

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type	LIGO-T970160-06 -	E	10/24/97
LIGO Data Analysis System Conceptual Design.			
LIGO Integration Group			

Distribution of this draft:

**ELECTRONIC
COPY**

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project - MS 18-34
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: <http://www.ligo.caltech.edu/>

TABLE OF CONTENTS

1	Introduction	5
	1.1 Purpose	5
	1.2 Scope	5
	1.3 Definitions	6
	1.4 Acronyms	7
	1.5 Applicable Documents	8
2	General description	9
	2.1 Product Perspective	9
	2.2 General Requirements	10
	2.3 Functional Overview	10
3	Design	13
	3.1 System Architecture	13
	3.2 On-line LDAS Data Analysis Flow	16
	3.2.1 Data Distribution Unit (DDU)	16
	3.2.2 Signal Conditioning and Preprocessing Unit (SCU)	17
	3.2.3 Data Visualization, Control & Monitoring Server Unit (DVU)	17
	3.2.4 Compute Server Unit (CSU)	18
	3.2.5 LDAS Layout at the Observatories	18
	3.3 Off-line LDAS Data Analysis Flow	19
	3.3.1 Data Archival Unit (DAU)	20
	3.3.2 Data Distribution Unit (DDU)	21
	3.3.3 Data Ingestion Unit (DIU)	21
	3.3.4 Signal Conditioning and Preprocessing Unit (SCU)	22
	3.3.5 Control & Monitoring Server Unit (DVU)	22
	3.3.6 Compute Server Unit (CSU)	23
4	LIGO Data Analysis System Software	24
	4.1 General Design Concepts	24
	4.1.1 Software Standards	24
	4.1.2 Software Development	24
	4.1.3 Software Code Management	25
	4.2 LDAS Software Components	25
	4.2.1 Data Types	25
	4.2.2 I/O Libraries	25
	4.3 LDAS Interfaces	26
	4.4 Distributed Computing Environment	27
	4.4.1 User Interfaces	27
	4.5 Computational Data Analysis (Filtering)	28
5	LIGO Networks	30
	5.1 Networking and data access	30
	5.2 LIGO Laboratory	30
	5.3 LIGO Scientific Collaboration	31

6	Installation Approach	31
6.1	Benchmarking	34
6.2	Prototyping Activities	34
6.3	Implementation	35
6.3.1	Observatories	35
6.3.2	Universities	35
<i>Appendix A</i>	<i>LIGO On-line data analysis flow</i>	<i>36</i>
<i>Appendix B</i>	<i>LIGO Off-line data analysis flow</i>	<i>49</i>
<i>Appendix C</i>	<i>LIGO LAN Architecture Details for Universities and Observatories</i>	<i>51</i>

LIST OF TABLES

Table 1	Applicable Documents	8
Table 2	Data Distribution Unit (DDU) -- On-line	16
Table 3	Signal Conditioning Unit (SCU) -- On-line	17
Table 4	Control & Monitoring Unit (DVU) -- On-line	17
Table 5	Compute Server Unit (CSU) -- On-line.....	18
Table 6	Data Archival Unit (DAU) -- Off-line.....	20
Table 7	Data Distribution Unit (DDU) -- Off-line	21
Table 8	Data Ingestion Unit (DIU) -- Off-line	22
Table 9	Signal Conditioning Unit (SCU)-- Off-line.....	22
Table 10	Control & Monitoring Unit (DVU)-- Off-line.....	23
Table 11	Compute Server Unit (CSU) -- Off-line.....	23
Table 12	WAN/LAN Connectivity among LIGO Laboratory Sites.....	31

LIST OF FIGURES

Figure 1	LDAS Functional Units. LAN or WAN refers to the local area networks either within a Laboratory site or between sites.....	9
Figure 2	LDAS Functional Block Diagram - On-line Analysis.....	12
Figure 3	LDAS Functional Block Diagram: Off-line Analysis	12
Figure 4	LDAS architecture for the on-line component.....	14
Figure 5	LDAS architecture for the off-line LIGO Laboratory University Sites (CACR configuration will change on a regular basis)	15
Figure 6	Layout of on-line LDAS components within the OSB at the observatories.....	19
Figure 7	Block Diagram of the LDAS Software Components.....	29
Figure 8	NSF's vBNS Infrastructure will allow LIGO Livingston to link up with other Laboratory Sites via NCSA and Louisiana's regional access via SEPSCoR to vBNS (LIGO and SEPSCoR have been added to map; NSF map is from mid 1996)	32
Figure 9	DOE's ESnet Infrastructure will allow LIGO Hanford to link up with other Laboratory Sites via PNNL (LIGO has been added to map).....	33

Figure 10 Data analysis flow for binary inspiral detection48
Figure 11 Data analysis flow for periodic source detection50
Figure 12 Planned LIGO General Computing LAN at Caltech showing the link to the Caltech LDAS LAN and also to MIT and Sites.....51
Figure 13 Planned LIGO General Computing LAN at MIT showing the link to Caltech and to the Sites (TBD)51
Figure 14 Planned LIGO General Computing LAN at the Observatories showing the link to the Universities and local LDAS LAN52
Figure 15 LAN Interfaces for the Observatories.....53

1 INTRODUCTION

The LIGO observatory systems will generate large amounts of continuous data (approximately 12MB/sec for the Hanford site). At the observatories, these data must be processed in such a way that they may be distributed to various on-line analysis and diagnostic systems. At Caltech, the recorded data will be stored in the LIGO data repository or archive. From here, the data will be reduced for long term archival storage; they will also be available for distribution across the members of the LIGO I scientific collaboration for the purpose of supporting a variety of off-line research and searches. The off-line component will provide a level of parallel processing capacity provided through a collaborative agreement with Caltech's CACR. Data analysis and distribution will be possible via the LIGO wide area network (WAN) that will be available to interconnect LIGO Laboratory sites and sites of institutions who are members of the LIGO I Scientific Collaboration. These functions will be provided by the LIGO Data Analysis System (LDAS).

1.1 Purpose

This technical note presents the conceptual design for the LIGO Data Analysis System (LDAS). This design is in accord with the LIGO LDAS Design Requirements Document (DRD), LIGO T970159. Since it is early in the prototyping phase of data analysis, it is expected that new insights or experience will lead to an evolution of the requirements over time. This is particularly true for those components of LDAS that provide data reduction methods. It is the intent of this document to present a baseline design, based on the present requirements, which is flexible enough to incorporate new features in the future.

All designs presented in this document will identify technology available at the time of writing or known to become available within the design and development phase of the system. However, since the field of data analysis and signal processing is rapidly expanding as new technologies become available, this will open the design to better, faster, cheaper options in the future. Therefore, where possible, flexibility has been included in the design to allow future developments to be incorporated wherever they are applicable. While specific hardware components are shown in the baseline design presented in this document, these will not necessarily be the ones implemented.

