RECEIVED
JUN 17 1997

OSTI

SLAC-PUB-7455

May 1997

Performance Issues, Downtime Recovery and Tuning in the Next Linear Collider (NLC)*

F. Zimmermann, C. Adolphsen, R. Assmann, K. Bane, D. Burke,
F.J. Decker, P. Emma, R. Helm, L. Hendrickson, S. Hertzbach ‡,
J. Irwin, H. Jarvis, P. Krejcik, M. Minty, N. Phinney, P. Raimondi,
M. Ross, J. Spencer, H. Tang, P. Tenenbaum, K. Thompson,
D. Walz, A.D. Yeremian

Stanford Linear Accelerator Center
Stanford University, Stanford, CA 94309, USA

‡ University of Massachusetts, Amherst, USA

The Next Linear Collider (NLC) consists of several large subsystems, each of which must be operational and tuned in order to deliver luminosity. Considering specific examples, we study how the different subsystems respond to various perturbations such as ground motion, temperature changes, drifts of beam-position monitors etc., and we estimate the overall time requirements for tuning and downtime recovery of each subsystem. The succession of subsystem failures and recoveries as well as other performance degradations can be modeled as a Markov process, where each subsystem is characterized, e.g., by its failure rate and recovery time. Such a model allows the prediction of the overall NLC availability. Our mathematical description of a linear collider is benchmarked against the known performance of the Stanford Linear Collider (SLC).

*Presented at the 1997 Particle Accelerator Conference (PAC97),
Vancouver, B.C., Canada, May 12-16, 1997*

*Work supported by the U.S. Department of Energy, contracts DE-AC03-76SF00515 and DE-FG02-92ER40715.

HH
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

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.

PERFORMANCE ISSUES, DOWNTIME RECOVERY AND TUNING IN THE NEXT LINEAR COLLIDER (NLC) *

F. Zimmermann, C. Adolphsen, R. Assmann, K. Bane, D. Burke, F.J. Decker, P. Emma, R. Helm, L. Hendrickson, S. Hertzbach [†], J. Irwin, H. Jarvis, P. Krejcik, M. Minty, N. Phinney, P. Raimondi, M. Ross, J. Spencer, H. Tang, P. Tenenbaum, K. Thompson, D. Walz, A.D. Yeremian
Stanford Linear Accelerator Center, Stanford University, CA 94309, USA

[†] University of Massachusetts, Amherst, USA

1 INTRODUCTION

The Stanford Linear Collider (SLC) is the only linear collider in the world. One of the primary operational limitations encountered at the SLC is the need to constantly tune up all the different subsystems, *i.e.*, injectors, damping rings, linac, arcs and the final focus. The tuning involves, for example, the optimization of the interaction-point (IP) spot size [1], and the minimization of the linac emittances. The latter are determined by a delicate balance of beam orbit, rf phases and wakefields, which is very sensitive to perturbations, such as temperature or bunch-length changes.

The preservation of a tuned-up state is difficult, as illustrated by Fig. 1 and Table 1. Figure 1 shows the average normalized (or specific) luminosity as a function of day-time for the 1996 SLC run. We observe steady luminosity increases during the day and swing shifts as well as clear drops at the 8-am and 4-pm shift changes. The latter are presumably the result of reduced attention and/or of different tuning strategies. The highest luminosity is achieved during owl shift (from 0 to 8 am), with a peak value reached around 5 am. In Table 1, we present the fraction of time

Figure 1: The diurnal luminosity during the 1996 SLC run. Shown is the luminosity normalized to 10^{10} particles per bunch. The actual luminosity would be ~ 12 times higher. A rate of 1 Z/hr is equal to about $10^{28} \text{ cm}^{-2} \text{ s}^{-1}$.

during which the SLC luminosity exceeded a certain minimum value. The table shows that the highest luminosity was achieved over less than 20% of the time, with remarkably similar numbers for the last two SLC runs.

* Work supported by the U.S. Department of Energy under contract DE-AC03-76SF00515.

