

ISTK Overview and the SDS Document

RHIC AP Note 29

Chris Saltmarsh

salty@bnl.gov

June, 1994

The two documents here were produced some time ago to give an overview of the ISTK components and a vision of how they were to work together. Although technical details of the components have changed, the broad concepts are constant. Glish is fully documented in its user manual by Vern Paxson (vern@ee.lbl.gov, RHIC AP Note 30) and the SDS man pages are being updated for the next release. An appendix to this document contains the SDS man pages as they are at present. The graphics components, being based on an old version of a C++ graphics class library called InterViews, are becoming obsolete, although several applications are still of real use and the broad structure contains useful ideas. In particular, the internal communication classes which allow transfer of the event-oriented programming of Glish into process space will continue to evolve and be used.

Contents

1. The Integrated Scientific Toolkit: An Overview
2. The SDS Manual

The Integrated Scientific Tool Kit: An Overview of the System and its Components

Ericka Lutz and Chris Saltmarsh
August, 1992

Table of Contents

Introduction to the Integrated Scientific Tool Kit	4
Document Format.....	5
The ISTK Approach Track 1	
A Cost Effective Approach.....	6
Deciding Which Components to Use	6
Approach Track 2	
More About ISTK and the Unix-Style Approach	7
ISTK Communications Systems	
Glish and Communistks.....	8
What are Communistks?	8
What is Glish?.....	8
Communications Track 2	
The Glish Language and the Glish Executive.....	9
Glish Script	10
Interfaces Controlled by Glish Script	11
Tuning an Accelerator Beam	12
The Tune	12
Measuring the Tune.....	13
Correcting the Tune	13
Tuning the Accelerator Beam Illustrations	15
Altering Sequences.....	16
Using Communistks to Alter the Beam Tune	16
Using Glish to Alter the Beam Tuning Sequence	16
Communistk Elements.....	17
The Focus.....	17
Communistk Elements.....	17
Communistk Constraints.....	17
Change of Focus	18
The Self-describing Data Standard	19
SDS Track 2	
Data Transfer and Memory Allocation	21

The sid Editor.....	21
SDS Enables Transfer.....	23
Two "Navagation" Aids	23
The Graphic Library for the Integrated Scientific Tool Kit.....	24
The GLISTK Family Tree.....	24
Credits and References.....	26
ISTK Implementation	26
ISTK Information and Documentation.....	26

Introduction to the Integrated Scientific Tool Kit

The Integrated Scientific Tool Kit (ISTK) is a computer software tool and component kit developed for use by experimentalists. ISTK's development has originated in software design for particle accelerators but its uses are applicable to many experimental physics control and analysis systems.

ISTK's three main components are designed to manage major areas of software systems which are often not sufficiently served by commercial or industrial products. By first looking at the tasks involved to carry out a project, it is easy to determine which ISTK components are appropriate for the tasks at hand.

SDS (the Self-describing Data Standard) facilitates data flow: the movement of data between computers and between applications. If any data will be used in more than one process, it should be put in an SDS.

The Communication Systems (Glish and Communistks) manage the control flow of operational sequences. The kind of control flow performance needed, between elements of a single application or between applications, determines whether Communistks or Glish (the internal and external control flow components) should be used.

GLISTK (the Graphic Library for the Integrated Tool Kit) handles the interface between people and computers.

Some of the currently available ISTK tools include:

sid, an interface to numeric and textual data that provides customized access and editing mechanisms.

latview, a graphical view of accelerator lattices that is able to signal user queries to other programs.

conv, a program that converts binary data files to different machine architectures.

gasp, a pre-processor for basic SDS program writing.

The Glish executive, a program to carry-out complex and adaptive multi-processor sequencing.

The developers of ISTK are involved in the ongoing process of developing multi-use software tools that can be assembled and modified for the user's own purposes. ISTK is also a collaborative project. Tools and ideas are being developed at the Superconducting Super Collider Laboratory, Lawrence Berkeley Laboratory, FermiLab and other physics laboratories world-wide. People are encouraged to add their own multi-use ISTK tools to the existing ISTK libraries and utilities.

Document Format

This document exists on two "tracks". Those readers desiring a general overview of the Integrated Scientific Tool Kit, its approach and its components, should read Track One. Those desiring more in-depth, technical information should supplement Track One with Track Two.

The ISTK Approach Track 1

A Cost-Effective Approach

There is a problem inherent in designing software: design and costing procedures often become a large part of a project. One common dilemma is that, unless a software project in development is already well understood, there is considerable financial risk in the early stages yet full understanding of software needs often comes only once the software development has been done. This dilemma makes it difficult and expensive to make a commitment to full software implementation.

The ISTK components have been designed to aid in the preliminary stages of product or system delivery (beyond the basic proof-of-principle tests). Using the ISTK components, ideas can be explored, explained, made more precise, and discarded (if necessary) without great additional cost.

The ISTK system can quickly respond to changes because the realm of application in which it was born (technically challenging experimental control and analysis) is always in such a phase of development. In large physics experiments, software flexibility is required at every stage; by the time a system has been "hardened" into a final form, other factors (equipment, realm of applicability) have changed and the software must be altered to meet the new needs.

ISTK's strengths are its flexibility and its faculty in filling the gaps between available software programs. ISTK cannot replace all software but it can be used as an excellent supplement to existing commercial software. In a constantly fluctuating experimental situation, ISTK's flexibility provides a great advantage. In a market where factors are more stable, the ISTK use-approach and design emphasis on flexibility for the user can also give significant advantages during the appraisal and start-up phases.

ISTK provides a variety of tools for building various kinds of systems. The components can be used as prototyping tools to create applications in a short period of time, for experiment preparation (simulations, test experiments, etc.), controls systems, and for preliminary data analysis. Much of the ISTK system is sufficiently hardened to form the basis of production code decisions; decisions that can be made effectively only after a job analysis process.

Deciding Which Components to Use

The appropriate way to use ISTK is to break each overall job into basic, general tasks. In order to decide which tools are needed for a particular job, the tasks should be broken into generic bits which are as dissociated from the particular application as possible. Some of the possible tasks inherent in a jobs might be: organizing data, displaying graphics, listing datasets, etc. This conceptual analysis includes determining the needed processes, defining the flow of data between these processes, and defining how these processes should be connected by control flow. All of this analysis should occur *before* any implementation takes

place. The design of the major ISTK components; SDS, the data discipline and Glish, the control-flow manager, are, in a sense, a bottom-up response to top-down and object-oriented analysis techniques.

Implementation Decisions

The ISTK components can be used for developing large scale systems as well as for designing simple and complex single applications. The conceptual approach is basically the same; only the implementation differs. ISTK tools are not specific to any one application. Strong attempts have been made to insure that the tools are well-designed enough to be useful, yet abstract enough to be used in a variety of circumstances.

CAVEAT:

It is important to remember that the ISTK tools and techniques do not, in themselves, solve problems. Their role is to ease the ability of skilled people (programmers, engineers, physicists) to encapsulate and pass on their problem-solving experience to end-users, end-users who do not want or need the depth or breadth of knowledge contained in the software structures in order to perform their tasks. In this, ISTK is like all other software systems. To pretend otherwise is wrong; to expect more will lead to disappointment.

Approach Track 2

More About ISTK and the Unix-Style Approach

The ISTK package contains general all-purpose applications; components that can be used alone, together, or as a supplement to pre-existing software to aid in clarifying and organizing experimental data and operations. The ISTK system of using inter-connected components owes a large debt to the principles used to build Unix, that is, independent modules each designed to do a single job well, working together in conjunction with each other. The result of this approach for Unix was an immensely capable and coherent operating system.

ISTK's various programs are not connected until they are meshed together in the net in any way that the operator chooses. The ISTK applications can be thought of as Tinker Toys, components which remain the same, hooked together by changeable, flexible, connections. By merely altering the way the components are connected, the user can build elaborate programmatic structures without making internal program changes.

ISTK extends the concepts of the Unix system. Whereas Unix is a programming environment, ISTK extends it to components that provide great flexibility for people doing physics experiments. Conceptually, the SDS data discipline and the Glish control flow description parallel the general rules of the Unix operating system and allow the building of general purpose tools.

Communications Track 1

ISTK Communication Systems

Glish and Communistks

The ISTK Communication Systems, Glish and Communistks, are the ISTK components in charge of control flow within and between applications. Within ISTK applications, connections are managed by the Communistk, the ISTK communication class. Between applications, ISTK connections are managed by Glish. Conceptually, the Glish and Communistk connections are similar. However, their implementation varies greatly.

What are Communistks?

Communistks are the communication class for the ISTK libraries. A Communistk is an object which focuses on a value and has a list of other data objects to notify when this value, or the Focus itself, changes. Communistks exist only in the object-oriented languages. They enable and simplify communication between objects in a program and, by using the Glish connection, between different programs within a processor and between programs running on physically distinct processors.

What is Glish?

Glish is both a rule-based language to describe high-order sequencing and an executive to carry out this sequencing. The Glish language explicitly states the control flow so that it can be managed and properly understood. The Glish Executive is also a multi-task controller and inter-task communication system.

The Glish system separates the high-level control from the mechanical, pre-determined software and hardware operations and connects the control flow to an up-front sequencing language that coordinates interdependent sequences of widely distributed processes. Because of this separation, the Glish sequence of operations is protected from non-vital hardware failure.

Glish enables tasks to be broken into small components for clarity and organization. The components run as small programs orchestrated by Glish into a larger program with a common set of rules and conditions. Any program, even one completely unaware of Glish, can be run and monitored by Glish.

Glish combines the benefits of computer sequencing with the flexibility of on-line human modification. It allows quick, immediate input while freeing people from redundant or petty operating decisions.

Communications Track 2

The Glish Language and the Glish Executive

Glish establishes, in a few words, rules on how to respond to an infinite set of circumstances. To do this, Glish has two parts, the Glish language and the Glish executive. The Glish Language is used to express the sequencing and the Glish Executive carries out that sequencing. The user establishes the cause and effect relationships by writing the Glish script. The Glish Executive then orchestrates and sequences events between small programs and organizes them into a larger program with a common set of rules and conditions. The Glish Executive is much like a switch-board operator who establishes the connections between appropriate "parties", allowing the transfer of information.

(The relationship between the Glish language and the Glish Executive is paralleled by the relationship between SQL as a data definition standard and the DBMS's that implement all or part of that standard, and the relationship between FORTRAN as a language standard and compilers that implement all or part of that standard.)

Glish is not restricted to one Executive, and in different circumstances, different Executives may be appropriate. For example, in a particle accelerator, there is a range of control flow, from hard-wired power supply settings, through theoretical and mathematical abstractions, to subjective decisions (like determining the full validity of the results). At times, a combination of Executives might be installed on the same system; the high-level Executive interfacing with the human operators and the low-level real-time system Executive in charge of the hard-event driven functions, both using the same sequencing rules, expressed in the same language.