1.2 Scope

The LDAS is to be developed to meet the requirements set forth in this document. The LDAS shall have the means to:

[A] Provide on-line data analysis capability to each of the LIGO Observatories. This capability includes the following:

- A means to extract physical strain from the interferometer output(s) and to utilize relevant ancillary channels (e.g., PEM) to remove instrumental or environmental signatures.
- A means to process strain data through real-time detection algorithms for both performance monitoring and scientific purposes. Sufficient computing power to allow processes to keep up with the incoming datastream shall be provided. Sufficient margin shall also be provided to accommodate maintenance down times and other system inefficiencies.
- A means to cross-correlate data (either time series or event lists) from multiple interferom-

eters.

- A means to store data frames and analysis results (local to the Observatory LAN) to short term storage media. This functionality will be provided by the LIGO DAQS resources, with possible augmentation by LDAS.
- A means to access both “live” and short term archived data via the Observatory LAN and the LIGO WAN. Access shall be subject to available bandwidth and demand.
- Means to retrieve, concatenate and extract specific channels of recent data from the on-line storage system.
- Sufficient automation to run continuously and autonomously during periods of normal operation.
- A means to display and visualize results of analyses over the Observatory LAN.

[B] Provide off-line data analysis capability. This capability is likely to be concentrated at one LIGO Laboratory site but shall be available “seamlessly” throughout the Laboratory. This capability includes:

- A means to reduce the raw data to science data representing calibrated GW strain data and a reduced subset of ancillary data and a data quality descriptor.
- A means to archive, retrieve and distribute reduced datasets acquired over a period of time at least 5 years in duration.
- A means for duplicating reduced datasets either for backup or for distribution.
- A means to access the data archive via LIGO WAN by the LIGO Laboratory and LIGO Scientific Collaboration with sufficient bandwidth to support database manipulation at the off-line site by remote users.
- A standardized interface for visualization tools to allow observers to see and interpret results from various analyses.
- Sufficient computing margin to enable multiple analyses to be conducted in parallel.

Specifically not considered to be within the scope of the LDAS are:

- Data analysis functions performed at centers other than the LIGO Laboratory Facilities.
- The on-line diagnostics system used for stimulus-response characterization, transfer function determination, and calibration functions. However, it is expected that software developed for the LDAS will find utility within the diagnostics system and vice-versa.
- Simulation capability shall be provided separately from, but coordinated with, the LDAS. The simulation environment is being developed also within LIGO.

1.3 Definitions

LDAS LAN: The local area network (LAN) dedicated to supporting LIGO Data Analysis Systems. Access is managed through user accounts and passwords.

Observatory LAN: The local area networks (LAN) associated with the observatories. Access to the Observatory LAN and the associated network resources are managed through user accounts and passwords.

LIGO WAN: The wide area networks (WAN) used to connect the observatories, the LIGO institutions, the off-line analysis system and the broader LIGO community. Access between sites on the LIGO WAN will be through telnet, rlogin, ftp and WWW protocols.

On-Line: Data are considered to be “on-line” if they are readily available to clients via the LDAS networks from short term storage.

On-Line Analysis System: The system which processes the data stream in real-time for performance monitoring and for the detection of astrophysical events which are time-critical. There shall be two similarly configured systems (although not necessarily performing identical operations at the same time or having the same capacities) at both LIGO Observatories.

Off-line Analysis System: The system which is used to reduce, archive, retrieve, analyze and duplicate datasets after these have been collected and transmitted to the data repository site. The system shall also provide sufficient computational capacity to members of the LIGO Scientific Collaboration.

Real-time data: The LIGO datastream written to disk cache by the CDS DAQS. The actual latency in the data is determined by the length of frames and access time to read these off the media.

Real-time analysis or processing: Analysis of data so that the information contained in them can be extracted on a time scale sufficiently short so that it is possible to influence either the improvement or maintenance of detector operational performance or the collateral detection of a potential astrophysical event by other detector systems (outside of LIGO).

Time-critical:

[i]Scientific: having the potential to influence the operation of other (non-LIGO) astrophysical or astronomical detection systems so that these instruments may be employed to observe the same phenomenon detected by LIGO (or vice-versa). The time scale for *time criticality* will range from fractions of an hour to fractions of a day or even longer.

ii]Operations: having the potential to allow recovery of a LIGO interferometer from off-nominal operation by virtue of the information extracted from the LIGO datastream.

Metadata: The descriptive data about data; metadata typically consists of (but is not limited to) catalog, index, content, configuration, citation, signature data that is used to describe data.

1.4 Acronyms

- CACR Center for Advanced Computing Research (Caltech)
- CDS Control and Data System
- CSU Compute Server Unit
- DAQS Data AcQuisition System
- DDU Diagnostic Distribution Unit
- DIU Data ingestion Unit

- DRU Data Reduction Unit
- DVU Data Visualization Unit
- GUI Graphical User Interface
- IFO Interferometer
- LAN Local Area Network
- LDAS LIGO Data Analysis System
- LIGO Laser Interferometer Gravitational Wave Observatory
- MPI Message Passing Interface
- MTB(C)F Mean Time Before (Critical) Failure
- MTTR Mean Time To Repair
- NFS Network File Services
- OSB Operations & Support Building (at Observatories)
- PEM Physical Environment Monitoring (System)
- RH Relative Humidity
- SCU Signal Conditioning Unit
- SNR Signal to Noise Ratio
- TBD To Be Determined
- WAN Wide Area Network
- WWW World Wide Web

1.5 Applicable Documents

The documents listed in Table 1 contain relevant design and specification data for the LDAS

Table 1 Applicable Documents

<i>Document Identifier</i>	<i>Description</i>	<i>Comments</i>
M970065	White Paper Outlining the LDAS for LIGO I	
T970140	LIGO LDAS Software Specification and Design Requirements	In process
T970100	LIGO System Software Design Issues	
	Paraflo User's Guide	Author: R. Williams, CACR
T970078	Interferometer Diagnostics Tests and Tools	Draft
T960009	CDS Data Acquisition System Design Requirements Document	
T950054	CDS Control and Monitoring Design Requirements Document	
T960010	CDS Data Acquisition System Conceptual Design	
T960108	Interferometer Diagnostics Conceptual Design	
T960107	LIGO Interferometer Diagnostics System Design Requirements	
VIRGO-SPE-LAP-5400-103	Frame Library Users Manual	
T970130	Specification of a Common Data Frame Format for Interferometric Gravitational Wave Detectors (IGWD)	
T970159	LIGO Data Analysis System (LDAS) Design Requirements	

