflops-sec / event	1.0	10	0.005	0.3
Continuous off-line CPU need (Gflops)	65	10	2	30

12.2 Requirements

Basic Parameters for Estimation of Requirements

The initial phase in the design of the PHENIX off-line system was an evaluation of the physics and functional requirements. The two most important and expensive requirements by far are: (1) to have sufficient CPU power for the initial event reconstruction of the raw data, and (2) to be able to store all the recorded data (raw and simulated). Table 12.1 summarizes the basic parameters for the estimation of these requirements for PHENIX, STAR and PHOBOS [1] and the proposed SSC detector SDC [2]. Both PHENIX and STAR expect a recording rate of 20 MBytes/s. Due to STAR's very large event size, they will only record one event per second in contrast to PHENIX's rate of 100 events per second. In these estimates it is assumed that all events are central Au + Au events with charged particle multiplicities around 6000 and that RHIC will operate approximately 10^7 seconds per year.

Data Storage Requirements

The assumptions given above lead to an estimated 200 TeraBytes of raw data each year. It can be assumed, that additionally 300 TeraBytes of simulated data and reconstructed data will need to be stored resulting in a *total storage need for PHENIX around 0.5 PetaBytes each year*. This requirement is 2-3 orders of magnitude larger than current experiments and corresponds to $\approx 100,000$ 8mm videotapes. An additional requirement for the data storage system will be to store as much of the data as possible with access times less than a few minutes. The proposed hierarchical data storage for all RHIC experiments will be described in Section 12.4.

CPU Power Requirements

In order to estimate the total need for CPU power it is necessary to estimate the CPU consumption per event. With the current tracking algorithms in PISORP 2.2 a full recon-

struction of a central Au+Au event takes approximately 1-2 minutes on a IBM RS6000 workstation, corresponding to 1000 Gigafllops-seconds. With an event rate of 65 Hz this implies a *continuous CPU need around 65 Gigafllops for event reconstruction*. If the event reconstruction facility has an "uptime" of 90% the CPU need reduces to 40 Gigafllops, but this reduction is compensated by the additional need for production of simulated data and repeated complicated data analysis. The magnitude of this CPU power requirement can be put in perspective by the fact, that the largest commercially available systems currently (1993) is around 10-20 Gigafllops. Due to financial limitations it might not be possible for PHENIX to obtain access to the required 65 Gigafllops of CPU power at RHIC and we might have to seek additional solutions (see Section 12.3.1).

Data Analysis Requirements

For the many users of the off-line system it will be mandatory that they can get *easy access* both to the data and to processing power. We estimate, that once the system is fully operational it will have to be able to support up to 50 concurrent users with up to half of them present at RHIC and the other half situated at their home institution. It is envisioned, that the bulk of the data will be stored and processed at RHIC, but the remote users will need access to a powerful *Wide Area Network* (WAN) with transmission speeds of 1 Gbit/s in order to function as well as the on-site user.

In order to maximize the productivity of the user in the simulation and/or analysis phase, he/she should be able to communicate with the off-line system through a *simple and uniform user interface* with full online help facilities and the same look-and-feel for all parts of the off-line system. To the extend it is possible the same GUI should also be used for the online system.

The user should have full *flexibility* in the use of the off-line analysis system. The system should operate across a reasonable set of *platforms*, both on- and off-site. Both *serial* and *parallel* computers should be supported. On-line *documentation* of all aspects of the system should be available. The system should make as many tasks as possible *automatic* and *transparent* to the user, while at the same time providing the advanced user the option of customizing his/hers programs.

Finally it should be emphasized that financial constrains might make it impossible to fully implement all the requirements. In order to ensure, that a fully functional system and optimized system will be ready from "day one", the off-line system will be designed and implemented through a set of more and more elaborate prototypes, which will allow a *design-to-cost* philosophy to be implemented.

12.3 Computing Models

12.3.1 The RHIC Computing Facility

In order to meet the requirements outlined in the previous section a simple model for PHENIX's off-line computing system has been developed. A fundamental assumption within

the model is, that the bulk of the computing will be done at a central RHIC computing facility (RHICCF), which will provide both hardware and manpower for a large fraction of PHENIX's needs. The current "strawman" model for RHICCF is shown in Fig. 12.1 obtained from a report by the RHIC Off-line Computing Study Group chaired by B. Gibbard [3].

The model operates with 4 main clusters of computing facilities: production clusters for PHENIX, STAR, a common cluster for the small experiments, and a common cluster for I/O intensive tasks. Each production cluster will consist of a) a set of CPU servers scaled to the CPU needs of the particular experiment, b) a file server for short time storage of data and experiment specific databases, and c) a tape drive system with an intermediate size tape robot. The I/O intensive cluster, which will be shared between all RHIC users, will contain a large scale disk and tape storage system. It is envisioned, that the data reconstruction is taking place at the production cluster and the data analysis is done at the I/O intensive cluster.

In addition to providing the bulk of the hardware needs for PHENIX's off-line system RHICCF will also provide a substantial amount of manpower and software for core tasks like a) general system support, b) network support, c) code development support, d) general application (CERNLIB, GUI, databases, publishing etc.) e) technical development and f) hardware support.

The current model assumes PHENIX's production cluster allocation to contain:

CPU Servers:	≈ 20 GFlops of CPU power based on loosely-coupled workstation-like servers at \$100/MFlops.
File Servers:	A total of 800 GBytes of disk space consisting of eight 100 GByte file servers, each providing 16 MBytes/s
Mass Storage Playback:	19 mm tapedrive providing 16 MBytes/s.
Mass Storage Robot:	2 TByte robot supporting 1 tapedrive.
Network:	ATM operating at 16 MBytes/s connecting all ≈ 60 nodes (CPU servers, file servers and tape drives).

In addition PHENIX will have shared access to the I/O intensive cluster, which currently is envisioned to be configured as

CPU Servers:	≈ 5 GFlops of CPU power.
File Servers:	A total of 2400 GBytes of disk space.
Mass Storage Playback:	3 19 mm tapedrives, each providing 16 MBytes/s.
Mass Storage Robot:	60 TByte robot supporting 3 tapedrive.
Network:	ATM connecting ≈ 40 nodes.

Most of the technology choices will be postponed until two years before the estimated RHIC turn on in order to optimize the system and to take advantage of the rapid developments in the computer industry. For the CPU server it will be of special interest to follow whether a farm of loosely coupled workstations, as currently used at FNAL, or a massively parallel processor (MPP), like the PARAGON at ORNL, will be the most optimal choice.

A serious problem is, however, the severe discrepancy between the CPU power of 20 GFlops currently allocated to PHENIX at RHICCF and out estimated need of more than 65 GFlops. Since computer performance improvements larger than 1.4/year, as assumed in Ref. [3], seems unlikely, we are forced to consider other alternatives like: a) Obtaining additional

Figure 12.1: "Strawman" model of the proposed RHIC Computing Facility. The center will contain 4 clusters. PHENIX will have its own cluster for event reconstruction and have shared access to the I/O Intensive cluster for data analysis.

compute cycles at major collaborating institutions like ORNL, LANL and/or LLNL, which currently operate some of the worlds largest supercomputers, b) perform a larger fraction of the total operations on an event at the highest trigger level, c) develop more efficient event reconstruction algorithms or d) reduce the total amount of operations on each event. Since the last solution could imply, that only a few especially interesting events would be fully reconstructed, it is our hope, that a combination of a), b) and c) will solve PHENIX's CPU shortage.

12.3.2 Main Code Modules of the PHENIX Off-line System

Of particular importance, especially in the construction phase of PHENIX, is the interaction between the off-line and the simulation system. Fig. 12.2 shows the main modules of the two systems and their interconnection. The two systems will interact via a set of standardized data structures stored in the event data base. If sufficiently powerful computers will be available at a later stage, the communication between the two systems can take place without intermediate storage. This later mode of operation can easily be implemented on any parallel computer system.