Communications Track 2

```
oswald := client("oswald -T -geometry +1+1");
sscsel := client("sscsel -geometry +660+1");
catch := client("Ccatch -geometry +190+380");
sid := client("sid");
whenever sscsel->SdsFile do
{
  oswald->SDSFILE($value);
  sid->Display($value);
}
whenever oswald->MoveTo do
{
  catch->ElementIndicator($value);
}
whenever sscsel->SRange do
{
  oswald->SRange($value);
}
```

Glish Script

The Glish language uses some unusual constructs ("whenever", "if and whenever" etc.) to handle sequencing situations. A simple description is written to describe what is expected to occur, the events that should be generated by those occurrences and under what conditions those events should be channeled to other processes. Because Glish is rule-based, this description of conditions can be written in any order. When the conditions defined in the Glish script are met, Glish sends named events to specified processes ("clients"). Glish can start and stop its clients and can provide inter-process communication between various clients. Any two (or more) processes running anywhere on the network can use Glish to communicate.

Communications Track 2

The image displays four distinct windows from a software application:

- Top Left:** A table listing parameters for a device. The data is as follows:

name	slot_length	out_diameter	height	beam_ssp	aperture	weight
bs	13.299	0.856	0.706	0.000	50	12000
- Top Right:** A graphical selector showing a circular diagram with a vertical line and a horizontal line. Below it is a list of regions: SouthWest Arc, Instructors, NorthWest Arc, SouthEast Arc, Collators, NorthEast Arc, North Abort, South Abort, and Injection Straight. A 'Selection complete' button is visible.
- Bottom Left:** A plot showing a series of rectangular pulses on a horizontal axis. Below the plot is a table:

Lat Index	11651	11651	Name	q5b	Level	0
S	39473.410	39473.410	Type	drift	Length	86.824
- Bottom Right:** A detailed data window for a lattice structure. The data is as follows:

lattice	24932	1302.075	299
element_index	96		
type_index	1		
occurrence	18		
level	0		
hook_index	0		
sense	0		
element	570	96	
name	ch		
type	drift		
sub_type			
length	0.325		
strength	0.000		
ift	0.000		
type_index	1		
npars	0		
length_index	0		
strength_index	0		
ift_index	0		
more_index	0		
parameter	439	0	
name	Nothing		
value	0.000		
legat_type	28	1	
name	drift		
has_length	1		
has_strength	0		
has_ift	0		
has_more	0		
row	9	0	
name	collimator		
cols	2		
column	2	0	
has_size	xsizc		

A screen dump of the interfaces controlled by the Glish script of (). Accelerator regions are chosen by the graphical selector (center top) and displayed graphically (bottom left) and textually (right). Selection of a magnetic devicee from these latter two displays feeds an SQL query to the main database. The results of this query, giving more detail on the device, are shown at top left. Each display is from a seperate program. They are controlled by Glish events.

Internal communication also exists within each display. In the display at right, for instance, the lattice structure 299 (at the top of display) is connected to the appropriate element type 96 (just below) by Communistsks. These internal connections may also be sent out of each program via Glish.

Communications Track 2

Tuning an Accelerator Beam

The example of tuning a particle accelerator beam demonstrates how the ISTK sequencing components, (Glish and Communistks) allow human intervention in the computer sequences that run the machine without interfering with the overall sequence. The example below shows how processes can be internally "re-wired" using Communistks and externally "re-wired" between each other using Glish commands.

The Tune

Setting the "tune" of a particle accelerator is crucial to its performance. When a beam of particles spins around in a circular accelerator, magnetic fields try to focus it into the middle, forcing that beam to oscillate. This beam of particles oscillates at a particular frequency (number of times around the ring), based upon the strength of the focussing magnetic field.

The exact number of oscillations the beam performs per cycle is vital, as various things can conspire to make the oscillations bigger and bigger, eventually destroying the beam. An accelerator is never, for instance, set to an integer tune. A beam set to an integer tune will oscillate exactly that integer number of times and cycle back in exactly the same place in exactly the same direction. If there are any imperfections in the magnetic field (a situation impossible to avoid), the beam set to an integer tune will be kicked slightly sideways by these imperfections every time it cycles the accelerator, and very soon the beam will be kicked completely out of place. (Imagine the actions of a child on a swing pushed by somebody else. If the pusher pushes at the same place in every oscillation, even if the pushes don't have much force, the oscillations will get bigger and bigger and bigger, until the child falls off. If, however, the child sitting on the swing is hit randomly, she and the swing will stay at the bottom of the curve, oscillating only slightly.) It is therefore vital to set the beam to a non-integer tune so that the first time the beam cycles the ring it will return to the center of the ring, the second time around it will have gone many oscillations plus a little bit, etc. This argument is true for oscillations that repeat after 2, 3, 4, ... times and can get quite complex in interactions with other forces. The principle is the same in all cases. The tune must therefore be able to be measured and adjusted slightly as the accelerator runs.

Measuring the Tune

A beam pick-up is located at one point in the accelerator ring and measures the beam position on successive turns. (In this example, the important part of the measurement is the post-decimal number. The integer part of the number of oscillations is not important for this example.) From these measurements a graph is obtained. The graph will show a wave pattern that may damp down as the oscillations get smaller and smaller or it may increase in magnitude, indicating unstable operation.

A Fourier transform is performed (usually using a mathematical technique called a Fast Fourier transform or FFT) to translate the wave form into a graph of a musical tone, showing pitch, movement and timbre. If a beam movement is analogous to the guitar string motion when the guitar string is plucked, the FFT transform (no information is added or deleted) is the mathematical representation of what the human ear apprehends from the guitar string motion; not the change in position of the string, but a perception of the range of speeds of movement.

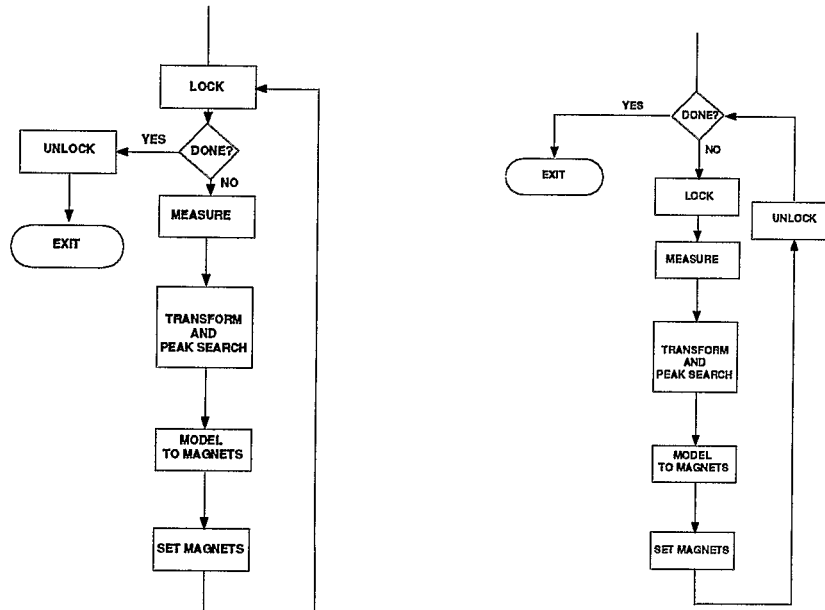
The beam should not measure as a pure tone, as a too-exact tone is at risk for becoming unstable. The standard beam will plot with some noise at low frequencies, a big peak (not a delta function, but a band of peaks), and more noise at the high test frequencies.

Doing a peak search to obtain the top number of the peak (fitting a curve to the highest points), the number of the tune is finally obtained. After the tune number is obtained, changes can be made to it, the power supplies can be altered accordingly, and the tune remeasured.

Correcting the Tune


As the particles are accelerated toward top energy (the speed at which the experiments will work), the magnetic bending and focussing must change in concert to control the progressively "stiffening" beam. The tune of the accelerator may change as the beam is ramped up from low injection speed to high field where the experiments can take place. As the beam accelerates, the tune is measured every sixty milliseconds and these measurements are plotted time against tune. The graph of the tune should show smooth, well-understood and expected movements. At times, small glitches around the curves occur and this is one place where the power supplies often need to alter the tune to flatten the glitches. Once the particle beam has been accelerated up to speed, other small adjustments may need to occur as well.

Because these tune measurements need to be taken so often to get an accurate picture of the beam, it is appropriate that computer programs (rather than people) be designed and implemented to do these measurements automatically with human intervention only when necessary. Adjustments to the beam also occur at rapid speeds and it is appropriate that computer programs perform these other functions as well.




This is a highly simplified block diagram illustrating tune measurement and feedback. One problem: as the machine is locked while an extensive group of measurements is done, operators may be unable to work for a considerable time...

... a change at the Glish script level will allow other users to touch the machine between measurements...




...to break control within a program (here to inspect data halfway through analysis), one must inform internal communication links to reattach to outside processes. This is also possible dynamically and without recompilation.

Tuning the Accelerator Beam




Tevatron beam oscillating back to central position after a sideways kick...



...the graphical browser "kasper" has been asked to plot only every thirteenth point. The plot demonstrates that the beam is on a thirteenth order resonance...

transform of data before kick

transform of data after kick



...a Fourier transform shows a strong peak at a fractional tune of 4/13.

Altering Sequences

Both ISTK sequencing components, Glish and Communistks, have the flexibility of allowing sequences to begin operation with a small number of rule-based functions and having other control flow statements injected later.

They also allow the functions to be disassociated from each other so that if one part of the sequence goes wrong, the entire sequence will not crash. They allow human intervention in the computer sequences that run the machine without interfering with the overall sequence.

Using Communistks to Alter the Beam Tune

In the instance of the particle accelerator's tune, a program can be written to take beam measurements and transform them into a graph. Sometimes though, for instance, the original beam measurement is inaccurate and the tune/time analysis will show a wayward analysis. In this case, the FFT might show "fuzz" and this noise will lead to a false peak during the peak searching analysis. If this occurs, the accelerator operators need access to the program during the time between the FFT and the peak search, a time where there is a single control flow. Communistks enable a person to plug in a viewer and control panel at this point. Only when the operator approves of the FFT will the peak search continue.

Using Glish to Alter the Beam Tuning Sequence

It is necessary to prevent other input during the period of time between measurement of the beam tune and alteration of the power supplies that control the beam tune. For this reason, the machine must be locked during these procedures. The long length of time taken by locking the machine (thus monopolizing it), performing the measurements, altering the power supply, remeasuring, re-altering the power supply, and completely correcting the beam before unlocking the machine can become a problem when other people and functions need access to the machine. Glish enables the "wiring" to be changed with a statement of rule change, moving the lock/unlock functions into the loop cycle so that there is time between each alteration cycle for others to "play through".

Communications Track 2

Communistk Elements: The Focus and The Constrainer

The Focus

The Focus of a Communistk is essentially a "Thing of Interest" to a program object; in computer terms, a memory location into which a value is written. The Communistk is responsible for policing what happens to that Focus and for routing advice of any changes to the Focus' value to data objects which have expressed an interest in the Focus. The Constrainer class holds constraint information about the value of the focus.