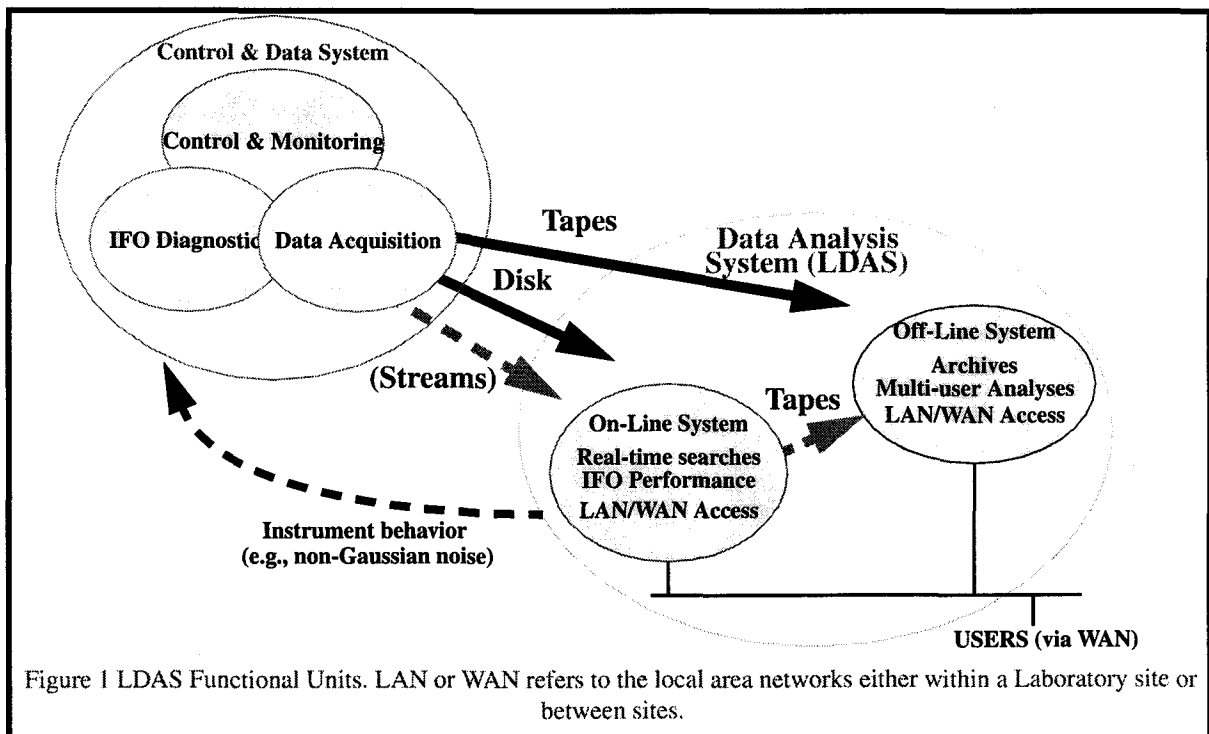
2 GENERAL DESCRIPTION

2.1 Product Perspective

The LDAS is divided into two primary functional units as shown in Figure 1.

These are:

- **On-line LDAS:** The on-line segment consists of two functionally identical but independent units located at the LIGO Observatories (Hanford, WA and Livingston, LA). Each provides the capability to run real-time detection algorithm, and also the ability to provide end-to-end insight into interferometer behavior for specific signal types. The system interfaces to the LIGO DAQS for accessing the real-time data. It also has a limited one-way interface (data displays for operators) for interferometer diagnostics to provide performance metrics (e.g., based on non-Gaussian noise characteristics). The site LAN may be used to access the data cache by scientific analysis workstations at the site. There will also be the ability to access the real-time data remotely by other LIGO Laboratory sites.
- **Off-line LDAS:** The off-line component will likely reside at Caltech's CACR and provides several functions: data reduction and compression for long term archival; data retrieval; and multi-user (independent) refined analysis of LIGO data. Access to the archive will be via a wide area network (WAN) capable of providing high throughput access to the archive. The baseline provides for regular transfer of data tapes from the LIGO Observatories to the off-line analysis center. Once data tapes are received, they will be processed in a yet to be defined manner to extract/compress/refine the science data for the permanent archive.



The on-line and off-line systems shall be independent; however, where it is required, those critical databases or components will be “mirrored” between them. The systems will be highly integrated so that, for example, a remote user will be able to access either the on-line or off-line system with nearly identical user interfaces and commands.

2.2 General Requirements

The specific requirements which this system must meet are given in the DRD, LIGO-T970159. The primary requirement on the system, which drives the design which follows, is to be able to:

- Provide on-line analysis at the observatories.
- Process and reduce the raw LIGO datasets at the off-line center to prepare the data for archival storage and retrieval.
- Provide computational and storage resources for off-line analysis using the archived data
- Provide a flexible design which can be reconfigured to reflect new analysis or computational requirements as they evolve.
- Provide access to LIGO data from all LIGO Laboratory sites and also from member institutions of the LIGO Scientific Collaboration for the LIGO I search.

2.3 Functional Overview

Figure 2 depicts the functionality provided by the on-line LDAS. The on-line LDAS processes the data stream provided by the CDS/DAQS at a rate sufficient to keep up with the acquisition rates. Processing involves data distribution to remote users qualified for access to the datastream. It also involves performing signal conditioning operations on the GW strain channel to calibrate it and to improve the SNR, if feasible, by regressing out instrumental or environmental signatures.

As a minimum, the on-line LDAS will also process the GW (strain) channel to look for one or more astrophysical source types which, if detected, may be correlated with other (non-LIGO) astrophysical detectors. These detection algorithms will be run by one or more processes running at real-time rates to filter the datastream for source signatures. A complementary product of such analyses will be data characterizing the instrument performance in terms of false alarms rates, noise floor sensitivity and spectral variations, etc. Users will be able to access such event and health/status information using X-windows and web technology based interfaces into the LDAS.

The on-line LDAS will have its own local area network (LAN) which is interfaced with the observatory general computing LAN and is also interfaced to the LIGO wide area network (WAN) interconnecting LIGO Laboratory Sites.

The LDAS also has an off-line component resident at Caltech (most likely CACR) which services the entire Laboratory and also the collaboration for the LIGO I searches. Figure 3 shows this component's functional elements. This component serves several functions. It provides access to additional computational resources not feasible to be provided at the observatories. These resources may be utilized and configured to perform more refined analyses than may be possible or warranted at the two sites. It is also here that the full correlated use of data from multiple interferometers will likely be exploited. The computational resources will include dedicated LIGO resources and, on a much larger scale, institutionally provided resources which are available a significant fraction of the time.

The off-line LDAS also provides for the management and creation of the LIGO data archive. Data reduction and compression will be utilized to reduce the data to a volume suitable for commitment to long-term storage and retrieval. Data distribution is provided by the archive both for the (local) types of processing discussed above and also for direct distribution to other Laboratory or Collaboration sites where additional data analysis make take place.

Data ingestion provides for incorporating into the archive new data delivered on media from the observatories. Tape media from the observatories are not expected to be directly compatible with the archive media, hence a translation will need to be performed.

Note that the figures provide functional views generated by the DRD, and they do not translate into distinct and separate hardware (or software) systems. The LDAS is designed such that, to the greatest extent possible, the coupling between the functional units is weak. Functional units will couple over single, well defined interfaces. This will provide a flexible system which may be reconfigured with few problems (e.g. porting a function onto new hardware, moving a functional block from one computer to another).