The off-line event handling package will be designed so the user will have full control over the in- and out-put level of the data structures. This implies, that PHENIX will *not* operate with separate reconstruction and data analysis packages with different user interfaces etc., but the event reconstruction will take place by running the event handling programs with raw data as input and reconstructed events as output. A data analysis user will then run the same program with reconstructed events as input, but will also have the option of starting at a lower data structure level and redo the tracking with other tracking modules than used in the official event reconstruction.

Software Layer Model

In order to more clearly differentiate between the function of different modules of software and the responsibilities of the creators of these modules a simple layered model for the PHENIX software has been generated. As show in Fig. 12.3 we envision three fundamental layers of software within PHENIX, each with their own distinct functions and creators.

Physics Applications: All modules operating on real and/or simulated event data with the purpose of transforming the data from one data structure level to another level. Examples: Tracking modules, detector simulation modules. These modules will be written by PHENIX physicists (graduate students, post-docs, ..) with as few restrictions as possible concerning programing language in order to enhance the productivity.

Core Applications: This layer will consist of core modules and interface modules for the event handling package. Examples: I/O modules for the data structures, general software bus for the offline system, event display, etc. These modules will be designed and programmed by the off-line group in collaboration with the simulation and on-line group. The programming language is likely to be C or C++.

The Main Event Handling Modules in The Off-line and Simulation Systems

Figure 12.2: Conceptual layout of the main modules of the event handling package and the simulation package. The data flow through the simulation system will be going from the event generator module at the lower right up to the digitization module in the upper right corner. The output from this module and the on-line system will be accepted by the off-line system for calibration, event reconstruction and data analysis. The main modules will exchange information through a set of data structures (labeled DS). This design will provide a high degree of modularity.

Software Layer Model

Figure 12.3: Conceptual layered model for the design of PHENIX software. The purpose of this model is to focus attention on a) who has responsibility for the various layers of software, and b) what requirements should be imposed on the software.

System Applications: All system modules which might be either common to all RHIC experiments or which will require advanced programming know-how. Examples: Network services, System management, Data base management etc. We envision this layer being the responsibility of the RHIC Computing Center.

As indicated on Fig. 12.3 the three layers will communicate through two sets of Application Program Interfaces (API). The first will be designed by PHENIX. The second will have to be designed by RHICCF and all RHIC experiments in collaboration.

12.4 Data Storage and Management

Probably the most important task of the off-line system is to organize and provide access to the multitudes of data collected by the collaboration. The data can be subdivided into three different classes: a) Simulated or real event data, b) detector and instrumental data and c) general information data. This division of all the PHENIX data into 3 distinct classes is shown in Fig. 12.4, which illustrates a conceptual layout of the general PHENIX multidatabase system.

12.4.1 Event Storage and Management

As described above the event data volume to be handled by PHENIX is in the order of 500 TeraBytes/year. The handling of this gigantic amount of data is likely to be the most difficult issue facing the PHENIX off-line system. The experience from other experiments indicate, that already the handling of much smaller data volumes can be a serious bottleneck.

All RHIC experiments will be faced with this data volume problem and it is therefore important, that the storage problem will be solved as a collaborative RHIC effort. The RHICCF committee [3] has defined the following criteria for the RHIC primary mass storage:

- A single recording format should be used for all detectors.
- The recording system must be supported by a large scale robotic system.
- The recording rate for a single drive should be a significant fraction (>50%) of the anticipated raw data collection rate.
- The capacity of the individual media must be as large as practical.
- The cost must be as low as possible.

Currently the best option seems to be 19 mm tapes from either Sony or Ampex. These tape drives support recording rates up to 32 MBytes/s and each tape can record up to 100 GBytes. Due to the large capacity of each 19 mm tape the associated tape robots have a very large capacity at a low price of $\approx \$13\text{k/TByte}$. But even this low price will only allow a small subset of the annual data volume to be stored in the robots. The amount and types of

Conceptual Layout of PHENIX's Multi-database

Figure 12.4: Conceptual layout of PHENIX's multidatabase. The different modules are described in more detail in the main text.

data stored in the mass storage robots for access within few minutes will therefore probably be limited to very recent data and to a variety of data summaries.

The mass storage will likely contain a large disk system serving as a rapid access cache. This will allow writing to the tape drives at a uniform speed, but more important it might allow the off-line system to perform a variety of operations on the data before they are stored on tape. These operations can range from simple calibrations to the actual event reconstruction.

An effective file caching between the disk and the mass storage robot is of major importance for the throughput of the off-line system, especially for users performing data analysis on a large set of different files. Several commercial solutions for the file caching are available, all derived from the UniTree product developed by the National Storage Laboratory at LLNL and IBM.

The location of the experimental and simulated data will be stored in the following database:

Event Database: Typical information stored in this database will be: How far has each run proceeded through the complete event reconstruction and data analysis process and where are the different parts of the complete data structure for a given run located in the data storage system. In addition the database will store information on the location of mixed event runs, 2- and 3-dimensional spectra, ntuples etc.

12.4.2 Detector Database

The detector multidatabase will contain all information pertaining to the detector and the environment under which the experiments are performed. The multi-database will consist of several databases with a common user interface:

Configuration database: All parameters describing the complete configuration of the detector: geometrical data, hardware configuration of all data acquisition channels, all production software for triggers and off-line system. All entries will be time stamped, so the complete detector configuration at any time can be reconstructed. Estimated size: 1-5 GBytes with an update rate of 10 MBytes/day

Calibration Database: All calibration parameters obtained either in earlier calibration runs or obtained as feedback from off-line analysis. The instantaneous size (total size of all data active at any particular time) will be ≈ 1 GByte, but since many of the sensitive calibration constants might change on a daily bases we estimate the total database will grow with 100-200 Gbytes/year.

Electronic Log Book: Will be obtained from either manual or automatic entries from the online system. This database will be a multimedia database containing all records pertaining to the recording of data: personal comments, graphics, histograms, VCR recordings, audiotapes of communication etc. Estimated size: 100 Gbytes/year (this estimate is very dependent on the extend to which graphical data will be stored directly in the database).

Monitoring Database: All information related to the monitoring of the status of the detector. Part of this database might reside within the online hardware.

Software Database: All official software will be contained in a time- and version-stamped database. This database should either be fully integrated with the code management system or be the code management system itself.

12.4.3 Information Database

The information database will contain all the general information necessary for the collaboration.

People Database: A small database containing information about all PHENIX collaborators: addresses, telephone numbers, responsibilities etc.

Physics Results Database: A multimedia database containing transparencies, articles, PHENIX technical notes, histograms, movies etc. describing the results obtained by PHENIX. Included will also be similar data obtained by other collaborations.

Manuals and Reports: An easy accessible database (WWW?) containing on-line manuals, reports and documentation of all aspects of PHENIX.

The total size of the information multi-database is expected to be considerably smaller than the detector database, but might have a considerably more complex structure due to the larger amount of relations between the entries.

The planned capacity of 800 GBytes for the PHENIX file server seems to be adequate for containing the parts of the PHENIX databases, which need a rapid access time.

12.4.4 Data Structures

All the different software modules of the off-line system will be communicating with each other through a set of data structures. It is of great importance to design these structures as early and as flexible as possible. The data structures are usually organized in banks (ZEBRA terminology) or structures (C or F90 terminology) and the following set of minimum requirements for their implementation has been defined:

- All structures should be defined in selfdocumenting files.
- There should be *no* numeric pointers to the data.
- There should be full support for nested structures.
- There should be tools for defining more complex relations between banks (like one-to-many, many-to-one etc.).
- The structures should be highly integrated with the programming language used for the modules operating on them.

There exist already several product developed in the high energy physics community, that partially fulfills these requirements (ZEBRA, JACELLE, CHEETAH, ADAMO ...). Currently CHEETAH seems to be the best match to our requirements.

12.5 Software

12.5.1 Code Development Environment

The main part of the construction of the off-line system will consist of programming the hundreds of thousands of lines of computer code needed. It is therefore of outmost importance to create a highly structured and well functioning code development environment. In the following some of the most important ingredients in such an environment will be discussed.

Operating System

UNIX will without any doubt be the dominant operating system within PHENIX for the next 2-4 years. Currently most workstations and CPU servers envisioned in the PHENIX off-line system runs UNIX. Over this time period there will be a decreasing support of VMS. We will, however, follow the commercial developments and success of systems like NT, NextStep and openVMS in order to evaluate, if we at a later stage should adopt one of them.