When there has been a change to the Focus, the Communistk informs data objects which have an interest in the Focus about the changes. These data objects don't know or care which or how many data objects may change the Focus. They are only interested in knowing if it has been changed. The data object which changes the value of the Focus does not notify any other data objects of its action; it is the data objects who wish notification which ask the Communistk to inform them of any changes to the Focus.

Communistk Elements

The Present Datum device is thus notified by the Communistk of a change in the value of the Focus. It does not change the data, it only presents the data.

When the change datum device, the "catcher", is given a change request from a person, it requests information from a separate data source that holds information about the name and type of the focus.

The change request is then filtered through the Communistk and the Constrainer. Between the Communistk and the Constrainer is a forced internal link.

Communistk Constraints

Any changes to the focus are filtered through the Constrainer which holds pre-established allowable values; minimums, maximums, etc. The information about these constraints is held in a separate data source. Responsibility for changes to the focus lies with the Constrainer, the low-level wrapper. The change-datum devices do not need to take on that responsibility.

The Communistk creator, when creating space for the Communistk's Focus or when attaching to the Focus, can assure that attempts to write in a value of an incorrect value type will be blocked. Additional constraints on that value (hard minimum and maximum values, warning and alarm values, restrictions on single-hit changes, etc.), may be placed using the Communistk Constrainer class. These constraints are associated with the Focus itself, not with the data object that is inputting the value change. Any attempt to change the value of the Focus are properly checked and illegal actions will be thrown out by the Communistk. Such actions may be monitored without intervention of the object which tried the modification.

Change of Focus

In addition to the change in a Focus' value, the Focus itself may move from one data object to another. An analogy for this might be the "daily specials" list at the local restaurant:

Joanne's Restaurant has a list of dishes served only on Mondays (roast beef with potatoes, vegetarian casserole, shrimp salad), a list of dishes served only on Tuesdays (Chef's Salad, Sloppy Joe with boiled carrots, calves liver saute), etc. These lists remain the same; every Monday the specials are the same as every other Monday. Joanne is lazy, she doesn't like to rewrite the Specials Board every day, so she has seven boards in the closet, each one listing the specials for one day of the week.

Joanne's Special's Boards are analogous to the Foci. Foci value (like the daily lists) remain the same, but a move in the Focus resets which value is being looked at (Monday's, Tuesday's or Wednesday's Board). There are a number of methods which do no changing of Focus value but simply declare to all interested parties that the data Focus has shifted.

The Self-describing Data Standard

SDS, the Self-describing Data Standard, is the ISTK component that facilitates data flow: the movement of data between computers and between applications.

In fundamental terms, SDS is a description of the structure of a selected quantity of data. SDS binds a descriptive information header to this data; this layer of information helps facilitate data manipulation and data transfer between processors of diverse architecture and between various applications within a computer.

By using the SDS description, every application in the system, whatever its "world view", has the ability to utilize the same data. (SDS does not control or determine data formats; the person generating the data determines the appropriate format.) For applications designed to analyze general forms of data in general ways (displays, Fourier transforms, etc.), SDS aids by providing a full description of the data to be analyzed, and for applications designed to analyze specific types of data in specific ways, SDS aids by filtering out inappropriate data.

SDS' strength (like that of all the ISTK components) is in its ability to fill the gaps between other available software. SDS has been created to provide a common link between data management techniques in an experimental situation where it is effectively impossible to replace the myriad of processors and languages with a single system.

Transferring data between applications and between processors can present problems, especially for people in the experimental world who need a great deal of flexibility. Many of these problems are based on differences in data structuring and memory allocation in different systems.

SDS facilitates the data transfer between processors of different architecture. Perhaps a good way to think of SDS is as "glue", a substance that binds these components (systems, languages, applications) together at the level of bits and bytes, the level of memory allocation and data structuring.

SDS operates on a higher level of data description. One of the important SDS data description features is the SDSListk, a type of persistent list that requires no duplication of data items and no manual memory allocation.

SDS also aids data transfer between programs within a computer, using names and internalized pointers to sort many of the data coherency problems of shared memory that exist in compiled languages.

There are tools to aid in constructing an SDS description. Once the SDS exists, data can be examined, converted as binary data between various processor architectures, or moved to various places where it can live on file, in shared memory, or in network messages. Programs currently exist to convert the data to another format, edit the data, copy the data, delete the data, find out what architecture it was originally formatted for, etc.

Using SDS to describe data also allows for potential upgrade. Should a system become available that is more appropriate than SDS, SDS will allow for transition to that system. Without the use of SDS (or another data-standard), the transfer of information into a new environment will be difficult.

SDS Track 2

SDS - Data Transfer and Memory Allocation

SDS assumes that, optimally, an entire operation will be automatically run by computers and, in designing SDS, a decision was made to build a structure optimized for the systems with the highest rates of performance rather than building a system optimized for the slowest systems (people). In SDS, data remains in two parts: the basic data objects, and

--- Edit SDS source qcd100.data.sdsh ---

Base 10 Search >>> Mode Datasets

qcd100.data.sdsh from Wed Jun 10 18:00:11 1992

SDCEvent[1661]		< 203 >	1	P a g e s	
bytecount	59824				
record[4985]		< 49 >			
first	1.144e-02				
second	83.000				
third	2.000				
checksum	981024450				
checkcount	59824		0		
SDCEventAddress[1661]	-217773852	< 203 >			L i n e s
SDCEventCounts[1661]	4985	< 203 >			
SDCEventDesc					
filename[19]	/homee/qcd100.data				
mod_date	695067484				
offset	0				
size	85770240				

Save to file:

InstaQuit

Here, the sid editor is attached to a large (> 80 Megabyte) experimental data file. The header contains all type and size information, including the three lower objects SDCEventAddress, SDCEventCounts, and SDSEventDesc. These objects contain file information, event pointers, and sizes extracted from the body of the data by a previous scan.

The data itself is mapped from disks. Random access to each event is efficient. Editor customizing is done with a resource file. In this example, the three top objects are ganged together so that accessing SDCEvent [203] will automatically access the lower objects (also [203]). The number of records within each event has been read from the datafile and written explicitly in the SDCEventCounts object. In general, each event is of a different length.

the data structure description. This two part division assures that if data searching is needed it can be done as efficiently as by any interpreted language, and at times if and when data searching is not needed, data access is optimized for use by compiled languages. Because SDS stores data in a format appropriate for high-speed manipulation, the system would be useless without interface tools for the interaction between people and data. To this end, editors and viewers have been developed and are provided with the SDS package.

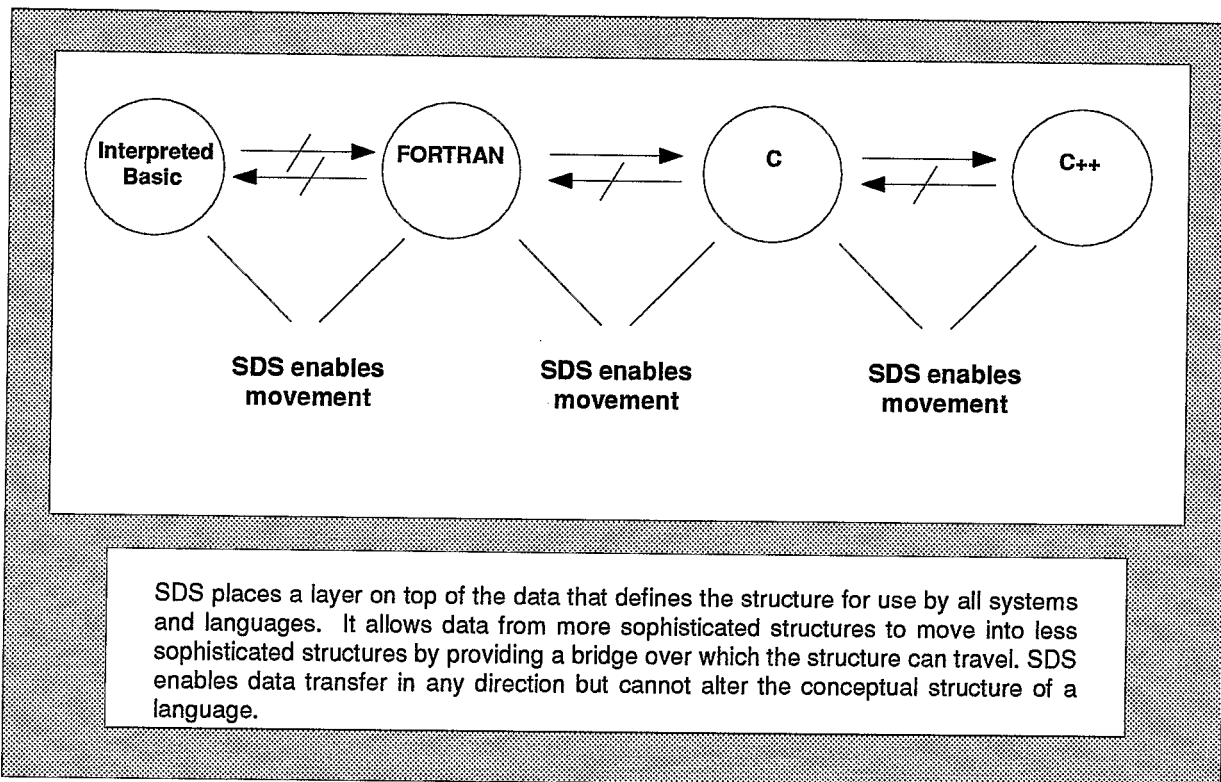
Sharing information between processors and between computer languages is a general problem in the computer world. Different computer systems interpreted languages embed data structuring information in their data, information that cannot easily be read by compiled languages. Some data structures are more sophisticated and complex than others. One of the major problems is that every computer stores its data structures differently.

In SDS, the meta-data that exists in a compiled language only during the compilation period remains as information directly associated with the data. Unlike in an interpreted language, however, no descriptive information is embedded in the actual data. And unlike in a database management system (which must also contain this meta-data), the full description may be moved with small data sets.

SDS formalizes memory allocation so that the programmer doesn't have to remember how much memory has been allocated. In order to keep track of memory allocated, part of the structure of the SDS header describes the size information of all of its data objects. The SDS description also contains the original memory allocation requests; it enables the programmer to wait until run-time to allocate memory, and the SDS description can be written either at compile time or at run-time.

SDS attempts to use the best facets of the memory allocation techniques of interpreted and compiled languages in the right order (keeping the speed of C and the informative power of the interpreted language, but not in a way that slows down C) and then attempts to make it possible to move these pieces of allocated memory from process to process and across networks.

SDS puts data onto the network in the format it has been programmed in and provides enough information about the structures, sizing, memory allocation, etc. so that any computer can translate the data into a readable format. By presenting the data in its own format, SDS avoids the problem of unnecessary translation.



The left screenshot shows a graphical/textual selector for SSC accelerators. It features a list of options: SouthWest Arc, Injectors, NorthWest Arc, SouthEast Arc (highlighted), Collisions, NorthEast Arc, North Abort, South Abort, Injection Straight, and Selection complete. Below the list are buttons for "To top" and "InstaQuit".