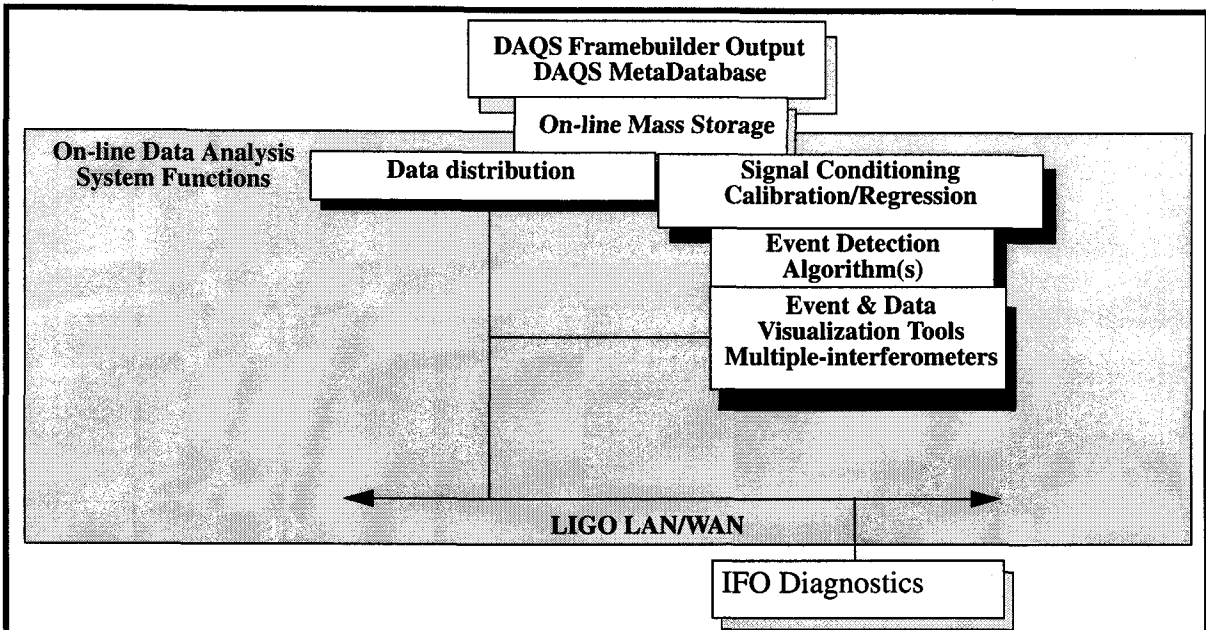


Figure 2 LDAS Functional Block Diagram - On-line Analysis

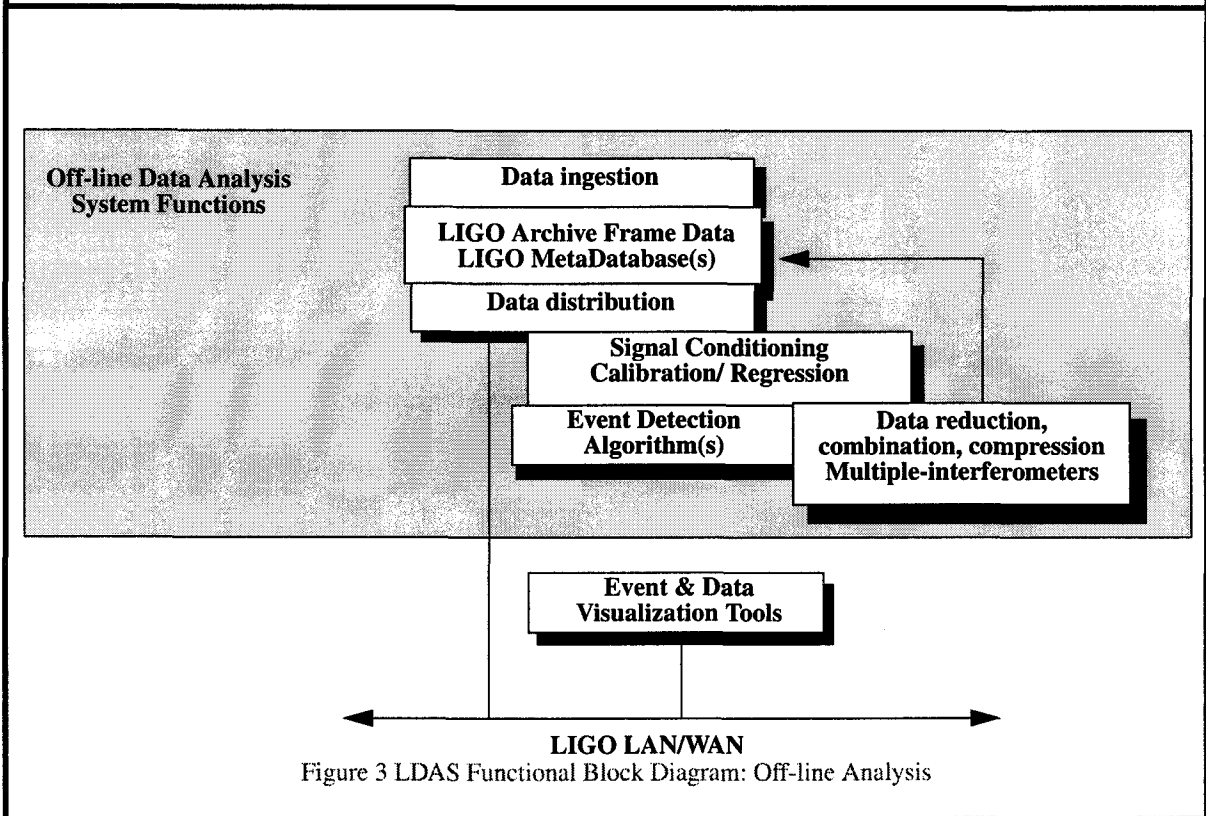


Figure 3 LDAS Functional Block Diagram: Off-line Analysis

3 DESIGN

3.1 System Architecture

Figure 4 depicts an overview of the LDAS architecture for the on-line component at both sites. The Livingston site will require half the processing and storage capacity because it accommodates only one interferometer.

At each site there will be a dedicated OC3 or 155 Mb/s LAN (nominally ATM) to support on-line data analysis. This LAN is linked via a router to the Observatory General Computing (GC) LAN which is similar in design. Both the GC LAN and LDAS LAN are linked via a gateway to the LIGO WAN. It is planned that the LIGO WAN at Hanford will be able to use DOE's ESnet services. There will be a gateway/router configured to transfer LIGO packets off ESnet and onto vBNS either at SDSC or at Caltech; it is expected that at a future date access to LIGO Hanford via ESnet will also be possible directly at MIT. LIGO Livingston is planned to be accessible via NSF's vBNS using LSU as a gateway.

There will be a limited interconnection from the LDAS LAN to the CDS LAN. A gateway will provide limited TBD services. The intent of the isolation is to limit the loads introduced by data distribution and analysis so that CDS is not affected and so that the GC LAN is also not affected.

LDAS units are interconnected either through Ultrafast or Fast Wide SCSI (F/W SCSI) ports to the on-line mass storage system or via LAN to servers and processors in the LDAS.