It is the policy of PHENIX to work closely together with the RHIC Computer Facility, and it is our hope, that any changes in supported operating systems or other major changes in the computing environment will be done simultaneous by all RHIC experiments through collaboration with RHICCF.

Programming Language

As discussed above in connection with Fig. 12.3, then it is envisioned, that several programming languages will coexist in the off-line system. Many high level applications programmed by physicists will still be done in FORTRAN77, whereas many of the core modules more likely will be done in C. It is especially the support for structures and dynamic memory allocation, that makes C attractive for core module programming. We see currently no particular advantages in adopting FORTRAN90, but if F90 should increase dramatically in popularity and should high efficiency compilers be commonly available, then there would be no serious problems associated with accepting modules in the off-line system programmed in F90. The question of object oriented programming languages will be addressed below.

Code Management

The basic function of the code management system is to provide a highly structured environment for the development of software. We have identified the following requirements to our code management system:

History FEature:	Autodocumentation of changes and easy regeneration of earlier versions based on either version number or date.
Auto Documentation:	Features for automatically documenting each individual module
Include Feature:	Easy insertion of common data structures and/or code.
File Structure:	Should support a UNIX like directory structure with all modules accessible outside of the code manager.
Module Generation:	Should provide easy scripts for generating a specific subset of the total code.
Configuration Control:	Configuration control should be supported at the source text level (of special importance when operating with code for both parallel and serial computers).
Network Distribution:	Support for automatic network distribution of the complete code.
Code Compression:	Support for code compression to reduce network transfer times.

At the start of the development of PHENIX's Integrated Simulation Application (PISA) it was decided to use the CMZ code management. CMZ has served well over the last 18 months. It does, however, not fully satisfy all the above requirements and a new code management system will be selected. CVS or its descendant RCVS seems to satisfy most requirements and does furthermore have the advantage of being used by STAR.

CASE Tools

Until now most major experiments have managed to develop and implement their off-line systems without full use of a CASE tool (Computer Assisted Software Engineering). Due to the increased complexity and worldwide distribution of the next generation of experiments at SSC and RHIC this is, however, likely to change. ALEPH and D0 used SA/SD (Structured Analysis / Structured Design). This provided a smart way of describing requirements, Data-Flow diagrams, Entity Relationship Diagrams and State Transition Diagrams. It was, however, difficult to follow all the fancy diagrams up, since the SA/SD tools did not easily translate into code. Several newer CASE tools does, however, promise greater power and flexibility.

It will be one of the first tasks in the design phase of the off-line system to investigate to what extent CASE tools should be used to develop our software. The question is closely coupled with the availability of moderately priced CASE tools running across many platforms and the likelihood of getting all people involved in software development to accept the stringent use of CASE.

Code Documentation

It is of outmost importance in a heterogenous collaboration like PHENIX, that all software is very well documented and written in such a way, that it can be easily understood by all interested parties. Many code management system provide some features for automatic code

documentation at check-in time and it is our intension to expand these capabilities, so no code module is not automatically fully documented. The documentation will consist of both header information for each module and extensive in-line comments as well as full online description in the information database described above.

Code Testing

Testing of software is probably the area in which the physics community is most far behind the commercial software developers. Whereas we understand, that it probably will be very difficult to reach the commercial quality assurance (QA) level, then it is our intension to develop a much more structured and detailed testing scheme than usually implemented in nuclear physics experiments.

Some of the elements entering in our QA program will be:

Design: Probably the most important aspect of QA is to create a well structured design before any coding.

Modularity All code should be highly modular. This allows easier testing and replacement of old modules without unwanted consequences.

Interfacing All modules should communicate through a set of well-defined data structures and should only be allowed access to the parts of the data structure absolutely necessary for its operation.

Visual inspection: All code will be periodically inspected by software guru's in order to check if it complies with PHENIX's set of programming rules, and if the modules only operate on the parts of the data structures, that they are intended to.

User test case: Programmers will be required to provide test cases for all their module packages. Software guru's will periodically check the code against these test cases.

Simulated data In collaboration with the simulation group a special set of simulated test data will be created by a small group of people, who are not directly involved in the coding. All major modules and packages will be tested against this set of simulated data before every major release. The simulated test data will continuously be refreshed and augmented with new "weird" special cases in order to test all corners of the off-line system software.

Object Oriented Approach

A major decision to be made for the off-line system will be the extend to which object oriented approaches are going to be used. We will identify one or two test projects for this approach. After the finish of the main prototyping (see Section ??) a decision will be made, if OOP is going to be used in the final implementation. This approach has the advantage, that if OOP is chosen, then our understanding of the system is much better after the prototyping and much better objects can be designed. OOP will in general only lead to superior code, if the objects are well-designed.

12.5.2 High Level Program Packages

The high level program packages are defined as all the software that performs "physics" operations on the data. The main ingredients of the off-line software will be (see also Fig. 12.2):

Event Generators: Currently several different event generators based on hadronic scenarios (FRITIOF, VENUS, HIJET) are being used in addition to a set of simple vectormeson generator based on Vogt's work [4]. We intend to strongly expand our interaction with theorists in order to obtain a larger set of Monte-Carlo event generators addressing physics of interest to PHENIX.

Detector Simulation The main simulation tool is PISA (Chapter 3). PISA performs a detailed detector response simulation based on GEANT. PISA is considered a prototype which is going to be completely restructured within the next 1-2 years. We furthermore plan to develop a much faster detector simulation code based on parameterized shower profiles etc. to be used for studying low cross section phenomena.

Event Handling The event handling package will contain all the aspects of the primary data flow, like calibration, event reconstruction and analysis. This part of the off-line code is likely to contain the most lines of code. It will be developed primarily by PHENIX physicists under the guidance of the off-line group. Currently a nearly fully functioning prototype (PISORP) has been developed.

12.5.3 Core Modules

The high level software modules will be attached to a software "bus" consisting of a large set of core modules:

Off-line frame:	All basic routines providing the backbone of the event handling package.
I/O modules:	All modules interfacing to the data structures and the users.
Database manager:	All modules providing interface and control over the databases.
Graphic modules:	All modules providing graphical display of the data flowing through the off-line system: Event display, histogram display, GUI etc.
Production control:	Command procedures for handling the automatic flow of data through the system.
PHENIX utilities:	All modules performing utility tasks specific for PHENIX.
General utilities:	External application packages performing general utility functions like the CERN libraries, GNU software etc.

on- and off-line systems [5], and GEM estimates their combined computing needs to be in the order of 500 man-years.

As stated earlier in Section 12.2, the off-line system will be designed based on a *design-to-cost* philosophy, which implies that the functionality and degree of sophistication will depend strongly on the amount of allocated manpower and hardware. A system designed completely by captive labor without any professional help would barely meet the minimum requirements, and would in the long run be very expensive due to the lack of reliability and increased maintenance problems.

The overall schedule for the off-line system estimates that the main design will be finished by late summer 1994. The first fully operational prototype will be complete early 1996, and the final implementation is finished by the end of 1997. Integration and testing will occur in 1998.

Bibliography

- [1] *Off-line Computing for RHIC Experiments*, Bill Love et al., RHIC Detector Note 8, 1992.
- [2] *SDC: Technical Design Report*, SSCL-SR-1215, 1992.
- [3] *RHIC Off-line Computing, Study Group Interim. Report*, J. Featherly et al., BNL, June 1993.
- [4] R. Vogt, *Atom. Nucl. Data Tabl.* **50**, 343 (1992).
- [5] Bruce Gibbard, BNL; personal communication.
- [6] *GEM Technical Design Report*, SSCL-SR-1219.

Chapter 13

Systems Engineering and Integration

13.1 Overview

The PHENIX experiment comprises 11 different detector systems using a variety of different technologies. These subsystems are being developed and constructed by separate groups within the PHENIX Collaboration. Integration, installation, testing, and commissioning for these systems are tasks critical to the success of PHENIX.