The right screenshot shows a dataset selector. It displays a list of datasets: "exsds.sds" and "exsds.sds:1?". Below the list is a text field containing the path "/usr/local/src/ISTK/Data/HP68K/". At the bottom, there are buttons for "Read disk", "Read memory", and "InstaQuit".

Two "Navigation" Aids

The illustration on the left shows a graphical/textual selector for SSC accelerators. The illustration above is a dataset selector which displays file directories and SDS datasets.

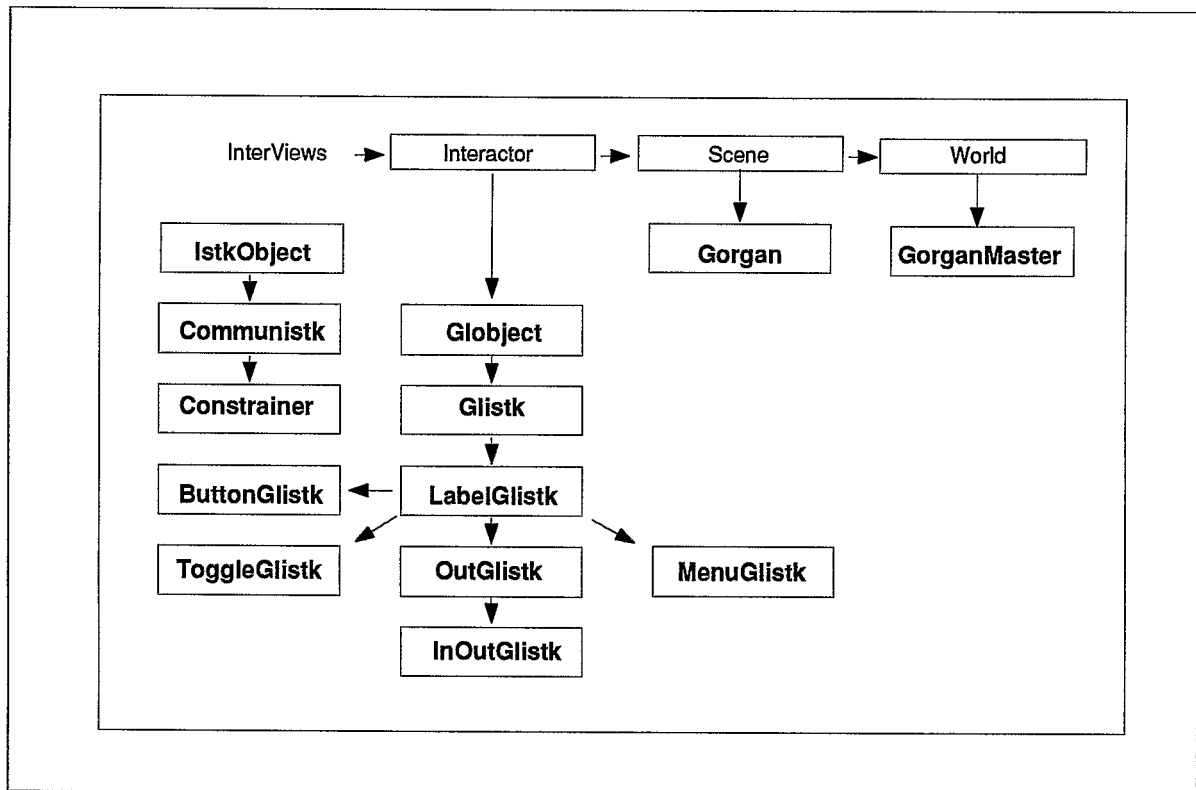
GLISTK Track 1

The Graphic Library for the Integrated Scientific Tool Kit

GLISTK is a C++ library of graphic tools developed to enable rapid creation of scientific computer interfaces for unobtrusive data-viewing and data-manipulation. Because GLISTK is designed specifically for experimentalists, it shrinks the time between the conception and execution of working, connected scientific interfaces.

A good interface is a communication feedback loop between a person instructing the computer what to do and the computer telling the person what it has done. GLISTK provides subtle but important features to strengthen and clarify the feedback connections between the person and the computer through its "no-frills" aesthetic display, through highlighting, through mouse button simplification, and through the lack of icons labeling GLISTK objects.

The Glistk Family Tree



A primary difference between GLISTK and other interface-building tool kits lies in GLISTK's ability to communicate between graphic objects and its ability to extend that communication into the outside world. By having a non-graphic object as its base, a GLISTK interface can be operated by other processes as well as by people. Communication is facilitated by the ISTK communication class, the Communistk. Communistks make the interior connections between GLISTK objects as well as provide linkage to the outside world via the Glish executive.

All of GLISTK's features have been designed with "data purity" in mind. GLISTK objects provide a view of a Communist's data focus and provide the means for altering data, but they do not affect the data by its display.

GLISTK objects are basic building blocks from which other, more intricate and specific interface tools can be easily derived. While decisions have been made about the look and function of every object in the tool kit, almost every object method is "virtual" so that those functions and appearances can be changed.

The Gloss library is an extension of GLISTK. Gloss objects are derived from GLISTK general objects to form more complex objects such as: Plistk, a graphic plotter; DisLis, which displays the ISTK list class; and Glider, which combines decrement and increment ButtonGlistks with an InOutGlistk. Gloss is a collaborative project and people are encouraged to add their own derived objects to the library.

Credits and References

The Integrated Scientific Tool Kit (ISTK) was developed directly out of the problems associated with running particle accelerators. Development began in the early 1980's by Christopher Saltmarsh at the CERN Super Proton Synchrotron in Geneva. SDS, developed by Saltmarsh, was used at the E778 FermiLab in which non-linearities were added to the Tevetron magnetic field to simulate potential S.S.C. control problems. Other ISTK components used at this experiment were Glish written by Vern Paxson at the Superconducting Super Collider Central Design Group, and extended to its present version at LBL, and GLISTK, the ISTK graphic library, developed primarily by Matthew Kane. SDS (and the other ISTK components) are under continuing development and are currently being used at various sites internationally. Continuing ISTK development includes work by Matthew Fryer, Andrea Young, documentation by Ericka Lutz, and contract work from Soam Acharya of Cornell University and Sriram Chittathoor of the University of Connecticut. There have been considerable contributions from user groups such as the S.S.C. Magnet Test Division. The current list of ISTK collaborators (see the mail exploder on largo.lbl.gov) includes people from: the Superconducting Super Collider Accelerator System and Magnet Divisions; Lawrence Berkeley Laboratory; FermiLab; Cornell University; Lawrence Livermore Laboratory; Texas University; and the Texas Accelerator Center.

ISTK Implementation

- The base computational environment for ISTK is Unix.(*), although the SDS component is ported to VMS and PC systems.
- The current SDS implementation is in C with a C++ wrapper and a Fortran binding.
- Glish functions are available from C++ and C.
- GLISTK is presently based on the InterViews graphics package.

ISTK Information and Documentation

All code and documentation is available via anonymous ftp from [largo.lbl.gov](ftp://largo.lbl.gov). The main postscript manuals are:

Lutz, Ericka, and Matthew Kane. The GLISTK Manual . March, 1991

Lutz, Ericka. The Integrated Scientific Tool Kit: An Overview of the ISTK System and its Components. Superconducting Super Collider publication S.S.C./man 012. July, 1991

Lutz, Ericka. The SDS Document, A Conceptual Basis Towards Understanding the Self-describing Data Standard. December, 1991

Paxson, Vern. The Reference Manual for the Glish Sequencing Language. (vern@ee.lbl.gov)

Paxson, Vern and Chris Saltmarsh. Glish: A User-Level Software Bus for Loosley-Coupled Distributed Systems. Extended Abstract July 21, 1992. Submitted to the Winter 1992 USENIX Conference.

Man pages, notes and technical documents are also there. Enquiries may be addressed to
salty@largo.lbl.gov
matt@largo.lbl.gov

* UNIX is a registered trademark of UNIX System Laboratories, Inc.

The SDS Document
A Conceptual Basis Towards Understanding the
Self-Describing Data Standard

Ericka Lutz and Chris Saltmarsh
December 1, 1991

TABLE OF CONTENTS

THIS IS SDS

THIS IS A TEACHING DOCUMENT

WHAT SDS DOES

EXPERIMENTAL SYSTEMS AND DATA FLOW

Dealing with Data Flow Management

Integrated Systems - Ideal If They Could Be Universal

Interpreted and Compiled Languages

The SDS Approach

WHY SDS?

Data Transfer / Memory Allocation / Data Structure

MEMORY ALLOCATION IN COMPILED LANGUAGES - C

Automatic Allocation

Run-Time Allocation

Common Memory Allocation Errors in C

A Note About Memory Allocation in FORTRAN

Lost Information and Overwrites in Compiled Languages

MEMORY ALLOCATION IN INTERPRETED LANGUAGES

Demonstration of a Crude Search Mechanism

SDS' MEMORY ALLOCATION TECHNIQUES/SYSTEMS

DATA TRANSFER BETWEEN PROCESSORS - LIFE BEFORE SDS

Data Structures

Varying Data Structure Storages in Compiled Languages

Dealing with Data Transfer - External Data Representation

THE SDS APPROACH TO DATA TRANSFER

TRANSFER WITHIN A SYSTEM

ALLOCATION WITHIN PROCESS AND SHARED MEMORY

Sharing Data Originating in Process Memory

Shared Memory

How Shared Memory Functions

SDS AND SHARED MEMORY

DATA DESCRIPTION

LISTS AND ARRAYS

LISTS IN THE C++ IMPLEMENTATION - THE SDSLISTK

Persistent SDSListks

HIGHER ORDER NAMING CONVENTIONS

SDS COMMUNICATION - COMMUNISTKS AND GLISH

The Focus of a Communistk

Communistk Constraints

Change of Focus

GLISH

ISTK Information and Documentation

Suggested Reading and References

THIS IS SDS

SDS, the Self-Describing Data Standard, is a component of the Integrated Scientific Tool Kit (ISTK), a computer software tool and component kit developed for use in large experimental control systems. The SDS component of ISTK facilitates data flow: the movement of data between computers and between applications.

SDS' strength (like that of all the ISTK components) is in its ability to fill the gaps between other available software. SDS has been created to provide a common link between data management techniques in an experimental situation where it is effectively impossible to replace the myriad of processors and languages with a single system.

THIS IS A TEACHING DOCUMENT

To understand why SDS has been developed and how it operates so that it can be used effectively, it is necessary to understand the overall concepts of data management on a deeper level. This document has been prepared for that purpose. This text has not been designed to be used in a computer science course; it is a simplified and generalized explanation that provides the conceptual framework within which SDS can be understood.

SDS is a system that can be used at various levels for progressively more complex tasks. For simple data description and transfer, SDS can be used without much explanation. However, for more complex processes, a stronger background is useful. To understand what SDS can and cannot do, any users or potential users should read this document before using SDS. (For very experienced programmers, much of the material in this document will be review and can be skipped at your discretion).