Figure 5 shows an overview of the LDAS architecture for the off-line component. This component is principally implemented at Caltech/CACR but has an independent component at MIT.

At CACR the network dedicated to intensive data transfer and computation will be the extant CACR HiPPI networking infrastructure. The dedicated LIGO resources within CACR will be linked into the HiPPI in an identical manner to the rest of CACR's resources. In this manner, it will be possible to share resources within CACR transparently for parallelized computation based on MPI. As additional resources become available, the LDAS will be able to schedule access to them in order to provide the optimal quality of service consistent with available resources. The CACR LDAS will be accessible on the local LIGO LAN at Caltech using ATM as an interface. much of this interface already exists in a prototyped environment.

The LDAS component at CACR is similar in function to the on-line component at the observatories. The major difference is in the scale of computational resources that are implemented, the hardware dedicated to data ingestion and to data reduction for archival storage, and the archival storage system itself. Moreover, these resources are accessible transparently to all LIGO Laboratory and Collaboration sites via the LIGO WAN and internet in general.

MIT will have a local LDAS component consisting of high end unix workstations for local MPI-based computation (on a more limited scale) and also consisting of a mass storage system sufficient to contain approximately 600 GB of data at any given time. This allows intensive analysis to be done locally when internet-based access to LDAS at CACR proves inadequate.

Figure 4 LDAS architecture for the on-line component

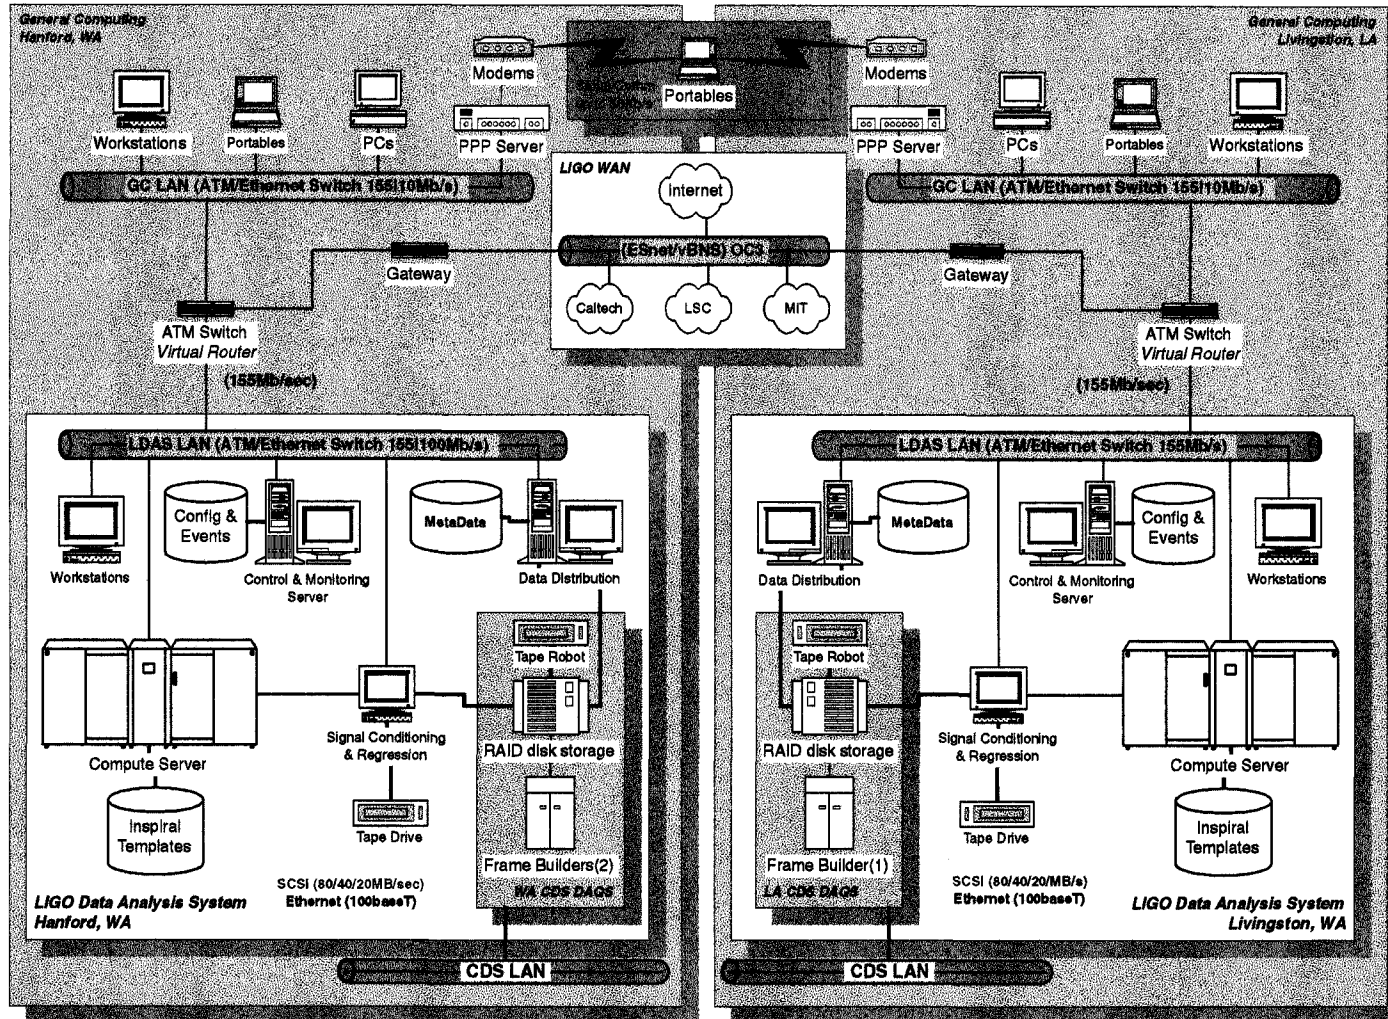
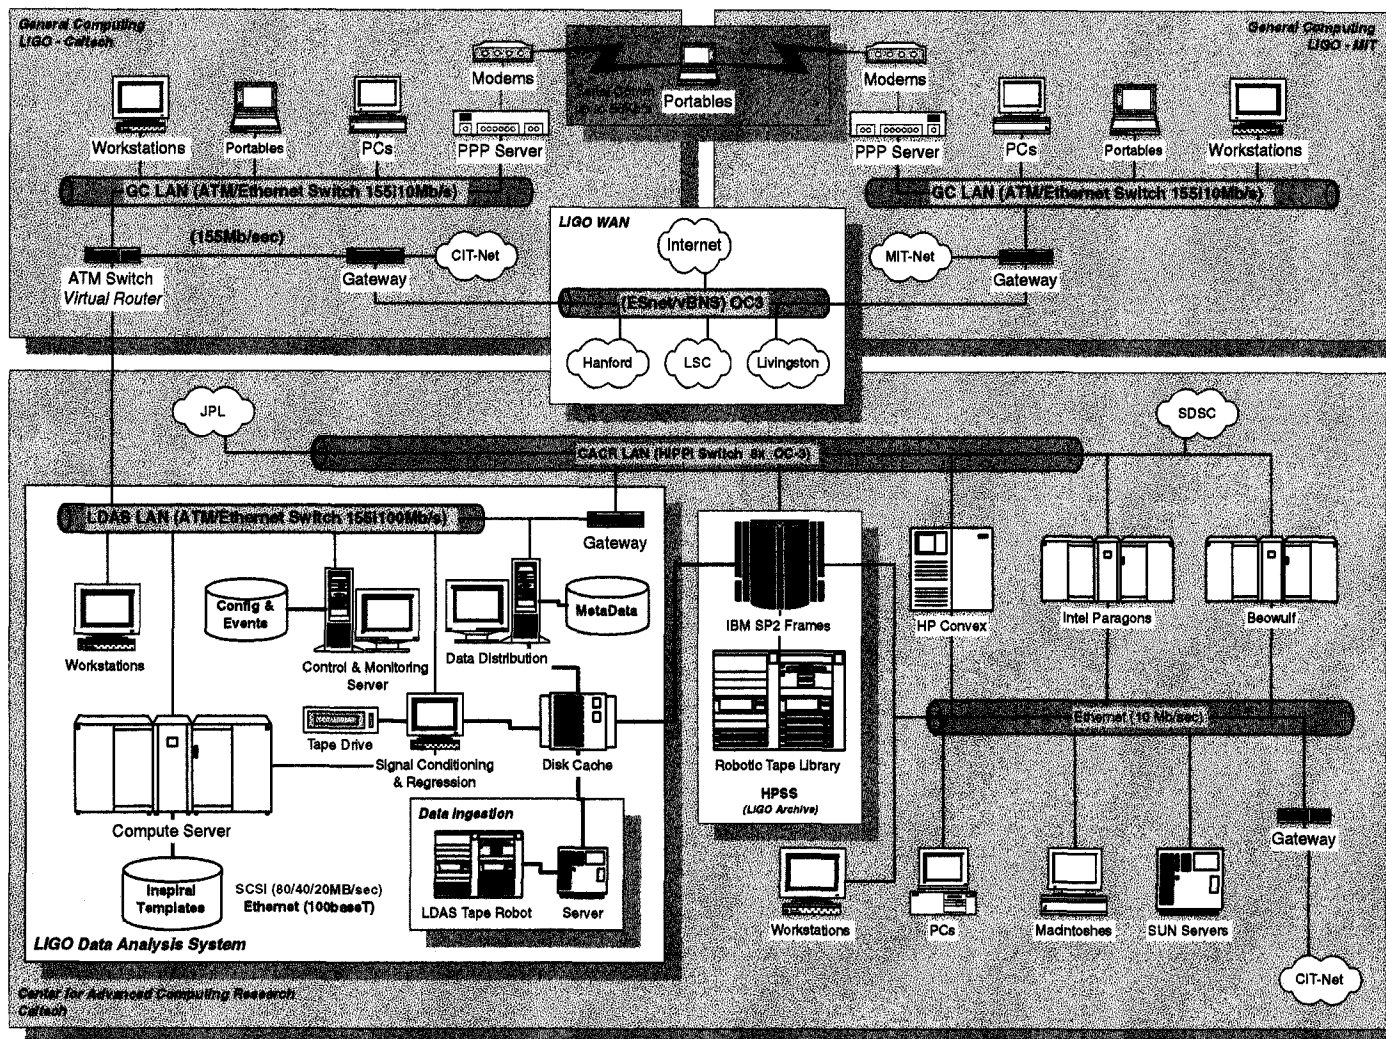