The Systems Engineering (SE) group at BNL is responsible for these tasks. Additionally this group designs the support structures, common systems, and facilities. The SE group also coordinates the interface with the accelerator, and manages documentation, quality assurance, and safety issues.

13.2 Changes

The SE group's responsibilities have been broadened since the Conceptual Design Review. Chapter 12 in the CDR (titled Structures and Facilities) is expanded here to address these changes. There are three new major headings in WBS section 5.5. They are Integration, Installation, and Test and Commissioning. There are significant changes in the detector design that affect SE planning. The major changes and progress since the CDR include:

1. Detectors and Carriages

- The configuration of the outer tracking and EMCAL modules is changed from two to four segments per arm. There is a corresponding decrease in ϕ coverage per segment from 45° to 22.5° .
- The ToF will be placed only in front of the PbGL.
- Radial space allocations have been adjusted.

2. Facilities

- The new counting house is now planned as a single story building with reduced floor space. Specifications have been defined and submitted to BNL Plant Engineering and design is well underway.
- Utility requirements have been defined.

3. Integration

- The SE group coordinates mechanical and electrical specifications to guarantee compatibility and fit of the subsystems.
- A system of control drawings is in effect.
- An archiving system for drawings from all participating groups has been implemented.
- A parameter database is being prepared.
- A review program has begun to insure compatibility of the various systems and minimize duplication.

4. Installation

- The installation plan, while not significantly changed has been considerably refined.
- The mechanism for moving the central magnet and the detector carriages is improved to reduce the downtime required for maintenance.

5. Test and Commissioning

- The SE group is responsible for coordination of testing and commissioning of system components.

13.3 Detector Systems and Central Arm Carriages

Figure 13.1 shows a "beams-eye" view of the central arm detectors with the latest segmentation. The most significant change is that the outer detectors (DC, PC, TEC, and EMCal) are now divided into four segments, instead of two. Space is provided for the TOF only in front of the PbGl modules of the EMCal. The PbSc modules are now located slightly closer to the beam, just outside the respective PC3 Pad Chambers. The structural-design concept for the support carriages remains basically unchanged, except that the carriage shapes follow the new detector outline. Diagonal-stiffening bars (square steel tubing) have been added to increase the natural frequency of the loaded structure and reduce deflections.

For several of the detector systems the new segmentation plan results in smaller individual modules, which are more easily handled. In particular, the smaller EMCal walls are now within the capacity of the 30-ton crane in the high bay of the Multi-use Laboratory, where PHENIX subsystems are to be assembled.

Figure 13.1: Central magnet and central arms on the RHIC beam line showing detector support carriages with one central arm in data-taking position (left) and one retracted (right).

The muon ID is now five layers of steel, instead of the two layers of concrete shown in the CDR. This change does not affect the placement of the Muon Magnet (MM) or other detectors.

13.4 Facilities

13.4.1 Counting House

Figure 13.2 shows the current design for a single story counting house. Floor space has been trimmed to about 485 m², which is the minimum necessary to house the data acquisition electronics, computing systems, and a control room with a reasonable amount of work space. Removing the second story resulted in a major cost savings since the building is now compatible with the footings of the MFH. Additional needs are to be provided by temporary structures and trailers, or in existing buildings at BNL.

13.4.2 Utilities

AC Power

The total AC power requirement for the detector systems in the MFH is approximately 300 kW. This includes detector high voltage, front-end electronics, data acquisition electronics through the DCM's and LVL-1 trigger, as well as control electronics. This AC power must be an isolated feed, which is independent of magnet power and utility power for air conditioning, pumps, and general use. The utility-power requirements are specified at 500 kW. This includes an air conditioning and cooling load of 400 kW. These loads plus the magnet load of 1200 kW are within the capacity of 2500 kW already available in the MFH. However, provision must be made for the isolated feed. As information pertaining to the specific power requirements for each detector becomes available, distribution panel capacities and locations will be specified. Power requirements for the magnets is discussed further in Chapter 4.

The power requirements for the data acquisition electronics, computing facilities, and control consoles in the counting house are 250kW. This power must also be isolated from the utility feeds. These requirements along with the air conditioning loads are folded into the specifications for this new building.

Cooling Water

Cooling water at or below 20 C° is required for the magnets. With all coils operating, the cooling load is about 1.2 MW at a maximum flow rate of 800 L/min. The RHIC project will provide cooling-tower water including pumping. The cost of a supplementary chiller was estimated, but it is not included in the baseline detector.

Climate Control

The air conditioning load used in the cost estimate is 40 tons. Most of the detectors expect to use air cooling with a resulting heat load increase on the MFH air conditioning requirements. Additional units will have to be installed to bring the net available capacity of the MFH up to the levels required to accommodate both detector and personnel demands. Capacities in the counting house are sized accordingly as part of the construction specifications. Additional individual chilled water units may be used for several of the detectors and to cool enclosed

Figure 13.2: The proposed PHENIX counting house.

equipment racks in the counting house. Because these will exhaust heated air outside the facility they are not expected to be a concern for air conditioning.

13.5 Integration

The SE group will provide a number of integration functions within the PHENIX project.

1. Establish a basis for common terminology.
2. Define tolerance specifications and requirements.
3. Manage distribution of technical information such as parameters, and specifications.
4. Integrate physical envelopes, interface requirements and resolve conflicts.
5. Ensure that tasks are performed and completed in an organized fashion to prevent time/schedule and spatial conflicts.
6. Determine support requirements and ensure that provisions are made to meet them.
7. Maintain record of progress, new requirements, changes, and problems as they develop in order to ensure that a timely solution can be found.

Detector integration is coordinated by the SE group which negotiates, generates, and distributes the system level design and layouts. These are in the form of a series of envelope and interface control drawings. Drawing numbering and design standards are monitored and controlled to assure compatibility with RHIC standards. All mechanical drawings are electronically filed and archived with Autocad being the standard CAD system. Electronic drawings must be compatible with View-logic, P-CAD, or Autocad. Figure 13.3 shows the steel section envelope control drawing currently used to allocate space in the r - z plane between the magnets.

A database of detector parameters is being developed to centralize and make readily available the major specifications and physical characteristics of the experiment. This database will be updated as frequently as necessary to provide the collaboration with a current source of information about the apparatus. It will be kept on-line and accessible via the network to all group members. This database will be a basic reference document for PHENIX.

For consistency and compatibility, and to minimize duplication, the SE group reviews all specifications and requirements for common needs such as facilities, utilities, survey, alignment, etc.

13.6 Installation

While each subsystem group is responsible for, and directs the installation of their detector, the SE group at BNL plans and coordinates the overall installation process. Many components, from massive magnet parts to fragile detectors, will be arriving at the site over the

Figure 13.3: Envelope control drawing showing magnets in r - z plane.

two year installation period. The installation sequence and the myriad of activities in and around the MFH must be well choreographed to insure that PHENIX is ready for beam on day one. In addition, the group will provide support, as required, for rigging, survey, and technician labor to work under the supervision of detector subsystem personnel.

Figure 13.4 shows that the concept for moving the CM and detector arms has been considerably refined since the CDR. A system of permanent and temporary tracks with guides and a hydraulic drive system allow these large assemblies, supported on Hilman rollers, to be moved with relative ease, both for installation and maintenance. A major purpose of this system is to allow access to all detectors for normal maintenance with minimal impact on the running schedule.

Permanent tracks, fixed to the floor in the collision hall, serve for both installation and normal maintenance. Temporary track sections are used for installation and removal of the detectors or CM from the collision hall to the assembly area.

The drive and lift system provides the horizontal travel and vertical lift actuation required

Figure 13.4: Illustration of the concept for moving the CM and detector carriages.

to move the CM or detector carriages. The drive system comprises pairs of actuators which drive horizontally from incremented attachment points on the tracks to corresponding points on the carriages or CM. Hilman rollers running in lateral guides on hardened steel tracks provide accurate low friction action. The lift system comprises jacks which lift the carriages or CM from base corner points thus allowing rotation of the Hilman rollers to change rolling direction at intersections in the tracks. This rotation operation is necessary only for removal of the detectors or magnet from the collision area.