Luminosity	1994/95	1996
$L > 10 \text{ Zs/hr}$	56 %	57 %
$L > 30 \text{ Zs/hr}$	42 %	50 %
$L > 40 \text{ Zs/hr}$	27 %	34 %
$L > 50 \text{ Zs/hr}$	15 %	18 %

Table 1: Fraction of time (not including overall downtime) during which a certain luminosity was exceeded in the 1994/95 and 1996 SLC runs.

2 PERTURBATIONS, TUNING AND RECOVERY

The SLC experience suggests that tuning and stabilization will be very important also for the NLC.

Magnet displacements due to ground motion or temperature changes could reduce the NLC luminosity, by steering the beams out of collision or by increasing the IP spot size. The most critical magnets are those in the final-focus system. In Ref. [2], it was shown that ground motion does not significantly impact the luminosity, because, first, it is highly correlated and, second, an orbit feedback can correct the accumulative effect of a small random component. The main sources of temperature variation in the NLC final focus are synchrotron radiation and energy loss at beam-pipe transitions, which both cause a local power deposition of about 1–3 W. A proper design will ensure that, for this heat load, the magnet temperature changes by less than 0.02 K over a few seconds or by less than 1 K over 15 minutes. In this case, fast and slow orbit feedback loops can maintain orbit and spot size, provided the offsets of the beam-position monitors (BPMs) are sufficiently stable over the above time periods. FFTB experience has demonstrated that this is possible. Since most of the tuning and stabilization is based on beam information, *e.g.*, BPM readings, the largest perturbations to the tuned-up accelerator are scheduled or unscheduled downperiods without beam.

We, thus, distinguish three types of tuning: 1) continuous tuning under normal operating conditions (*i.e.*, which has to be performed regularly to compensate for slow drifts of the rf phases, BPM readings, or power supplies), 2) recovery from a one-hour down time (DT) where the accelerator was in a stand-by mode, and 3) recovery from a 24-hour repair opportunity day (ROD) with access, repair work and major temperature changes. Tables 2 and 3 illustrate this classification for the NLC final-focus system. Table 2 presents a list of recovery steps and recovery times after a down period. Only the time spent on tuning is listed, and no time for actual repair work is included. Table 3 com-

recovery from	t_{DT} [min]	t_{ROD} [min]
check BPM polarity & offset	NA	5
activate orbit feedbacks	5	5
close FF collimators	0	0
feedb. & orbit for 90 bunches	5	5
match incoming dispersion	NA	5
measure FF emittances	5	5
coupling corr. & beta-match	0	0
turn on & phase crab cavity	NA	5
establish collisions	2	2
turn on detector	NA	5
correct IP aberrations	5	5
total	22	42

Table 2: Performance recovery times of the NLC final focus after a 1-hr down time (DT) and after a 24-hr repair day (ROD).

procedure	t [min]	T [hr]	$\Delta L/L$ [%]
multi-bunch steering	0.5	0.08	0
dispersion (x&y)	0.12	0.25	0.8
waist (x&y)	0.12	0.25	0.8
skew1 (x'y')	0.06	0.25	0.4
IP divergence	0.017	1	0
skew sexts. (x' ² y', y' ³)	0.12	1	0.2
skew2 (xy')	0.06	1	0.1
skew3 (x'y)	0.06	1	0.1
multi-bunch y-disp.	0.06	8	0.03
multi-bunch waist x&y	0.12	8	0.03
adjust FF main collimators	5	24	0.35
orbit re steering	60	100	0.25
BPM align. & offsets	30	170	0.1
sext. (x' ³ , x'y' ²)	0.12	170	0
chrom. x&y	0.12	170	0
chrom. skew (x'y' δ)	5	170	0.05
2nd order y-disp.	0.6	170	0.01
crab angle (xz')	—	170	0
match inc. dispersion	5	170	0.05
total			3.27

Table 3: Continual tunings procedures in the NLC final focus: required time t , tuning period T , and estimated luminosity impact $\Delta L/L$.

piles the continuous tuning procedures and their estimated impact on luminosity.