Figure 5 LDAS architecture for the off-line LIGO Laboratory University Sites (CACR configuration will change on a regular basis)



3.2 On-line LDAS Data Analysis Flow

Figure 10 and the preceding discussion in Appendix A describe how data will flow through the analysis process for the specific search for binary inspiraling chirped waveforms. Data are served from the on-line data mass storage system by a data distribution system. Because most presently envisioned analysis algorithms for other sources implement similar optimal filtering techniques, these analyses will also follow a similar flow.

A set of preprocessing steps take the raw data and condition the gravitational wave channel by dropout correction, suitable calibration, unwhitening, regression with other channels, and transformation into the Fourier domain in preparation to hand off the particular epoch of data to a massively parallel processing system designed to perform optimal Wiener filtering for the detection algorithm.

The following discussions rely on details of the data analysis flow which are presented in Appendix A and from which follow the hardware specifications presented in the tables which follow.

3.2.1 Data Distribution Unit (DDU)

The DDU the front end component of the on-line LDAS. Refer to Figure 4 in the discussion which follows. The DDU sits at the interface between DAQS and LDAS. The primary function of the DDU is to make data available to remote users at other LIGO Laboratory or Collaboration sites. The DDU is connected to the LDAS ATM LAN and through this LAN to the site LAN via an ATM switch which establishes a virtual network for the Observatory. Access to the GC LAN or CDS LAN is via gateways which serve to limit ethernet traffic on the respective networks. The GC LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The DDU will serve frame data to clients situated either locally, on the CDS or GC LANs). The specifications for the DDU are listed in Table 2.

Note data distribution to the on-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the on-line processes will not be affected by client loads on data distribution and also so that the on-line process may be isolated from inadvertent interference by off-site remote users.

Table 2 Data Distribution Unit (DDU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
RAID System (shared with CDS/DAQS)	-	-	-	>500 GB striped(TBD)
Ultra Wide SCSI Port	-	-	> 40 MB/s	-
Data Server	>200 MFLOPS	>256 MB RAM	-	-
Network	-	-	- LIGO LDAS LAN (100BT/100Mb/s) - LIGO Site LAN (OC3/155Mb/s)	-

3.2.2 Signal Conditioning and Preprocessing Unit (SCU)

Referring once again to Figure 4, the SCU has four interfaces. The input stream is taken from the DDU on a dedicated high speed F/W SCSI port. Refer to Table 3 for the specifications.

Table 3 Signal Conditioning Unit (SCU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS	>128 MB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 2 ports	-
Tape Drive/Robot	-	-	>2.5 MB/s >5 tape storage	>25 GB/tape
Network	-	-	- CSU direct (≥ 100BT) - LIGO LDAS LAN (OC3/155Mb/s)	-

3.2.3 Data Visualization, Control & Monitoring Server Unit (DVU)

Referring to Figure 4, the DVU provides the user API to configure the on-line LDAS, monitor its health and status, and monitor event lists and other performance metrics derived from on-line processes which are useful for interferometer performance assessment. It consists of one or more high-performance work stations, monitors, and dedicated disk. Access to this system will be provided in the control room in the vicinity of the detector operational monitors so that console operators may have access to scientific information relevant to maintaining acceptable interferometer sensitivity.

The DVU may also be the web server for the LIGO Metadatabase. Refer to Table 4 for the specifications.