Prior to the initial installation of the CM and detector carriages, the muon magnet and the muon ID absorber will already be in place. The permanent track system will be permanently affixed to the MFH floor. The temporary tracks are put in place and the CM is rolled into position. While the fields of both magnets are mapped the individual detectors are tested and installed into the carriages in the MFH assembly area. Prior to rolling the carriages into the collision hall, the CM must be removed to the assembly area to allow the far side carriage access to its position on the outside of the beam line. Then the CM is rolled in, followed by the near side carriage. The CM and carriages are locked down whenever they

are not actually being moved. Finally all detector components are surveyed in place.

For normal maintenance, the permanent section of track in the collision hall will allow each detector carriage to be retracted laterally from the Central Magnet by 0.82 m. With the carriages retracted, the CM will be free to travel away from the Muon Magnet by 1.5 m in z . These large assemblies can then be rolled back into their data-taking positions with an accuracy of better than 1 mm. Such good tolerances eliminate the need for realignment of individual detectors. Only a quick survey and a straight-through-tracks run will be required to resume data taking. It will be possible to retract and replace the carriages and the magnet including required alignment and survey in one day. This will allow reasonably quick access to the BB counter on the muon end and Station 1 of the muon tracking system. Also both BB counters can now be mounted symmetrically to the CM.

To minimize the time required to accomplish retraction of the arms and central magnet, and to facilitate testing, the utility and data connections to both the CM and the carriages are maintained through umbilicals. Since the Hilman rollers rest on hardened tracks, vertical jacking is not required prior to moving these systems. This allows z movement of the CM without removing the MVD or BB.

For major maintenance or upgrades requiring removal of one or more detectors, the carriages will be moved from the collision hall to the assembly area. To accomplish this, the shield wall must be removed, the beam pipe disassembled, and temporary track segments laid in place. Then the near carriage, the CM, and the far carriage are rolled out in sequence. During the maintenance period, the shield wall and the beam pipe can be re-assembled and accelerator operation can resume. All of this can be accomplished within a standard 4 month shutdown.

13.7 Test and Commissioning

The subsystem groups are responsible for testing and commissioning their respective detectors. To minimize interferences and schedule/manpower conflicts, the SE group will coordinate the post-installation testing and commissioning process. Close collaboration with each subsystem group is required to plan the testing schedule, establish test parameters and specify acceptance criteria for the respective detector systems. The test and commissioning plan will be approved by the DC members, Project Management, Engineering, and Safety/QA.

13.8 Cost and Schedule

The cost estimate and schedule for systems engineering and integration is presented in Table 13.1. In this table the carriage (WBS 5.5.1), facility (WBS 5.5.2), integration (WBS 5.5.3), installation (WBS 5.5.4), test and commissioning (WBS 5.5.5), and safety (WBS 5.5.6) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

Because there are three additional sections (Integration, Installation, and Test and Commissioning), as discussed above, the costs are not directly comparable to the numbers in CDR. Installation costs for the magnets and the muon ID absorber are not included in WBS 5.5, but rather in WBS 5.1 under magnets (see Chapter 4).

Table 13.1: Cost Estimate and Schedule for PHENIX Systems Engineering and Integration

WBS	Description	Total Cost	1993			1994				1995				1996				1997				1998	
			Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2
5.5.1	Detector Carriages	\$554,787.16																					
5.5.1.1	Design/Analysis	\$257,314.91																					
5.5.1.2	Prototype-Not Req'd	\$0.00																					
5.5.1.3	Fabrication / Assembly	\$258,714.11																					
5.5.1.4	Testing/Debugging	\$38,758.14																					

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998				19			
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2		
5.5.2	Facilities	\$1,545,614.30																										
5.5.2.1	Design	\$84,565.53																										
5.5.2.2	Prototype - Not Required	\$0.00																										
5.5.2.3	Fab./Procure	\$946,162.91																										
5.5.2.4	Install	\$480,153.21																										
5.5.2.5	TEST	\$34,932.65																										

WBS	Description	Total Cost	83																											
			1993			1994			1995			1996			1997			1998			1999									
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3			
5.5.3	Integration	\$1,244,610.46																												
5.5.3.1	Technical Oversight	\$573,596.66																												
5.5.3.2	Interface control	\$90,689.83																												
5.5.3.3	Parameter Book	\$25,355.79																												
5.5.3.4	RHIC/Phoenix Interface	\$213,034.12																												
5.5.3.5	Travel/Liaison	\$18,633.31																												
5.5.3.6	Coordination	\$6,189.76																												
5.5.3.7	Facility Requirements	\$39,030.14																												
5.5.3.8	Survey/Alignment	\$25,954.30																												
5.5.3.9	QA Review	\$252,135.55																												
5.5.3.10	Survey/Align Completed	\$0.00																												

[illegible][illegible]

WBS	Description	Total Cost	1994			1995			1996			1997			1998			1999					
			Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	
5.5.6	Safety & Environmental Protection	\$161,456.66																					
5.5.6.1	Procedures & Control	\$161,456.66																					

Chapter 14

Project Management

14.1 Introduction

The basic structure of PHENIX Project Management (WBS 5.6) is unchanged since the CDR was submitted. There is some reduction in the cost estimate for project management, which results from two changes. First, some scientific manpower at BNL, which was formerly carried under project management, is now moved into the appropriate WBS areas. Second, there is expansion and redefinition of the area of effort now called System Engineering & Integration (WBS 5.5). Some tasks formerly listed under Project Management, including some related to safety and quality assurance, are now in WBS 5.5.

The main components of the Project Management cost are several categories of manpower, including scientists, engineers, administrative personnel, and technical personnel, as well as project-related travel and communications. These are discussed in more detail below.

The current cost estimate for project management is based on a 69-month (July 1993 through March 1999) construction and installation phase. Construction funds are expended on this effort starting in July 1993.

14.2 Personnel

14.2.1 Scientists and Engineers

The Assistant to Project Management, who is a Ph.D. physicist, is supported on PHENIX construction funds as of October 1993. There are a number of engineering functions not associated with specific WBS elements in the detector subsystems. For example, the position of Project Engineer, who is also the director of the Project Office, is described below.

14.2.2 Technical and Administrative Personnel

Technical personnel include a PHENIX Building Manager and a Technical Supervisor, whose jobs are to oversee the preparation of PHENIX space on site and to supervise the technical effort associated with detector installation. The project also requires a full-time equivalent

Table 14.1: Cost Estimate and Schedule for the PHENIX Project Management

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998				1999		
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3
5.6.1	Scientific Management and QA	\$1,712,553.50																									
5.6.1.1	Baseline definition	\$688,430.24																									
5.6.1.2	Project documentation	\$547,920.00																									
5.6.1.3	Administrative assistance	\$426,160.00																									
5.6.1.4	Computing Equipment	\$50,043.36																									

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998				1999		
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3
5.6.2	Site Management	\$397,486.16																									
5.6.2.1	Coordination of site preparation	\$178,200.00																									
5.6.2.2	Consumables	\$219,286.16																									

WBS	Description	Total Cost	93		1994				1995				1996				1997				1998				1999		
			Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3
5.6.3	Budget tracking and reporting	\$521,526.00																									
5.6.3.1	Internal budget tracking	\$343,326.00																									
5.6.3.2	Tracking & reporting	\$178,200.00																									

(FTE) for document and drawing control, budget and schedule tracking, and general project software maintenance.

A number of other administrative functions, including procurement, contracts, and secretarial duties, are expected to be supported through existing departments and divisions of BNL, including RHIC and the Physics Department.

14.3 Cost Estimate and Schedule

The cost estimate and schedule for PHENIX project management is presented in Table 14.1. In this table the scientific management and quality assurance (WBS 5.6.1), the site management (WBS 5.6.2), and the budget tracking and reporting (WBS 5.6.3) cost are shown at PHENIX WBS level 4. The detailed cost estimate, along with the schedules and profiles, is available in the PHENIX WBS book. The dictionary, basis of estimate, and contingency analysis are presented in the PHENIX WBS Notes book.

The total cost estimate for project management is \$2,360k in FY 1993 dollars, although nearly half this sum is covered by BNL operating funds. The cost in terms of RHIC project funds is \$1,298k. The manpower resources for all the positions described above are in place.