Similar analyses have been performed for the other NLC subsystems. The results are summarized in Table 4, which shows that the recovery time of each subsystem is on average 15 minutes after a short down period and about 40 minutes after a ROD. Therefore, the tune-up of the entire NLC takes about 2 or 6 hours, respectively. In addition, there is a luminosity loss due to continual tuning of roughly 20%.

3 MARKOV MODELS

The overall availability A_{NLC} of the NLC is the product of the availabilities of the electron and the positron system: $A_{NLC} = A_{e+} \times A_{e-}$. For an overall availability of 90%, the availabilities $A_{e\pm}$ must be 95%. Both the electron and positron system can again be subdivided into 8 distinct subsystems. To assess the availability A_{e-} (or A_{e+}), we have

subsystem	t_{DT} [min]	t_{ROD} [min]	$\Delta L/L$ [%]
systemwide	—	15	—
injectors	4	45	2.5
damping rings	16	64	2.4
compressors	15	70	3.2
main linac	17	45	4.6
collimation	25	25	4.3
IP switch/b. bend	10	15	0.9
final focus	22	42	3.3
extraction line	9	21	0
total	118	342	21.2

Table 4: Recovery times and luminosity reduction due to continual tuning for all NLC subsystems.

developed a simple Markov model [3], consisting of 9 different states (see Fig. 2). The 0th state means the beam is brought to the IP. The other 8 states refer to a failure in one of the 8 subsystems. For example, in the 8th state the injector is off. Each subsystem is characterized by three parameters: a failure rate λ_i describing the frequency at which this subsystem fails, its tuning recovery time τ_i , and an average repair time ϵ_i . The recovery from a failure of the i th subsystem proceeds at a rate $\mu_i = 1/(\sum_{j<i} \tau_j + \epsilon_i)$.

Ignoring multiple subsystem failures, the NLC Markov model consists of 9 equations:

$$dA_i/dt = \lambda_i A_0 - \mu_i A_i = 0 \quad (i = 1, \dots, 8) \quad (1)$$

$$1 = A_0 + \sum_{i=1}^8 A_i \quad (2)$$

where A_i denotes the probability that the system is in state i , and A_0 is the availability A_{e-} or A_{e+} . Equation (1) describes transitions between state 0 and state i . The total time derivative was set to zero, since we are interested in the steady state solution. Equation (2) is a normalization condition, which ensures that at any given time the system is in one of the states. Solving Eqs. (1) and (2) yields the availability A_0 . Introducing the total failure rate $\bar{\lambda} = \sum_{i=1}^8 \lambda_i$ it is

$$A_0 = 1/(1 + \bar{\lambda}\bar{\tau}) \quad (3)$$

where $\bar{\tau}$ represents an effective recovery time for the entire ($e+$ or $e-$) system. In general, the availability A_0 is a func-

Figure 2: Markov model of the NLC. 'OFF' indicates a failure of the respective subsystem.

tion of three 8-component vectors: $A = A(\vec{\lambda}, \vec{\tau}, \vec{\epsilon})$. In the following we assume, for simplicity, that all 8 subsystems have the same repair and recovery time ($\epsilon_i = \epsilon, \tau_i = \tau$).

As for the failure rates λ_i , two special cases are of interest. First assume that each subsystem has the same failure rate $\lambda_i = \lambda$. Then $\bar{\lambda} = 8\lambda$ and $1/\bar{\mu} = ((\sum_{i=1}^8 i\tau)/8 + \epsilon)$, i.e., the effective recovery time is the sum of the average tuning recovery time and the repair time. The resulting availability is $A(\epsilon, \lambda, \tau) = 1/(1 + 36\lambda\tau + 8\lambda\epsilon)$.