In this document, we discuss how SDS has been developed in response to the experimenter's problems and to problems in existing techniques, and we review some of the fundamentals of low-level computer function; memory allocation and data structuring in complete computer environments, within the compiled languages of FORTRAN, C, and C++, and within the interpreted language of Basic.

We also review higher-order naming functions and memory allocation in process and shared memory, SDS general organizational structure objects (SDSListks), persistent objects, and object-oriented programming (with a discussion of two ISTK components: Communistks and GLISH).

WHAT SDS DOES

In fundamental terms, SDS is a description of the structure (byte ordering, name, floating point coding, etc.) of a selected quantity of data (a dataset). SDS binds a descriptive information header to data; this layer of information about the data helps facilitate data manipulation and data transfer between processors of diverse architecture and between various applications within a computer. SDS does not control or determine data formats; the person generating the data determines the appropriate format.

By using the SDS description, every application in the system, whatever its "world view"¹, has the ability to utilize the same data. For applications designed to analyze general forms of data in general ways (displays, Fourier transforms, etc.), SDS aids by providing a full description of the data to be analyzed, and for applications designed to analyze specific types of data in specific ways, SDS aids by filtering out inappropriate data.

¹Every computer language has its own "world view" towards viewing data and there are often dramatic differences between the world views of different programming languages. This "world view" is often one of the reasons to choose to work with a particular language.

EXPERIMENTAL SYSTEMS AND DATA FLOW

There have been a number of commercial (often business-oriented) systems designed to address the problems that SDS addresses: how to manage the flow of data within and between computers. These commercial systems are not always appropriate for experimenters working with large experimental control systems; their emphasis is geared to the commercial world which has different data management requirements. Many commercial systems too closely couple the data structuring problems with the applications.

During the operations of complex, experimental equipment, vast quantities of raw data are collected by numerous widely-distributed computers and instruments. These megabytes of incoming data must be clearly understood as well as be accessible for short and long term analysis. Keeping track of this data is one of the experimental programmer's biggest concerns, especially because it is difficult to predict exact technical specifications or to predict what problems will arise that will require data restructuring.

Because it is difficult to ascertain which data are pertinent before some analysis has occurred, much of the incoming data must be stored until the preliminary processing is complete.² It cannot be pre-determined who will need access to particular pieces of data or what type of computer and applications they will use for analysis.

Whatever is in charge of managing these quantities of incoming data must clearly understand what information has been gathered, must identify its origins, and must be able to transfer that information from processor to processor for analysis.

Dealing with Data Flow Management

There are several existing ways to approach the problems of data flow management. In an experimental situation, people have generally used existing interpreted or compiled computer languages, changing existing tools to fit the job rather than custom creating tools for the situation. This isn't always effective; it is always better to have the "right tool for the job" than to make do with tools designed for other purposes. More recently, complete environment and object-oriented languages have been used. These integrated systems, while very effective in certain circumstances, are also not necessarily appropriate for the experimenter.

Integrated Systems - Ideal If They Could Be Universal

For keeping control of data moving from process to process and across processors, a single system seems desirable.³

"The advantage is constancy, which shows up throughout the system. The whole integrated system is organized by a single principal, so a user need understand only this principal to explore any part of Smalltalk-80."⁴

In complete computer environments, the communication system works well, programming is easy, the debugging tools are effective, task-appropriate tools can be built and a graphic interface is included.⁵

²It is sometimes not possible to store all the incoming data.

³One example of a good single system is Smalltalk-80, which provides both a complete environment and an object-oriented language.

⁴Cox, Brad J. and Andrew Novobiliski. *Object-Oriented Programming : an evolutionary approach*, 2nd Ed., Reading, Mass.: Addison-Wesley, 1991, pg. 33

⁵Integrated systems typically don't easily lend themselves to sequencing but even that can be done.

If using a single system to facilitate data flow were feasible, SDS (and all the other Integrated Scientific Tool Kit components) would not be needed, as an integrated system provides a more coherent interface (and easier memory allocation, data structuring, data transfer) than the general purpose interpreted and compiled languages can provide.

However, integrated systems have their drawbacks. They are not yet efficient, perhaps too user-friendly (people can use them without a strong conceptual background) and costly. The reality is that people will continue to use other languages in various manufacturer's computers, all of which have slightly different ways of doing things, and will continue to need to share data generated in various languages and systems with each other. While the number of languages and processors used can be limited, it is not realistic to imagine that everything can be standardized.⁶

Unfortunately, using a single integrated system is not an option at this time. Given that processors of different types and languages of different types will continue to be used, the Smalltalk-80 type of complete environment can be effectively eliminated for overall control; it can only be used as a component.

Interpreted and Compiled Languages

Broadly speaking, there are two kinds of computer languages: interpreted languages and compiled languages. We will be concentrating on BASIC (interpreted) and C (compiled).

(NOTE: For the purposes of this demonstration we use an extreme approach. In reality, the differences between various compiled and interpreted languages are less stark; as with most things, they exist in shades of gray.)

⁶It is good to remember the ADA story. When the Department of Defense determined that 420 programming languages were being used in their facilities, they decided to develop ADA as a universal system that could handle every possible application. The result: the Department of Defense has 421 languages to contend with.

Interpreted and compiled languages translate into machine code in different ways:

Interpreters	Compilers
<i>Time 1</i> Read a line Check for errors Convert to machine language Execute line	<i>Time 1</i> Read entire file Convert to machine language Optimize Try to execute compiled program Check for errors
<i>Time 2</i> Read a line Check for errors Convert to machine language Execute line	<i>Time 2</i> Execute compiled program

Differences Between Compilers and Interpreters: Interpreters convert each program line to machine language each time the program runs. Compilers convert to machine language only once and then on subsequent times run the already converted commands.⁷

In a world of infinitely fast computers, interpreted languages alone might be sufficient for experimenters and SDS (or compiled languages and integrated computer environments) might not be necessary. An interpreted language provides quick feedback (informing the programmer of errors): all the data in its memory are named, easy to locate and easy to edit; everything can be saved; and prototype building and creation of other small programs is quick and effective.

At this point in time though, interpreted languages tend not to be fast enough for the experimenter's needs. In an interpreted language, data are searched sequentially and systematically every time the program is run. There are ways to speed up an interpreted language's performance and give it more efficient search methods⁸ but, whether an interpreted language is used in its most basic form or uses accelerated search methods, information about the data is embedded into the data themselves. This embedded information becomes a problem when the data are transferred between processors and from interpreted languages to compiled languages.⁹ To transfer data created by an interpreted language into a compiled language program requires moving the data through the separate processes of reformatting it, changing it, and filtering it. Typically, the data must be copied. These extra steps all slow the data transfer process.

Compiled languages are generally faster than interpreted languages because the compilation process has a better opportunity (the optimization process) to promote efficiency. This same advantage can also be a disadvantage: in order to make any alterations, however small, in a compiled language program, elements of the program must be re-compiled and linked. Another serious drawback with compiled languages is that information about the data (meta-data) is lost after compilation and it becomes difficult to reference what is contained in a block of data without searching through the code. Compiled languages, however, provide more data-structuring flexibility than interpreted languages do and, in an experimental situation where data-structuring flexibility and speed are vital, an interpreted language is at a disadvantage to the more sophisticated compiled languages.

⁷Adapted with changes from: Wilson, Stephen. *Using Computers to Create Art*. Englewood Cliffs, N.J.: Prentice-Hall, 1986 pg. 161

⁸Programs can be speeded up by a number of methods but interpreted languages are intrinsically slow.

⁹Compiled languages make assumptions that there is no embedded description of the data in the code.

The SDS Approach

SDS assumes that, optimally, an entire operation will be automatically run by computers and, in designing SDS, a decision was made to build a structure optimized for the systems with the highest rates of performance rather than building a system optimized for the slowest systems (people). In SDS, data remains in two parts: the basic data objects, and the data structure description. This two part division assures that if data searching is needed it can be done as efficiently as by any interpreted language, and at times if and when data searching is not needed, data access is optimized for use by compiled languages. Because SDS stores data in a format appropriate for high-speed manipulation, the system would be useless without interface tools for the interaction between people and data. To this end, editors and viewers have been developed and are provided with the SDS package.

WHY SDS?

SDS has not been designed to replace integrated database management systems. SDS combines data management techniques of interpreted and compiled languages and provides management control when data are being manipulated outside the database management system. The comparison between compiled languages and interpreted languages was not meant to argue the benefit of using of one system over another but, rather, to point out that neither system is perfect.

Data Transfer / Memory Allocation / Data Structure

Transferring data between applications and between processors can present problems, especially for people in the experimental world who need a great deal of flexibility. Many of these problems are based on differences in data structuring and memory allocation in different systems. Because computers store their data in varying formats, data transferred between systems often can't be read without a time-consuming translation process. Often, the "world views" (the idiom that a language uses to describe data) are not compatible and the languages have different approaches to looking at data (making problematic, for instance, the transfer of complex data structures generated in a compiled language to a system using an interpreted language).

SDS facilitates the data transfer between processors of different architecture. Perhaps a good way to think of SDS is as "glue", a substance that binds these components (systems, languages, applications) together at the level of bits and bytes, the level of memory allocation and data structuring. To understand how SDS binds these components, we need to look at the differences in data transfer and data structuring techniques in different systems/languages. We also need to look at memory allocation techniques in compiled and interpreted languages to clarify how SDS operates.

MEMORY ALLOCATION IN COMPILED LANGUAGES - C

In an experimental situation where it is impossible to determine how much memory is needed until objects have been created, it is difficult to pre-allocate memory. In the higher level compiled languages, memory allocation is formalized to allow the existence of dynamic objects of complex different data structures whose memory allocation is determined at run-time.

The following example is a demonstration of both automatic and run-time allocation in C language.

Sample Fragment of C code¹⁰

```
int a;  
int *b;  
a = 3  
b=(int *)malloc (size of (int));  
*b = 4  
printf ("%d\n",a);
```

Automatic Allocation

In the sample code above, when the compiler reads the command, **int a;**, the compiler allocates four bytes of memory in the stack¹¹ as the integer **a**.¹² The compiler keeps a pointer that instructs the computer: should there be a reference in the program to **a**, point to this area of memory.

The command **a=3** writes the number **3** into the area of memory (the four bytes) that has been allocated in the stack.

The command **printf("%d\n",a);** does, indeed, refer back to **a**. It tells the computer to print out the integer **a**. When the computer receives this command, the computer then goes to the determined place in memory (those four bytes in the stack), finds the value of **a** (which we have already determined to be **3**), and prints it.

Run-Time Allocation

Something different happens when the compiler reads the command **int*b;**. Because ***** means address, the compiler understands that **b** is not an integer itself, it is a pointer to an integer. Since a pointer takes up four bytes of space just as an integer does, the compiler allocates four bytes of space in the stack. This pointer now points at some undetermined place in memory.