Under the assumptions used in estimating the cost of WBS 5.6, there is a continuous and uniform level of effort throughout the project.

Chapter 15

Schedules and Profiles

Cost estimates, schedules, and funding or resource profiles are presented in the WBS Book, which accompanies the PHENIX Management Plan. The WBS Book and the WBS Notes Book (dictionary, basis of estimate, and contingency analysis) will be available to the Technical Advisory Committee at the Cost & Schedule Review, which will be held at BNL on 8-11 November 1993. A version of the summary cost estimate will be sent in another mailing to the Technical Advisory Committee so that they have it before the Cost & Schedule Review. The numbers below represent the status of the cost estimate at the time of the printing of the CDR Update.

15.1 Summary Cost Estimate

The summary cost estimate shows the PHENIX WBS level 3 (RHIC WBS level 4) costs and funding sources. The total cost (assuming all elements of the baseline were built for the US-based cost estimate) exceeds \$71M. In terms of construction funds from the RHIC Project, the cost in FY1993 dollars is \$32.5M plus contingency. The bottoms-up contingency, described in the Management Plan and shown in the WBS Book, is 22% when averaged over the project. Therefore, with contingency the total cost in Project funds is \$39.6M. This is approximately \$4.1M over the budget; plans for addressing this shortfall are described in the Management Plan.

Additional funding sources include:

1. **Institutional Contributions:** Resources (primarily engineering and technical manpower) supported on the operating funds of several US and Canadian collaborating institutions.
2. **Japanese Funds:** The member institutions of the PHENIX-J collaboration are contributing several detector subsystems built with Japanese funds.
3. **Other Foreign Contributions:** Non-US collaborating institutions are building several detector subsystems with a combination of PHENIX project funds and their own resources. This results in a cost reduction (relative to US cost estimates for these subsystems).

The additional funding sources listed above account for approximately \$28.3M. In addition, there are deferred items totalling \$10.3M.

The estimated cost (with contingency) of \$39.6M is to be compared with \$34.8M at the time of the CDR submission in January 1993. Many factors contribute to this 14% increase; the major ones are:

1. The stretch-out of the project to a new completion date of March 1999 had the effect of increasing the cost of manpower.
2. Returning to a five-layer steel μ ID absorber increased the cost of this item.
3. Integration and installation tasks have been pulled together under WBS 5.5. In the process the cost increased, primarily due to the inclusion of previously overlooked efforts and resources.

Chapter 16

Upgrade Options

16.1 Introduction

The baseline detector configuration of PHENIX is presented in Chapters 4 through 12 of this document. It is capable of performing a very large fraction of the PHENIX physics menu, as shown in Chapter 3. The detector design concept is an open one, allowing for the easy integration of future upgrades.

Several possible upgrades to the baseline were identified in the CDR, namely:

- Conversion of the four-layer TEC to a six-layer transition radiation detector/tracker (TRD).
- Addition of a higher-resolution, finer-grained electromagnetic calorimeter.
- Addition of a hadron-blind detector for Dalitz rejection at small radius.

There has been no change in the list of potential upgrades since the submission of the CDR. There has also been no attempt at this time to make new cost estimates for the upgrade items along the lines of those prepared for the present update. Similarly, there has been no attempt to put a priority order to this list for the same reasons given in the CDR: "We expect the order of implementation of these and other possible upgrades to depend on a number of factors which cannot be determined at present. These include future developments in physics: Theoretical, experimental (both before we have data from RHIC and after), and technical. In some cases, particular upgrades will be the subject of specific proposals for additional funding; the fate of these proposals will also influence the order of implementation." Below we discuss briefly recent developments in each.

16.2 Tracking TRD

The upgrade of the PHENIX four-plane TEC to a six-plane tracking transition radiation detector (TRD) greatly expands the momentum region over which PHENIX has e/π separation capability. The tracking TRD can separate electrons from pions starting at 1.0 GeV/c and

continuing above 50 GeV/c. Also, the dE/dx particle identification of the TEC improves by almost one order of magnitude when going from four to six active layers. Since the PHENIX baseline detector has limited electron identification above 3.5 GeV, the TRD upgrade improves the experiment's ability to measurement the J/ψ , Υ , and Drell-Yan continuum.

The PHENIX TEC has been designed to accomodate the upgrade to a tracking TRD in a straightforward way. An open area exists between each TEC plane where the TR radiator packets will be inserted. Support structures for those packets are incorporated into the TEC frames as are structures needed to support the TEC windows under xenon gas hydrostatic pressure. The front-end electronics for the TEC are being designed to handle a dynamic range covering energy deposits from both minimum ionizing particle dE/dx and transition radiation x-rays (Chapter 6.5.2). The 5-bit flash ADCs instrumenting each anode wire will have 28 levels set in the lower dE/dx range and three in the higher transition radiation range. Additionally, the gas system for the detector will initially purge the TEC P10 gas mixture into the air but is designed to be upgraded to a fully recycling system when xenon gas is introduced into the TRD.

The scheme for upgrading from the TEC to the tracking TRD will involve the instrumenting of the final two TEC planes with electronics, installing radiator packets between the planes and changing the gas system into a fully recycling one. The detector gas will be changed from an argon-based mixture to a xenon-based mixture. At the present stage in the detector design, it isn't clear whether the TEC will have to be disassembled to allow for the installation of the radiators. The other aspects of the upgrade, however, involve little disturbance to the detector.

16.3 High Resolution Photon Detector

16.3.1 Overview

Since the completion of the CDR, a detailed investigation as been carried out to study the design and performance of the high resolution photon detector for PHENIX. The results of this study are summarized in detail in PHENIX Note 79 [1], and in the proposal [2] recently submitted to the National Science Foundation. Here, a short description of the status of the proposal will be given, along with a brief discussion of the technical progress which has been made towards developing the high resolution detector.

16.3.2 Changes

Progress has been made in several areas since the PHENIX CDR was submitted.

1. A proposal to build the high resolution PHENIX Plasma Photon Detector (P^3D) has been submitted to the National Science Foundation.
2. The calibration scheme has been firmly established.
3. The timing resolution of CsI and BaF_2 crystals has been measured and found to be excellent.

4. The first full-size BaF₂ crystals have been ordered and tested.
5. The CsI prototype detector has now taken data under real experimental conditions in the BNL-AGS Au + Au program.

NSF Proposal

A proposal to build the PHENIX Plasma Photon Detector (P³D) was submitted to the National Science Foundation in May 1993. The proposal is a collaboration of 11 institutions consisting mainly of PHENIX collaborators, along with several other groups which are not part of the main PHENIX experiment. Its purpose is to build the high resolution photon detector described in the CDR which will extend the p_T coverage for photons into the region below 1 GeV/c. This will substantially extend the search region for direct photons from the QGP and the mixed hadronic phase beyond what will be attainable with the baseline PHENIX detector.

The proposal was sent out for peer review early this summer by the NSF. The review process is continuing, but the initial reaction from the NSF has been encouraging. The total amount of funding requested for the detector and accompanying system integration into the PHENIX experiment is approximately \$7.1M.

Progress on Detector Design

Development work on the high resolution detector has proceeded with additional study and testing of the multicrystal prototype detector. A complete description of the entire set of test results is given Chapter 5 of reference [1]. We have firmly established that an energy resolution of $< 2\%/\sqrt{E}$ can be achieved under test beam conditions down to very low energies. The actual resolution measured over the range from 50 MeV to 1 GeV was:

$$\frac{\sigma(E)}{E} = 0.8\% \oplus \frac{1.8\%}{\sqrt{E}} \oplus \frac{0.3\%}{E}$$

It was also shown, as discussed in Chapter 4 of ref.[1], that the excellent energy resolution of the detector can be preserved even in the high multiplicity environment at RHIC. It was found that the energy resolution for an isolated photon in the 200–500 MeV range is increased from 1.8% to 2.2% (i.e., the increase is less than 50%), and its efficiency decreased by no more than about a factor of two, even in events with triple the average HIJET multiplicity.