Table 4 Control & Monitoring Unit (DVU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS /CPU ≥ 2 CPU	512 MB/ CPU	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Disk	-	-	-	> 50 GB (TBD)
Network	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

3.2.4 Compute Server Unit (CSU)

The CSU provides the computational capacity to process in parallel many optimal filter templates. It receives input data directly from the SCU and provides event lists and other reduced data via the LDAS LAN. Results will be post-processed, displayed, or otherwise manipulated separately using DVU components or other local or remote hardware. Refer to Table 5 for its specifications.

Table 5 Compute Server Unit (CSU) -- On-line

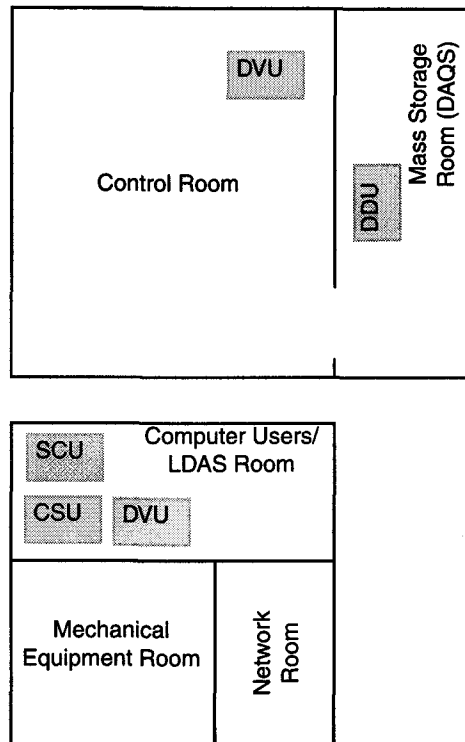
<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	32 nodes, ≥ 10 GFLOPS in aggregate	Per node, ≥ 128 MB RAM	100BT Ethernet	per node, 10GB disk
Fast Wide SCSI Ports	-	-	> 20 MB/s 1 port per node	-
Network	-	-	- LIGO LDAS LAN (OC3/155Mb/s) switched, point-to-point	-

3.2.5 LDAS Layout at the Observatories

Figure 6 depicts a schematic layout of where the principal components of LDAS will be located within the OSB.

The DDU resides in the mass Data Storage Area, in the vicinity of the CDS DAQS because it shares the RAID system with it. The SCU will be located in the Computer Users/LDAS Room. Similarly, the CSU will be located in this area. The DVU, of which there will be several, will be distributed between the Control Room and the Computer Users/LDAS Room. It is anticipated that there will be a need for several DVUs to support operations/diagnostics research and also on-line detection event analysis.

Figure 6 Layout of on-line LDAS components within the OSB at the observatories



3.3 Off-line LDAS Data Analysis Flow

Figure 11 and the accompanying discussion in Appendix B describes schematically how data will flow through the analysis process for the specific search for periodic source waveforms. This is considered at present the most demanding analysis task for which LIGO must prepare to perform off-line. There are several variants on the search task: directed searches looking for expected source waveforms (including spin down); directed searches looking for unknown waveforms; and large-area searches looking for unknown waveforms. The first two problems may also be implemented in part on line for performance monitoring purposes. The last of these problems is presently intractable with computing technology that is projected to be available in the next several years. The limited periodic search will serve as a paradigm for a data analysis flow that must be accommodated by the off-line LDAS component. Of course, the off-line LDAS will also be used to perform refined or deeper (lower chirp mass) searches for inspiraling binary waveforms; this flow parallels that already discussed in Section 3.2

Data are served from the off-line data archive by a system equivalent to a High Performance Storage System (HPSS), consisting of a tape robot and large volume disk cache system which function as the LIGO Data Archive. The disk caching system will be a partition on a shared disk RAID system also serving the data ingestion and reduction processes.

A set of preprocessing steps take the raw data and condition the gravitational wave channel by dropout correction, suitable calibration, unwhitening, regression with other channels, and trans-

formation into the Fourier domain in preparation to hand off the particular epoch of data to a massively parallel processing system designed to perform optimal Wiener filtering for the detection algorithm.

A significant role of the off-line LDAS will be to reduce and archive the raw data arriving from the Observatories. The data flow for reduction will be similar to the signal conditioning process, except that the end product is written back into frame format, possibly compressed, cataloged, and then transferred to the archive for long term storage and retrieval.

The following discussions rely on details of the data analysis flow which are presented in Appendix B and from which follow the hardware specifications presented in the tables which follow.

3.3.1 Data Archival Unit (DAU)

The DAU is the front end component of the off-line LDAS. Refer to Figure 5 in the discussion which follows. The primary function of the DAU is to store and retrieve data from long-term storage. Data are then served via the Data Distribution Unit (DDU, see below) to users located at any of the LIGO Laboratory or Collaboration sites. The DAU is connected to the LDAS ATM LAN and also to the CACR HiPPI network so that the data may be accessed by all CACR computational resources. Access to the LIGO ATM LAN is via gateways which serve to limit ethernet traffic on the respective networks. The LIGO ATM LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The specifications for the DDU are listed in Table 6.

Note data distribution to the off-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the off-line processes will not be affected by client loads on data distribution and also so that the off-line process may be isolated from inadvertent interference by off-site remote users

Table 6 Data Archival Unit (DAU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Tape Robot & Cabinets			≥ 5 Tape heads 6 MB/s	≥ 100 TB 3 cabinet
Disk System (Shared with DDU)	-	-	HiPPI	>250 GB striped for 5 users =2 * #users * vol/tape (25 GB) Partition on shared resource with DDU/DIU
Data Server/IBM SP2	2 node, >200 MFLOPS per node	≥ 512 MB RAM	HiPPI	-
Network	-	-	- CACR LDAS LAN - CACR HiPPI	-

3.3.2 Data Distribution Unit (DDU)

The LIGO Data Distribution Unit comprises the front end component of the off-line LDAS for users of archived data. Refer to Figure 5 in the discussion which follows. The DDU is comprised of a tape archive residing in cabinets accessible by tape robots; it consists of tape readers which dump frame-based data onto a large volume disk cache; it consists of a data server capable of sustained I/O rates sufficient to support up to 5 users at a time each of whom is capable of accessing data at rates up to 6 MB/s. More users may be supported, but a corresponding degradation of I/O throughput to each user will occur.

The primary function of the DDU is to provide access to the LIGO archive to users situated at all LIGO Laboratory or Collaboration sites. The HPSS is connected to the LDAS via the HiPPI backbone which forms part of the CACR network infrastructure. The HiPPI interfaces to the LIGO site LAN via an ATM switch which establishes a virtual network for the Laboratory. Access to the GC LAN is via gateways which serve to limit ethernet traffic on the respective networks. The GC LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The DDU will serve frame data to clients situated either locally or off-site. The specifications for the DDU are listed in Table 7.

Note data distribution to the on-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the on-line processes will not be affected by client loads on data distribution and also so that the on-line process may be isolated from inadvertent interference by off-site remote users.