The command **b=(int *)malloc (size of (int));** allocates memory for **int b** at run-time (as opposed to automatic allocation at compile time as is done with **a**). When the program actually runs, the computer will allocate enough space for an integer. The computer finds another four bytes, in a completely different part of computer memory (not on the stack), and determines it as the area of memory that **int*b** points to. The address of the new, allocated space is now written into the pointer **int*b**. Once the memory is allocated, ***b=4** writes the value of **4** into the new, allocated space.

In our example, using run-time allocation seems silly because it would be very easy to simply allocate space for the integer **b**. Imagine, however, an experimental situation where a piece of hardware is emitting numbers and measurements. Because the size and even the type of the numbers may be unclear until run-time, having the flexibility of being allowed to allocate the memory space as needed is a decided advantage.

¹⁰A final line in this code, much much later, would be: **free (b);** which frees the memory allocated for **b** and puts it back for future use.

¹¹Stack - according to legend, this data structure was inspired by spring-loaded stacks of cafeteria trays. A stack is a buffer - things are pressed down into the stack and they only come out of it in the reverse order they were in pressed in - last in = first out. Program stacks are used at run-time. When a program is run, addresses, variables, all the things that a program needs to know about are inserted into the stack.

¹²An integer takes up four bytes of space in most computers.

Common Memory Allocation Errors in C

A common mistake in the example above is to forget that space must be allocated before it is used, to use the pointer without saying what it is going to point to, and to accidentally write the value (in this case, 4), rather than the address of the value, into the pointer itself.

In the more sophisticated compiled languages, the programmer is allowed to allocate memory as needed, but must take responsibility herself for memory allocation. There are several steps involved but the basic principal remains the same: space must be allocated before it is used and still be allocated when it is used.

The chart below describes a common allocation error.

WRONG Example of Allocation in C

```
char*;  
getname();  
{  
char name [20];  
strcpy (name,"hello");  
return name;  
}
```

It is a char*.
It is a function called getname.
Start body of code.
Find space for 20 characters.
Copy into that space the literal string "hello".
Return with the pointer to that piece of space.
End code.

Unfortunately, this code MIGHT work (for a while). However, it is illegal. Why? It is not enough to simply name the object and set memory aside for it in the stack.

When the compiler reads the sample code `char name [20];`, it knows that the programmer has requested some memory. It allocates it in the stack.

At the command `strcpy (name,"hello");`, the computer writes "hello" onto that piece of memory. This memory is not reserved after the function returns. The code will only work until the compiler happens to use the memory for something else (as it is completely allowed to do). At that point, the memory will be overwritten, causing later problems.

The command `printf("%s",name);`, if given within the function, will print out the name: "hello" will appear. The scope of the object, however, exists only within the function it has been written in and if this command is referred to within another function, the computer will perform the command but important information inside that piece of memory may have been overwritten (and printing that data may be fatal to the program¹³).

¹³For example: in C, a character string is terminated with a 0. Printing the string searches for that 0. If it has been overwritten, the computer will continue to print and may cause a fatal error.

The chart below describes a correct example of a memory allocation code fragment.

RIGHT Example of Allocation in C

<pre>char*; getname(); { char*name; if(name=(char*)malloc(20))==NULL) exit(0); strcpy (name,"hello"); return name; }</pre>	<p>It is a char*. It is a function called getname. Start body of code. I need a temporary character pointer. "If" statement to trap malloc. Copy the literal string "hello into it. Return a pointer to the explicit piece of memory where the name is. End code.</p>
--	---

A Note About Memory Allocation in FORTRAN

In FORTRAN, a base-level compiled language, the programmer cannot allocate memory at run-time. To use FORTRAN, the programmer must allocate as much memory as she thinks will be needed. FORTRAN reads only as much data as memory has been allocated and no more, so usually, the maximum amount of memory is requested. This has some disadvantages: every program tends to become huge, and sometimes the programmer runs out of memory.

Lost Information and Overwrites in Compiled Languages¹⁴

In compiled languages, structuring information about the data (meta-data) is often stripped to make the program tighter after the compilation process. The result is that it is difficult to reference what is contained in the block of data without searching through the code.¹⁵

¹⁴ In interpreted languages, while all the data structuring information is available at run time, an unfortunate reality is that it remains either inaccessible or specific to the interpreter that created the data. In actuality, programmers using interpreters must write data structure information in two or more places.

¹⁵In FORTRAN, for instance, only a small portion of data structuring information remains after compilation. C retains slightly more information.

Referring back to our sample fragment of C code:

Sample Fragment of C code

```
int a;  
int *b;  
a = 3  
b=(int *)malloc (size of (int));  
*b = 4  
printf ("%d\n",a);
```

Once this sample code is compiled and the program is run, the meta-data about **3** (that the programmer named it **a** and that it is an integer) is never used; though it remains around (but not in!) the code unless it is stripped out, it is only used for debugging.¹⁶ Once the original information about the data structure is stripped away it is gone forever, leaving only a machine code pointer pointing to a memory location.

Stripping the code of meta-data can lead to data overwrites: in some compiled languages, pointer space allocated for an integer can be used to store other data structures. For instance, if the programmer specifies space for a floating point number but then writes in a double, the compiled language will do it but will overwrite her memory. This is less of a problem in interpreted languages.

MEMORY ALLOCATION IN INTERPRETED LANGUAGES

Interpreted languages work differently than compiled languages. In an interpreted language, all memory allocation is done at run-time and is automatic. Code is searched through sequentially every time the program is run. As it searches, the interpreted language checks to see if a data object that is referred to exists in memory. If the data object is nonexistent, the interpreted language will allocate memory for it.¹⁷

Sample Fragment of BASIC code

```
Let A=0.23  
Let I=1,047,233,823  
print A,I
```

¹⁶If something goes wrong, the debugger will use the symbols to help find the problem.

¹⁷The fact that code is syntactically correct doesn't mean that it is going to work. Because allocation in an interpreter is all done at run time, within a program there might be a sub-routine which very seldom gets attached. The program may hit the sub-routine and suddenly find that it is out of memory.

What the Interpreter Does

Let A=0.23 translates as:

There is something called A

It's a float (BASIC code for float is 7)

There's 1 of them

Allocate space for it; it equals .023



Let I=1,047,233,823 translates as:

There is something called I

It's an integer (BASIC code for integer is 8)

There's 1 of them

Allocate space, it equals 1,047,233,823



Demonstration of a Crude Search Mechanism

When the interpreter reaches the command, **print A,I**, it begins to search for **A** from the top of the data heap. Recognizing the first item as a float, specifically, the float **A**, it moves down to the next line to find its value, and prints out **1.792**.

After having found and printed **A**, the interpreter returns to the top of the data heap to search for **I**. The first thing it finds is the float **A**. As that is not the integer **I**, the interpreter continues to the beginning of the next item, skipping the rest of the description of **A**. It is able to do this because it knows that a float plus its description needs a particular number of storage spaces in memory. It now jumps down that many spaces and checks the next object. In this case, the second item is, indeed, the integer **I**. The interpreter moves down to the next space to find the value of the integer **I**. Finding that value, **32** it is now able to print it out.

This is a crude process. Many ways exist to speed up it up (binary searches, etc.), but the important thing to remember is that there is a lot of meta-data (that it is a float, that its name is **A**), deeply embedded in the data itself. In an interpreted language program, this meta-data must exist, unlike the symbol tables of a compiled language program which can be stripped away. In our example, the meta-data about the actual numbers **3** and **4** utilizes more memory than the numbers themselves.

SDS' MEMORY ALLOCATION TECHNIQUES/SYSTEMS

In SDS, the meta-data that exists in a compiled language only during the compilation period remains as information directly associated with the data. Unlike in an interpreted language, however, no descriptive information is embedded in the actual data. And unlike in a database management system (which must also contain this meta-data), the full description may be moved with small data sets.

SDS formalizes memory allocation so that the programmer doesn't have to remember how much memory has been allocated. In order to keep track of memory allocated, part of the structure of the SDS header describes the size information of all of its data objects. The SDS description also contains the original memory allocation requests; it enables the programmer to wait until run-time to allocate memory, and the SDS description can be written either at compile time or at run-time.

One philosophy behind SDS is the belief that all the techniques for memory allocation are fine, but that they have their limitations. It is very useful to have naming, structure, size and location information (not just in symbol tables, but up front, as in an interpreted language). However, the programmer can do without extraneous information being presented every time she looks at her data.¹⁸

SDS attempts to use the best facets of the memory allocation techniques of interpreted and compiled languages in the right order (keeping the speed of C and the informative power of the interpreted language, but not in a way that slows down C) and then attempts to make it possible to move these pieces of allocated memory from process to process and across networks.

DATA TRANSFER BETWEEN PROCESSORS - LIFE BEFORE SDS

Sharing information between processors and between computer languages is a general problem in the computer world. Different computer systems provide various models for memory allocation. Interpreted languages embed data structuring information in their data, information that cannot easily be read by compiled languages. Some data structures are more sophisticated and complex than others. One of the major problems is that every computer stores its data structures differently.

Data Structures

There is a short list of basic types of data structures (integers, floats, doubles, char strings, etc.). In FORTRAN and in most BASIC type languages, all data structures must fall within those parameters. In more sophisticated languages, (C, C++ etc.), those structures can be more complex.

```
int value [256];
int min_value;
int max_value;
```

This sample code outlines an array of 256 integers and a minimum and maximum value associated by the naming but not explicitly related to the array. A better way to associate these values would be to define a data structure:

```
struct val;
{
    int value [256];
    int min_value;
```

¹⁸Accessing this information about the data can slow the program down. When this information is not needed, the data can be accessed rapidly. Reminder: in SDS, the data remain in two parts: the basic data objects, and the data structure description. This two part division assures that if data searching is needed it can be done as efficiently as by any interpreted language, and at times if and when data searching is not needed, data access is optimized for use by compiled languages.

```
        int max_value;  
    }
```

This structure defines a new kind of object, called **val**, composed of an array of 256 integers, another integer called the minimum, and another integer called the maximum. These values are inextricably, explicitly mingled, not just by the naming as in the example before, but by the data structure itself. Once the generic structure exists, the programmer can request one and can define its **val min** and **val max**, determine the contents of the value array, etc.

Varying Data Structure Storages in Compiled Languages¹⁹

When data is transferred between compiled languages, the data structuring information is stripped after compilation and the data is transferred in binary form. These binary numbers can be interpreted in any way. The computer has no way to know how the data was formatted on the computer it was created on, and therefore has no way to differentiate pad bytes from significant bytes.