Additional test beam data have also been obtained on the time resolution for high resolution crystals at low energies. A series of beam tests were carried out at TRIUMF in March 1993 in which the time resolution of both undoped CsI and BaF₂ crystals were measured for electrons, muons and pions. Figure 16.1 shows some of the results obtained for electrons at a momentum of 150 MeV/c. The uncorrected time resolution for CsI was $\sigma = 159$ ps versus $\sigma = 110$ ps for BaF₂. After correcting for the time resolution of the start counter which defined the beam ($\sigma = 95$ ps), the corrected time resolutions were 127 ps and 55 ps for CsI and BaF₂, respectively. This measurement clearly demonstrates the much better time resolution obtainable with BaF₂, due to its much faster ($\tau < 1$ ns) decay time. It should

Figure 16.1: Time resolution for 150 MeV/c electrons measured at TRIUMF. Top: CsI ($\sigma = 159$ ps). Bottom: BaF₂ ($\sigma = 110$ ps). After unfolding the time resolution of the start counter ($\sigma = 95$ ps), the corrected values are 127 ps and 55 ps, respectively.

be noted that BaF₂ is one of two main options for the type of crystal to be used in the final detector, and additional prototype work is presently under way for further study the performance of BaF₂ in a broader range of tests. Moreover, high quality BaF₂ crystals up to 17 radiation lengths in length (35 cm) are presently available from China at an attractive cost ($\sim 3/\text{cm}^3$), and it is possible that suitable phototubes with the appropriate gain and window material may be available from Russia at a much lower cost than what is available from western suppliers. We intend to further explore all of these various cost saving options in order to try and reduce the overall cost of the detector.

A further test of the performance of the high resolution detector in the actual heavy ion collisions was performed using the 11 A-GeV gold beam at the BNL-AGS in October 1993. Our existing prototype detector, consisting of 30 undoped CsI crystals, was installed in experiment E877 and was used to measure photon production in central Au + Au and Au + U

collisions at forward rapidities. The detector took data with a p_T threshold below 1 MeV, demonstrating that it is possible to operate a detector such as this with good sensitivity at very low energies, even in a very high multiplicity environment. We intend to use these data to measure the inclusive photon spectrum in central gold collisions and try to place some limit on the level of direct photon production at very low p_T . This will be an important test of how the detector performs under actual running conditions, as well as our ability to extract real physics information from the data obtained. In addition, it will exercise other aspects of the detector, such as the laser calibration system, for which similar systems will be used elsewhere in the PHENIX experiment.

16.4 Hadron-blind Detector

16.4.1 Introduction

The hadron-blind detector (HBD) is a device intended to measure electrons at small radius inside the CM field region around the collision point. The goal is to improve rejection of the dielectron background due to random combinations of electrons from photon conversion and Dalitz decay of π^0 . Rejection would be improved by detecting the soft partners of such electrons that curl up in the field and are not seen in the central arms.

The HBD implementation under study is the same as presented in the CDR. It is a series of cylindrical gas volumes in which electrons produce Cherenkov light. The photons are detected with a parallel-plate avalanche chamber (PPAC) with a solid CsI photocathode at the outer radius of each gas volume.

Research and development on the HBD was actively pursued during the past year on this idea, with encouraging results which were submitted for publication to NIM. The prototype tested is quite insensitive to hadrons but did not generate sufficient photoelectrons for good efficiency. Several improvements are incorporated in the version to be studied during the coming year:

1. A better method of laying down the photocathode was used by the PHENIX group studying similar photon detectors for the RICH. This will be used in the next HBD prototype.
2. A new PPAC with a wire configuration giving a more uniform E-field will be used.
3. Several different gas mixtures will be studied to optimize the photoelectron yield.
4. Faster amplifiers will be used to take advantage of the different time structure of electron and ion signals in the PPAC.

A testbeam area is being set up at MIT's BATES Laboratory to provide better access for the HBD, and other PHENIX R&D efforts, to a high quality, low energy electron beam.

Bibliography

- [1] A High Resolution Photon Detector for PHENIX at RHIC, April 30, 1993, PHENIX Note 79.
- [2] Phenix Plasma Photon Detector (P³D), proposal submitted to the National Science Foundation, May 1993.

Chapter 17

Collaboration and Organization

17.1 PHENIX Collaboration

The PHENIX collaboration consists of over 350 scientists, engineers, and graduate students from 44 institutions in 10 countries. The updated institutional distribution of the PHENIX collaboration is shown in Table 17.1.

17.2 Management Structure

Figure 17.1 shows an updated organizational chart of the management structure for the PHENIX experiment. The function of each box in the organizational chart in Fig. 17.1 is described in the PHENIX Management Plan. The current members of the PHENIX management (Project Management, Executive Council, and Detector Council) are listed in Table 17.2. Institutional responsibilities are described in the Management Plan.

PHENIX Collaboration

Brasil:		Russia:	
U. Sao Paolo	3	IHEP-Protvino	24
		INR-Moscow	7
Canada:		ITEP-Moscow	8
McGill U.	10	JINR-Dubna	23
		Kurchatov Inst.	9
China:		PNPI-St. Petersburg	22
CIAE	13	Individual	1
IHEP	8		
Inst. Mod. Phys.	8	Sweden:	
Peking U.	4	Lund U.	8
Germany:		U. S. A.:	
U. Münster	7	U. Alabama	4
Individual	1	BNL	25
		UC-Riverside	5
India:		Columbia U.	14
BARC, Bombay	3	Florida State	3
		Georgia State	3
Japan:		Idaho NEL	4
Hiroshima U.	8	Iowa State / Ames Lab.	9
INS, U. Tokyo	5	LLNL	10
KEK	7	LANL	16
Kyoto U.	2	Louisiana State	4
Nagasaki Inst.	2	MIT	9
Nat. Inst. Rad. Sci.	1	SUNY-Stony Brook	12
U. Tokyo	6	ORNL	17
U. Tsukuba	11	U. Tennessee	4
		Vanderbilt U.	4
Korea:		Yale U.	9
Chung-ang U.	1	Individual	1
Korea U.	4		
Seoul Nat. U.	6		
Soong-Sil U.	1		
		Total	366

Table 17.1: Institutional participation in the PHENIX collaboration.

Figure 17.1: Organizational chart of the PHENIX experiment.

Project Management

Spokesperson:	S. Nagamiya (Columbia), nevis::nag
Project Director:	S. Aronson (BNL), bnlcl6::aronsons
Deputy Project Director:	G. Young (ORNL), orph01::young
Project Engineer:	L. Paffrath (BNL), bnldag::paffrath
Assistant for Project Management:	W. Kehoe (BNL), bnlcl6::kehoe

Liaison Physicist:	B. Johnson (BNL), bnl::brant
--------------------	------------------------------

Executive Council

P. Braun-Munzinger (Stony Brook)	Z. X. Sun (CIAE, Beijing, China)
J. Moss (LANL)	Z. P. Zheng (IHEP, Beijing, China)
F. Plasil (ORNL)	R. Santo (Münster, Germany)
R. Hayano (Tokyo, Japan)	I. Otterlund (Lund, Sweden)
K. Yagi (Tsukuba, Japan)	S. K. Mark (McGill, Canada)
A. Vorobyov (St. Petersburg, Russia)	S. S. Kapoor (BARC, Bombay, India)
V. Manko (Kurchatov, Russia)	

Detector Council

Magnet:	J. Thomas (LLNL)	Inner Det.:	B. Jacak (LANL)
Tracking:	E. O'Brien (BNL)	RICH:	H. Hamagaki (INS, Tokyo)
TOF:	Y. Miake (Tsukuba)	EM Cal:	S. White (BNL)
Muon:	W. Kinnison (LANL)		
On-line:	G. Young (ORNL) + W. Zajc (Columbia) + T. Carey (LANL)		
Off-line:	S. Sorensen (Tennessee)	Simulation:	C. Maguire (Vanderbilt)

Table 17.2: Members of the Project Management, Executive Council, and Detector Council for the PHENIX experiment.