In the second case, we choose the individual failure rates such that the failures of all subsystems have equal luminosity impact. This implies $\lambda_i(i\tau + \epsilon) = \text{constant}$, or $\lambda_i = \lambda(8\tau + \epsilon)/(i\tau + \epsilon)$ for all i , where λ now denotes the failure rate of the injector ($i = 8$). The availability is given by Eq. (3), using $\bar{\lambda} = \sum_{i=1}^8 \lambda_i$ and the effective recovery time $\bar{\tau} = \tau_8$ determined from the recursion relation

$$\tau_{i+1} = \tau_i + ((i+1)\tau + \epsilon - \tau_i)\lambda_{i+1} / \sum_{j=1}^{i+1} \lambda_j \quad (4)$$

Figure 3: Availability $A_{e\pm}$ versus $1/\lambda$ for an individual recovery time $\tau = 1$ hr, two different repair times ϵ and two different assumptions on failure rate scaling.

Figure 3 shows that, assuming a subsystem recovery time of $\tau = 1$ hr, an availability $A_{e\pm}$ of 95% is achieved if each subsystem fails less often than once per month for zero repair time, or less often than once every 3 months, if the average repair time ϵ is 8 hours.

To benchmark our mathematical description of a linear collider against an existing accelerator, a similar analysis was performed for the SLC. Figure 4 illustrates the greater complexity of the SLC model, which arises because the electron and positron beam share a common linear accelerator, and electrons are used for positron production.

We consider two different models for the individual SLC subsystem recovery rates. In the first model (I), we assume as before that all τ_i are equal. In the second model (II), we postulate, somewhat arbitrarily, a 7 times longer recovery time for the last subsystem ('IP') i.e., $\tau_0 = 7\tau$, which has the effect of weighting the downstream end of the machine more strongly. The recovery times for all other subsystems are still taken as equal to τ . Thus, neglecting the repair time, only two parameters are left: the subsystem failure rate λ (assumed to be equal for all subsystems) and the subsystem recovery time τ .

In the 1996 SLC run, there was a hardware failure or interruption about once every 2.5 hours. As a measure of the SLC or 'IP' availability, we take the fraction of time where both beams were actually in collision: about 57%. Using the total failure rate and the measured availability, we can deduce the subsystem recovery time τ . We attribute the total failure rate uniformly to all 10 SLC subsystems, i.e., $\lambda = 1/25 \text{ hr}^{-1}$, and find the recovery time τ for which the predicted availability equals 57%. This time is 30 or 10 minutes for the two different models, respectively.

Predicted and actual beam availabilities for different SLC subsystems are compared in Table 5, where the prediction is the sum $\sum_{i \leq j} A_i$, describing the probability that the beam is delivered either to the j th or to any downstream subsystem, and the actual beam availability was estimated by the time fraction for which beam was present. While the data appear to favor models with a longer recovery time for the last subsystems (model II), they would also be consistent with upstream systems being tuned over extended periods of time.

In conclusion, our approach to describe the tuning states of a linear collider by a Markov model looks very promising. In the future, the model presented here should be extended so as to include information about other beam parameters, e.g., emittances, energy and bunch lengths, and it should be further developed to understand the impact of multiple subsystem failures.

subsystem	May 96	model I	model II
e- damping ring	0.94	0.90	0.93
e+ damping ring	0.82	0.76	0.82
e- arc	0.89	0.78	0.83
'IP'	0.57	0.57	0.57

Table 5: Comparison of actual beam availability for various SLC subsystems with that predicted by two different models; the recovery-time τ was adjusted to give equal IP availability.

Figure 4: Markov model of the SLC

4 REFERENCES

- [1] P. Emma, L.J. Hendrickson, P. Raimondi, F. Zimmermann, "Limitations of Interaction-Point Spot-Size Tuning at the SLC", these proceedings (1997).
- [2] J. Irwin and F. Zimmermann, EPAC96, p. 516 (1996).
- [3] See, for example, Roger D. Leitch, Reliability Analysis for Engineers, Oxford Science (1995).