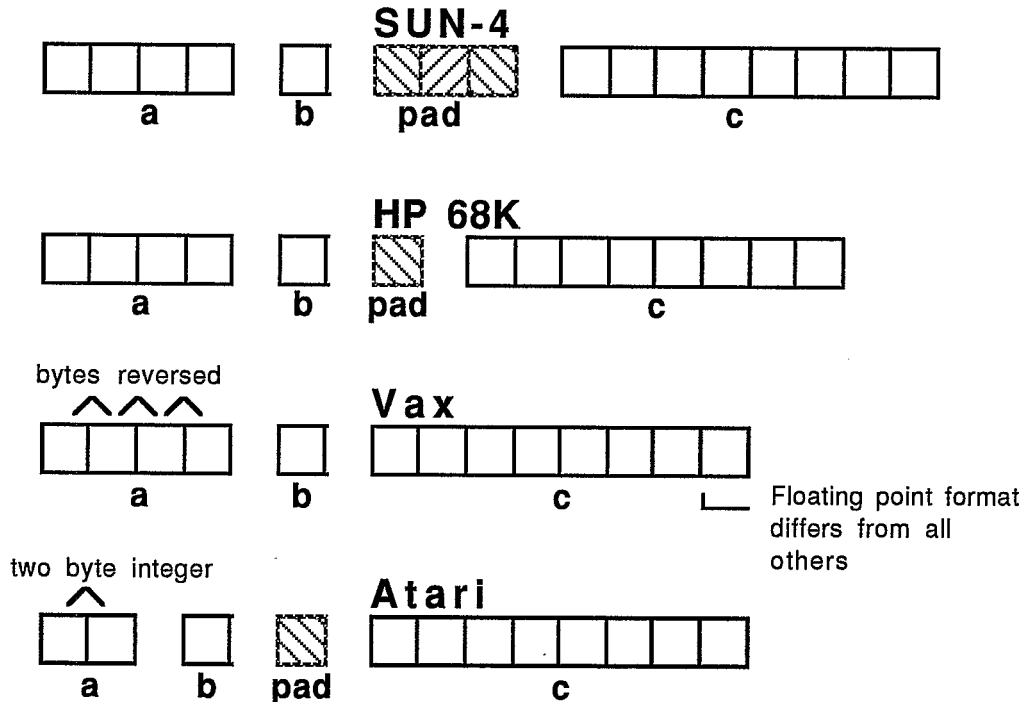
¹⁹The problem with varying data structures is not limited to compilers; data created with an interpreted language on one machine, when given to the same interpreted language on another machine, will not necessarily be understandable.

The graphic below illustrates some varying data structure storages for the following data structure:

```

struct fred
{
  int a;
  char b;
  double c;
}

```



Dealing with Data Transfer - External Data Representation

Data transfer between processors is not a new problem in the computer world and various ways have been invented to deal with it. One example: a division of Sun Microsystems promotes a technique called "External Data Representation." When transferring information between two unknown computers, it automatically translates data into a specific format as it is being put onto the link. Any computer receiving the information will know, for example, that there are no pad bytes in the data, etc.


This technique works. Its major drawback is that each process automatically translates its data into this format even if they are going to the same type of processor (an extraneous third step), and the data structure must be already known by both ends of the system; no meta-data exists with the data itself.

THE SDS APPROACH TO DATA TRANSFER

SDS has been designed with the belief that a data flow facilitation system should manage data as it has been written and not put structure or memory constraints on the programmer. SDS allows the programmer to determine what data structure model she wants to use for the data. The SDS header then simply describes that model.

SDS puts data onto the network in the format it has been programmed in but, unlike with the "External Data Representation" technique, provides enough information about the structures, sizing, memory

allocation, etc. so that any computer can translate the data into a readable format. By presenting the data in its own format, SDS avoids the problem of unnecessary translation.



SDS places a layer on top of the data that defines the structure for use by all systems and languages. It allows data from more sophisticated structures to move into less sophisticated structures by providing a bridge over which the structure can travel. Data cannot be directly transferred between Basic and FORTRAN. Data can be transferred from FORTRAN to C, but it can be difficult to transfer data from C to FORTRAN²⁰. Data can be transferred from C to C++ but because of differing data structure, transfer from C++ to C is difficult. SDS enables data transfer in any direction but cannot alter the conceptual structure of a language - data transfer of complex data structures to less sophisticated is difficult even with the SDS description.

SDS presents a lump of binary data which can be mapped efficiently, while allowing the data structuring information to remain accessible to the programmers by placing it outside the data. Any piece of published or transferred data keeps the descriptive information with it. Programmers seeking access to the data are presented with mechanisms which inform them what is inside the block of data.

For our data structure example,

```

struct val;
{
    int value [256];
    int min_value;
    int max_value;
}

```

the SDS description would explicitly state that this structure has an array of 256 integers followed by an integer called the min_value, followed by an integer called the max_value. It would specify that its name is **val** and that this data came from a computer with, for instance, a "R.I.S.C." padding strategy. By having all this information, a computer with, for instance, a "Motorola 68000" or a "Vax" padding strategy can understand how to translate the data to a format it can understand.

TRANSFER WITHIN A SYSTEM

SDS also aids data transfer between programs within a computer, using names and internalized pointers to sort many of the data coherency problems of shared memory that exist in compiled languages. The following discussion will review memory allocation techniques within process and shared memory in order to clarify how SDS operates in this area.

ALLOCATION WITHIN PROCESS AND SHARED MEMORY

²⁰To transfer a repetitive structure from C to FORTRAN without SDS, FORTRAN must allocate the memory. Typically in this form of data transfer, the data will have to be copied.

Within a computer's process memory, a program with all its instructions and data are stored in a piece of memory. Should the program need more memory, it can ask the computer to allocate more space. This additional allocation happens in one of three ways depending of the type of computer being used:

In computers which don't use a memory manager (real time systems, Apple II,²¹ etc.), the RAM is directly accessed. As a result, the memory is not protected and any process can affect it.

In computers with a memory manager, the program requests memory from the memory manager. When the memory manager accesses the memory, it gives the program a pointer to that memory. That pointer may be an indirection; it may not map directly to the actual memory location but, because the memory manager is policing the memory access, the actual location of the memory is irrelevant to the program. Use of the memory manager also insures that the memory allocated will be reserved for the program and isolates the program so that any wrongful memory access will not affect any of the other programs.

In computers using a virtual memory system as well as a memory manager, the program is fooled into thinking that the computer has more memory capacity than it really does. The virtual memory system does this by swapping bytes from RAM to disk space to RAM as is needed. In this way, an 8 megabyte RAM system will seem to the program as though it has 100 megabytes. Virtual memory systems work because most programs' data access tends to be localized.

Sharing Data Originating in Process Memory

Process memory provides fast data access but doesn't provide the ability to share data with other programs within the computer. Sometimes more than one program within a computer wants to share the same data. This can be done by filing the data generated in process memory on disk and allowing the programs interested in that data to make a copy of it for their own process memory. However, this can be wasteful and time consuming, many copies of the same information can use up valuable space, and there may be problems with data coherency; if one of the programs updates the data, it must inform all the other programs that have copies (via the disk) how the data has changed.

Shared Memory

Shared memory allows data access to more than one program within a computer. Technologically, it is immensely more effective than process memory when more than one program is using the data. Accessing the live memory is much faster than accessing files and, because there is only one copy of the information, the "dance" of loading from file, copying, altering, returning to file, etc. does not slow things down and data coherency problems can be more efficiently handled. (Of course, using shared memory is also more dangerous than using process memory; as there is only one copy of the data, mistakes are more difficult to remedy and there may be problems with access control²²).

How Shared Memory Functions

In the Unix system, the interprocess communications suite (ipcs) handles shared memory. When a program needs a piece of shared memory, it requests a particular size of it. If it wants to increase the size, the computer allocates more memory but (as in `realloc()`) the new piece of memory is not necessarily attached to the old piece of memory. The new memory may exist in a different location.

Thus, if two programs are connected to a lump of shared memory, one generating data and one reading it, and the program generating data realizes it needs more space, ipcs may allocate that space in a different

²¹In real-time systems, direct memory access provides the advantage of speed; there is no memory manager in the way. The Apple II has no memory manager because it is a simpler machine.

²²The memory manager has some control of access, and there are ways to put blocks on data, but there are always control flow problems and a piece of data might be accessed by a process before the generating process has finished its input.

place in shared memory. The generating program will continue to write into it. However, the program reading the shared memory data will still be attached to the piece of memory containing the old data and it will stay there until it is instructed to release it and the connection is broken. Only once it has been instructed to do so can a new connection can be made to the new piece of shared memory. In order to avoid this situation, a program-to-program control message structure is needed.²³

SDS AND SHARED MEMORY

Because of the intrusion of the memory manager, different systems cannot use the same addresses (absolute pointers) to access data in shared memory. If one program writes a pointer into a data object, that pointer will not point to the same piece of data in another program. In shared memory, the normal practice of leaving pointers to objects is extremely dangerous.²⁴

SDS sorts out many of the shared memory data coherency problems of compiled languages that exist because of the fact that absolute pointers can't move between programs. SDS allows data objects to be addressed within shared memory either by name or by guaranteed internalized pointers. (Internalized pointers are pointers that do not point directly to a data address, they are set offset to the data, usually relative to the beginning of the data.) For each program, the SDS code will give the correct pointer to the offset position.

WARNING FOR SDS USERS: When moving SDS datasets from from process memory to shared memory, please be **VERY CAREFUL!** A common error is to continue to use process memory even **AFTER** the data has been copied into shared memory.

DATA DESCRIPTION

SDS serves not only as "glue" between the bits and bytes of various systems, languages, and applications, but also operates on a higher level of data description. One of the important SDS data description features is the SDSListk, a type of list.

LISTS AND ARRAYS

A list is a very general organizational data structure, a way of associating data objects. When data objects are associated in a list, they can be easily searched for, reorganized, etc.²⁵ While many of the things done with lists can also be done using arrays, it is far easier to organize data objects with a list. (Arrays, in fact, are actually a simple type of lists; automatically built in to many computer languages.) While lists require the additional effort of extra code, they are ultimately far more flexible than arrays.

List code associates things in a linear or a multi-dimensional structure. The items in the list exist in memory connected only by list pointers that go from item to item. To add a new name to the list, all that must be done is to change the pointers, to break and reset the linkages. No duplication of data items is needed and, in the SDS system (SDSListks, the SDS list), there is no manual memory reallocation.²⁶ A list may be multiply linked and data items can be associated in more than one list at a time. General rules may be specified for all items on the list.

²³ipcs handles the basic tools for this (message passing and semaphores).

²⁴...and extremely funny, says Chris.

²⁵A common example of list usage: in many compiled language systems, the computer doesn't have any up-front information about the size and memory allocation of its data objects. In order to find how much data are stored, the user must keep requesting data until there is no more. This can be dealt with by producing a list; a directory structure, where pointers are kept to each new piece of data.

²⁶There are times when memory must be reallocated because the programmer doesn't know how many of a given type of data object she will need.


An array, in contrast, is a linear association existing in contiguous space in memory. To add an item to a linear array, a new space in the array must be created and the items already in the array must be shifted in order to insert the new item. In order to associate an item in more than one array at a time, it must be duplicated.

LISTS IN THE C++ IMPLEMENTATION - THE SDSLISTK

Lists are a highly effective organizational tool until the programmer wishes to pass the list outside the program. Because List code relies heavily on pointers, the code becomes unavailable.

The basic ISTK list class is a Listk. SDSListks add persistency to Listks. Whereas basic Listk code requests a pointer to a data object, SDSListk code requires more information about the data object so that it can be identified both internally and externally, allowing lists to be transferred, saved, examined and manipulated by different processes, wherever they might be.