Appendix A

Acronyms and Abbreviations

A.1 Physics and Simulations

Table A.1: Acronyms and Abbreviations for Physics and Simulations

Item	Description
FRITIOF	a Swedish Monte Carlo event generator (a Norse god)
GEANT	standard CERN simulation program (French for “giant”)
HBT	Hanbury-Brown-Twiss (Bose-Einstein correlation)
HIJET	Heavy Ion simulation program (derived from ISAJET)
ID	Identification
PID	Particle Identification
PISA	Phenix Integrated Simulation Application
PISORP	PISA Output Readback Program
QCD	Quantum Chromodynamics
QGP	Quark-Gluon Plasma
E_T	transverse energy
p_T	transverse momentum
p_{total}	total momentum
r	radial dimension from beam axis
T_C	critical temperature (for the transition)
y	rapidity
z	dimension along beam axis
$\Delta\Omega$	solid angle
η	pseudorapidity
ϕ	azimuthal angle
θ	polar angle

A.2 Detectors and Hardware

Table A.2: Acronyms and Abbreviations for Detectors and Hardware

Item	Description
AB	Analog Buffer
ADC	Analog to Digital Converter
AMU	Analog Memory Unit
AP	Analog Processing
ASIC	Application Specific Integrated Circuit
B	Buffer
BB	Beam-Beam Counter (Inner Detectors System)
BC	Beam Clock
CM	Central Magnet
CFD	Constant Fraction Discriminator
CMOS	Complimentary Metal Oxide Semiconductor
COB	Chip On Board
CTR1	em Fictitious Central-Arm Tracker
DAI	Data Acquisition Interface
DAQ	Data Acquisition System
DC	Drift Chamber (Tracking System)
DX	Name of a RHIC magnet closest to the collision point
DCM	Data Collection Module
DMU	Digital Memory Unit
DSP	Digital Signal Processor
EB	Event Builder
EMCal	Electromagnetic Calorimeter System
ENC	Equivalent Noise Charge
FADC	Flash Analog-to-Digital Converter
FEE	Front-End Electronics
FEUnnn	Russian PMT #nnn
FIFO	First In First Out
FOL	Fiber Optic Link
FOXFET	Field Oxide Field Effect Transistor
FPGA	Field-Programmable Gate Array

Table A.2: Acronyms and Abbreviations for Detectors and Hardware
(Continued)

Item	Description
GARFIELD	Drift-chamber simulation program
GL n	Global Level- n (trigger)
G-10	Fiber-glass-resin composite insulating board
heap	ASK GLENN
HBD	Hadron-Blind Detector
HDI	High Density Interconnects
HRX	High Resolution Crystal
HV	High Voltage
HVAC	Heating, ventilation and air conditioning
IC	Integrated Circuit
LL n	Local Level- n (trigger)
LSB	Least Significant Bit
LTS	Local Trigger Supervisor
LVL- n	Level- n for trigger
MB	Mega Bytes
MCM	Multichip Module
MCS	Monitor and Calibration Supervisor
MIMD	Multiple Instruction, Multiple Data
MM	Muon Magnet
MSB	Most Significant Bit
MVD	Multiplicity Vertex Detector (Inner Detectors System)
MUX	Multiplexer
μ ID	Muon Identification (Muon Arm)
μ P	Microprocessor
μ T	Muon Tracking (Muon Arm)

Table A.2: Acronyms and Abbreviations for Detectors and Hardware
(Continued)

Item	Description
NIM	Nuclear Instrument Module
PbGl	Lead-Glass (EMCal System)
PbSc	Lead-Scintillator (EMCal System)
PC _n	Pad Chamber # <i>n</i> (<i>n</i> = 1,2,3) (Tracking System)
PGA	Programmable Gate Array
PLA	Programmable Logic Array
PMT	Photomultiplier Tube
PPAC	Parallel Plate Avalanche Counter
PPMD	Programmable Prescalable Multilevel Discriminator
POPOP	organic compound (scintillator and wavelength shifter)
QVC	Charge-to-Voltage Converter
RICH	Ring Imaging Cherenkov Detector
RISC	Reduced Instruction Set Computer
ROC	Readout Controllers
TAC	Time-to-Amplitude Converter
TDC	Time-to-Digital Converter
TEC	Time Expansion Chamber (Tracking System)
TOF	Time-of-Flight System
TR	Transition Radiation
TRD	Transition Radiation Detector
TVC	Time-to-Voltage Converter
TTS	Transit Time Spread
UV	Ultra-Violet
VME	Versabus Module for Eurocard
VWVR	Voltage Write and Voltage Read
WLS	WaveLength Shifting
68K	Motorola 68 nnn Series Microprocessor

A.3 Organizations and Facilities

Table A.3: Acronyms and Abbreviations for Organizations and Facilities

Item	Description
AGS	Alternating Gradient Synchrotron
BARC	Bhabha Atomic Research Center, India
BNL	Brookhaven National Laboratory
CDF	Collider Detector Facility (an experiment at FNAL)
CERN	Centre European pour la Recherche Nucleaire (Switzerland)
CIAE	China Institute of Atomic Energy
DESY	Deutsches Electronen Synchrotron (Hamburg, Germany)
DOE	Department of Energy
ESB	Experimental Support Building
Ennn	Experiment #nnn at the AGS. Also, Experiment #nnn at FNAL.
FNAL	Fermi National Accelerator Laboratory
GAMS	a Russian lead-glass array at CERN
GEM	Gammas, Electrons, Muons (SSC experiment)
GSI	Gessellschaft für Schwerionen Forschung (Darmstadt, Germany)
IEEE	Institute of Electrical and Electronic Engineers
IHEP-Beijing	Institute of High Energy Physics, Beijing, China
IHEP-Protvino	Institute of High Energy Physics, Protvino, Russia
IMP-Lanzhou	Institute of Modern Physics, Lanzhou, China
INR-Moscow	Institute for Nuclear Research, Moscow, Russia
INS	Institute for Nuclear Study, University of Tokyo, Japan
ISR	Intersecting Storage Ring (at CERN)
ITEP-Moscow	Institute of Theoretical and Experimental Physics, Moscow, Russia
Idaho NEL	Idaho National Engineering Laboratory
JACEE	Japanese-American Collaborative Emulsion Experiment
JINR	Joint Institute for Nuclear Study, Russia
KEK	Institute for High Energy Physics (Koh Energy Ken), Japan
KEK-PS	Proton Synchrotron at KEK
LANL	Los Alamos National Laboratory
LEP	Large Electron-Positron collider at CERN
LBL	Lawrence Berkeley Laboratory
LLNL	Lawrence Livermore National Laboratory

Table A.3: Acronyms and Abbreviations for Organizations and Facilities
(Continued)

Item	Description
MFH	Major Facility Hall
MIT	Massachusetts Institute of Technology
NAnn	North-area experiment $\#nn$ at CERN
ORNL	Oak Ridge National Laboratory
PHENIX	Pioneering High Energy Nuclear Ion eXperiment
PHENIX-J	PHENIX-Japan
PNPI	St. Petersburg Nuclear Physics Institute, Russia
RHIC	Relativistic Heavy Ion Collider
SDC	Solenoidal Detector Collaboration at SSC
SUNY	State University of New York
SSC	Superconducting Super Collider
SSCL	Superconducting Super Collider Laboratory
UAn	Underground-area experiment $\#n$ at Sp $\bar{p}$ S at CERN
UC	University of California
US	United States
WAnn	West-area experiment $\#nn$ at CERN
ZEUS	detector (named for a Greek god) at HERA e-p collider at DESY

A.4 Management and Miscellaneous

Table A.4: Acronyms and Abbreviations for Management and Miscellaneous

Item	Description
CDR	Conceptual Design Report
DC	Detector Council
EC	Executive Council
EDIA	Engineering, Design, Inspection, Administration
ES&H	Environment, Safety and Health
MOU	Memorandum of Understanding
PAC	Program Advisory Committee
PM	Project Management
PO	Project Office
R&D	Research and Development
SE	Systems Engineering Group at BNL for PHENIX
SEAPPM	Safety and Environmental Administrative Policy and Procedure Manual
TAC	Technical Advisory Committee
TBA	To be Assigned
WBS	Work Breakdown Structure
pCDR	Preliminary Conceptual Design Report