SDSListks list Data Frags: the data being associated by the SDSListk²⁷. Around the Data Frags are the ListkItems, objects which serve as descriptive wrappers.²⁸



An Alphanumeric List

A ListkItem contains a primary and secondary pointer, the name of the Data Frag, and its identification number. The primary pointer in a ListkItem points to the actual code which gives access to the Data Frag. This code cannot be saved, only a summary of it (the descriptive name) can be saved.²⁹ The secondary pointer points to data which defines how the code object (Data Frag) will be used. This data can be saved and, once reconstituted from file or memory, these encapsulated data "specs" can be used in a variety of ways.³⁰

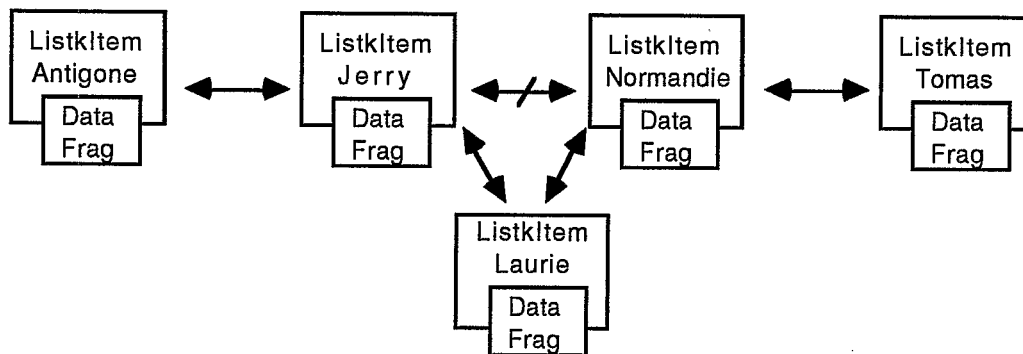
Data Frags in an SDSListk can be accessed by the name, ID number, or primary pointer. They can be run through (like an array); the user can delete or move them. Users don't create ListkItems directly; they are built when new user pieces are added, inserted, etc, to the list. The ListkItems supervise the Data Frags, make the list connections, and allow Data Frags to be added and deleted.

²⁷Like most list code, Listks cope with pointers to objects, not with the objects themselves. Listks do not allocate memory for Data Frags (data objects) that they list (they can't, they don't know what the Data Frags are). Also, they don't allocate memory for names UNLESS and UNTIL the user explicitly asks for a name heap to be created (in preparation for the SDSListk).

²⁸In an array, ListkItems are not necessary, they are contained in a rule that the Data Frags will be placed in contiguous spaces in memory.

²⁹If this were a complete Smalltalk environment, the code would be saved as well.

³⁰For instance, the user may want a textual description of a picture. She would use the data provided in the secondary pointer and create a different, textual object instead of the original visual object created and described by the original primary pointer.



Adding an Item to An Alphanumeric List

To add a new Data Frag to an SDSListk, the user presents the new Data Frag and requests that it be added to the list.

Persistent SDSListks

SDSListks are persistent, eliminating the problem of transferring pointers from one running program to another. To save an SDSListk for later use, or use by another program, the information referred to by the secondary pointer is collapsed and two new objects are created. The first item is a concatenated array of driving data about the Data Frags. The driving data can include such things as definitions of coordinates, colors, backgrounds, etc. The second object contains an array of the ID numbers, the name heaps, and the pointers; information about how the Data Frags were used. When these two objects are stored and brought back, all the necessary information is there to reconstitute the SDSListk. The only thing that cannot be done is to store the actual code.

GRAPHIC - SID Picture

HIGHER ORDER NAMING CONVENTIONS

There are additional conventions that can be used to augment the low-level SDS headers. An example of higher naming conventions follows:

When transferring data from an SDS to an SQL database, the SQL command **select * from table** will result in a database output that can be mapped directly into a single level C-structure and then loaded into an SDS dataset with all of its meta-data (row names types, etc.) intact. The command **select * from table where table.value >42**, will result in a subset of the data of the previous command (although the data structure is identical). The low-level SDS header has no way to contain the extra meta-data (the fact that the database output is a subset of the table and not the whole thing). Obviously, this can be problematic.

The SDS tool **db2sds** adds another layer of information, a simple higher naming convention, to the meta-data existent in the SDS header. If and when the "where" clause is present, an additional string data object will appear in the dataset named **table.where**. For the example above, the additional objects value would read "value > 42". This extra data could subsequently be used, for instance, to label a graph of "table" or to act as the qualifier when **sds2db** (the reverse operation, transferring data from an SDS to an SQL database) attempts to update the original table in the database.

SDS COMMUNICATION - COMMUNISTKS AND GLISH

The Integrated Scientific Tool Kit (of which SDS is a major component) has two high-order communication features, Communistks and GLISH, which SDS uses to access data objects and handle the control flow between programs/systems.

Communistks are the communication class for the ISTK libraries. A Communistk is a data object which focuses on a value and has a list of other data objects to notify when this value, or the Focus itself, changes. Communistks exist only in the object-oriented languages. They enable, and hopefully simplify, communication between objects in a program and, by using the GLISH connection, between different programs within a processor and between programs running on physically distinct processors.

For a successful communication set-up between data objects, several things are required: one or more senders, one or more receivers, and a message body. Communistks are built to take over the mechanics of communication checks and routing for a set of communication types. To this end, they take over responsibility for the message body: the Focus.

The Focus of a Communistk

The Focus of a Communistk is essentially a memory location into which a value is written. The Communistk is responsible for policing what happens to that Focus and for routing advice of any changes³¹ to the Focus' value to data objects which have expressed an interest in the Focus. The interested data objects don't know or care which or how many data objects may change the Focus. They are only interested in knowing if it has been changed.³² The data object which changes the value of the Focus does not notify any other data objects of its action; it is the data objects who wish notification which ask the Communistk to inform them of any changes to the Focus.

Communistk Constraints

The Communistk creator, when creating space for the Communistk's Focus, can assure that attempts to write in a value of an incorrect value type will be blocked. Additional constraints on that value (hard minimum and maximum values, warning and alarm values, restrictions on single-hit changes, etc.), may be placed using the Communistk Constrainer class. These constraints are associated with the Focus itself, not with the data object that is inputting the value change. Any attempt to change the value of the Focus are properly checked³³ and illegal actions will be thrown out by the Communistk.

Change of Focus

In addition to the change in a Focus' value, the Focus itself may move from one data object to another. An analogy for this might be the "daily specials" list at the local restaurant. Joanne's Restaurant has a list of dishes served only on Mondays (roast beef with potatoes, vegetarian casserole, shrimp salad), a list of dishes served only on Tuesdays (Chef's Salad, Sloppy Joe with boiled carrots, calves liver saute), etc. These lists remain the same; every Monday the specials are the same as every other Monday. Joanne is lazy, she doesn't like to rewrite the Specials Board every day, so she has seven boards in the closet, each one listing the specials for one day of the week.

Joanne's Special's Boards are analogous to the Foci. Foci value (like the daily lists) remain the same, but a change in the Focus resets which value is being looked at (Monday's, Tuesday's or Wednesday's

³¹The message body may possibly be null; the arrival of a message may be a signal in itself. The first release of Communistks notified everybody whenever ANY change happened to the Focus. Release 1.1 allows more restricted changes to be notified, such as passing warning limits. This code is not fully tested.

³² An object may, of course, be both a changer of the Focus and be interested if anything else has changed it. In that case, though the data object that changes the Focus is already on the list, the Communistk can be instructed not to notify it. Also, when a number of changes are to be made before response by other objects is required, the Communistk can be set to change the Focus value without notifying the data objects on its CommuList.

³³By the SetValue () method. The data object that is inputting the value change does not change the value of the Communistk's Focus directly.

Board). There are a number of methods which do no changing of Focus value but simply declare to all interested parties that the data Focus has shifted.

GLISH

Communistks only work within a process. GLISH, the ISTK component responsible for sequencing and control flow, facilitates communication between different processes and computers, carrying messages from one process to another via the GLISH executive.³⁴ For the programmer, the only difference between internal and external messages is that Communistks which will communicate outside their process must be named. For external communication, this Communistk name becomes the GLISH event name.

CREDITS AND REFERENCES

The Integrated Scientific Tool Kit (ISTK) was developed directly out of the problems associated with running particle accelerators. Development began in the early 1980's by Christopher Saltmarsh at the CERN Super Proton Synchrotron in Geneva. SDS, developed by Saltmarsh, was used at the E778 FermiLab in which non-linearities were added to the Tevetron magnetic field to simulate potential S.S.C. control problems. Other ISTK components used at this experiment were GLISH, a sequencing language rewritten to its present implementation by Vern Paxson at the Superconducting Super Collider Central Design Group, and GLISTK, the ISTK graphic library, developed primarily by Matthew Kane. SDS (and the other ISTK components) are under continuing development and are currently being used at various sites internationally. Continuing ISTK development includes work by Matthew Fryer and documentation by Ericka Lutz, and contributions from user groups such as the S.S.C. Magnet Test Division. The current list of ISTK collaborators (see the mail exploder on largo.lbl.gov) includes people from: the Superconducting Super Collider Accelerator System and Magnet Divisions; Lawrence Berkeley Laboratory; FermiLab; Cornell University; Lawrence Livermore Laboratory; Texas University; and the Texas Accelerator Center.

ISTK Information and Documentation

Computer files available via anonymous ftp from [largo.lbl.gov](ftp://largo.lbl.gov)

Document lists (including SDS and GLOSS man pages) available on [largo.lbl.gov](ftp://largo.lbl.gov) in file: [/usr/local/src/ISTK/Doc/index](ftp://largo.lbl.gov/usr/local/src/ISTK/Doc/index) (from anonymous ftp these will appear on [pub/ISTK](ftp://pub/ISTK)).

Paxson, Vern. *The Reference Manual for the GLISH Sequencing Language*. (vern@ee.lbl.gov)

Lutz, Ericka, and Matthew Kane. *The Glistk Manual*. March, 1991

Lutz, Ericka. *The Integrated Scientific Tool Kit: An Overview of the ISTK System and its Components*. Superconducting Super Collider publication S.S.C.-MAN-0017. July, 1991

Email to:

istk-counsel@psychosis.ssc.gov

istk-gripe@psychosis.ssc.gov

Suggested Reading and References -

Kernighan, Brian W. and Dennis M. Ritchie. *The C Programming Language*. Englewood Cliffs, N.J.: Prentice-Hall, 1978

Cox, Brad J. and Andrew Novobiliski. *Object-Oriented Programming : an evolutionary approach*, 2nd Ed. Reading, Mass.: Addison-Wesley, 1991

Wilson, Stephen. *Using Computers to Create Art*. Englewood Cliffs, N.J.: Prentice-Hall, 1986

³⁴As with Communistks, the sender is not aware of where that message may go - it is the job of the external communication router (the GLISH executive) to do this and to cope with exceptions.