Table 7 Data Distribution Unit (DDU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Data Server	2 node, >200 MFLOPS per node	≥ 1 GB RAM		-
Disk System	-	-		≥ 250 GB striped for 5 users Partition on RAID shared with DAU/DIU
Fast Wide SCSI Port			>40 MB/s	
Network access	-	-	- CACR LDAS LAN (OC3/155 mb/s) - HiPPI	-

3.3.3 Data Ingestion Unit (DIU)

Referring once again to Figure 5, the DIU provides the functionality of reading in raw data tapes arriving from the observatories for processing and transfer to the archive (DDU). Its server dumps the tape data to local disk cache which is then accessible either for data processing and analysis or transfer to the HPSS. The DIU comprises a single interface to the local disk cache based on F/W SCSI technology. Refer to Table 8 for the specifications.

Table 8 Data Ingestion Unit (DIU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Server/CPU	>200MFLOPS	512 MB		
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Tape Drive/Robot	-	-	>6 MB/s, 3 head 10 tape storage	CDS compatible design
Disk Cache	-	-	40 MB/s, 3 ports	≥ 250 GB Striped for 3 users (3 IFOs) Partition on RAID shared with DAU/DDU

3.3.4 Signal Conditioning and Preprocessing Unit (SCU)

Referring once again to Figure 5, the SCU has four interfaces. The input stream is taken from the disk cache on the DIU on a dedicated high speed F/W SCSI port. Refer to Table 9 for the specifications.

Table 9 Signal Conditioning Unit (SCU)-- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS/CPU ≥ 6 CPUs	1 GB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 2 ports	-
Tape Drive/Robot	-	-	>5 MB/s 5 tape storage	>25 GB/tape
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s) - CSU direct (100BT, 100 Mb/s)	-

3.3.5 Control & Monitoring Server Unit (DVU)

Referring to Figure 5, the DVU provides the user API to configure the off-line LDAS, monitor its health and status, and monitor event lists and other performance metrics derived from on-line processes which are useful for interferometer performance assessment. It consists of one or more high-performance work stations, monitors, and dedicated disk. Access to this system will be provided in the control room in the vicinity of the detector operational monitors so that console oper-

ators may have access to scientific information relevant to maintaining acceptable interferometer sensitivity. Refer to Table 10 for the specifications.

Table 10 Control & Monitoring Unit (DVU)-- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS/CPU ≥ 4 CPUs	2 GB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Disk	-	-	-	>50 GB
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

3.3.6 Compute Server Unit (CSU)

The CSU provides the computational capacity to process in parallel many optimal filter templates. The CSU will be a dedicated LIGO resource within CACR; however when they are available, all CACR computational resources are accessible for LIGO data processing and analysis. The CSU represents a minimum guaranteed capacity available to the LIGO Laboratory. The EPU is sized to accommodate at least the computational load represented by processing in parallel three interferometer data streams. Thus, as a minimum, it will have 3X the capacity of the individual on-line CSUs at the observatories.

The EPU receives input data directly from the SCU and provides event lists and other reduced data via the LDAS LAN. Results will be post-processed, displayed, or otherwise manipulated separately using DVU components or other local or remote hardware. Refer to Table 11 for the specifications.

Table 11 Compute Server Unit (CSU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	96(TBD) nodes, aggregate 30 1GFLOPS	Per node, 256 MB RAM	100 BT (100 Mb/s)	per node, 10 GB disk
Fast Wide SCSI Ports	-	-	> 20 MB/s 1 port per node	-
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

4 LIGO DATA ANALYSIS SYSTEM SOFTWARE

4.1 General Design Concepts

The LIGO Data Analysis System (LDAS) software will be designed and implemented in a manner that grants the system as much flexibility as possible. This will primarily be achieved by making software portability a high priority in the design. The software will be developed to run under POSIX compliant Unix operating systems with the added flexibility of having all graphical user interfaces (GUIs) supported in both the X11 environment and through a JAVA environment, thus providing a direct web technology interface into the system from remote locations.

4.1.1 Software Standards

Software code developed for the LIGO Data Analysis System will be developed using ANSI standard languages and POSIX compliant system interfaces to the operating system. The primary development language will be C++ which has a mature draft ANSI definition at the time of this writing. C++ is an object oriented language which extends the more traditional procedural approach to programming found in C and FORTRAN with concepts of inheritance, polymorphism and operator overloading which increase code reuse and reduce the cost of code maintenance. C++ also provides direct support for traditional ANSI C code, making binding the two languages straightforward.

C++ is a hybrid language supporting both traditional procedural coding practices and the newer object oriented coding paradigms. LDAS software developed in C++ will be designed and developed in the newer object oriented manner to facilitate code reuse, portability and maintainability.

The standard for GUI design will be based on X11 using Motif to provide a common look and feel to graphical widgets in user interfaces directly connected to the LDAS system. While JAVA will be used to develop functionally identical user interfaces for systems remotely accessing the LDAS system such as through web servers and from portable computers intermittently connected to the system.

To further extend the advantages of software standardizations, an LDAS software style guideline will be integrated into the development program, providing a common approach to classes design, code branching, comments and other aspects of software that offer multiple implementation options.

4.1.2 Software Development

Software development will primarily be carried out within the LIGO Data Analysis System group. The software will be developed with place holders (see FILTERS) where new methods for data analysis can easily be integrated. In addition a Software Developer's Guide will be provided explaining how others can integrate new models easily into the LDAS environment.

4.1.3 Software Code Management

Software configuration management includes the activities of configuration identification, change control, status accounting and audits. Software code development for the LDAS will be under software configuration management using Concurrent Version System (CVS). This provides a full history of the development. CVS is an interface to Revision Control System (RSC) with several additional functions to safely allow concurrent development of individual software components.

4.2 LDAS Software Components

The LIGO Data Analysis System software components are laid out in block form in Figure 7. This figure shows the conceptual functionality of the software as well as the relationship between data, analysis and the end user and the components needed to carry out LIGO data analysis.

4.2.1 Data Types

At the highest level, the data types that will be managed by the LDAS are the following:

- **Frame Format Data** - This is the common data format used in the LIGO Data Acquisition System and in the LIGO Data Archive.
- **Frame Meta Data** - This is the complimentary data associated with the raw data collected by the DAQS, registered during data archival, relational data, results of data analysis and configuration information. This data will reside under a database management system.
- **Event Data** - This is the data inputs and data outputs for the LDAS filters. It includes the LDAS configuration and results of searches by the LDAS. This event data is used as a staging area between event detection and event measurements. As event data matures to the highest level of characterization it becomes permanent Meta Data.
- **Light-Weight Format Data** - This is a simplified TBD data format used to allow importation of data from the LDAS into other software analysis environments such as MATLAB, Mathematica, IDL, etc.
- **DAQS Data Tape** - This is the media used by the DAQS to store the LIGO Frame data at the sites. This data will be transferred to the LDAS for ingestion into the LIGO archive.
- **Filter Data** - This is the data set of waveforms and parameters (including templates) needed by the filters used in gravitational wave analysis of LIGO data.

4.2.2 I/O Libraries

The LDAS will need the following input/output libraries to handle reading and writing of the top level data types used in LIGO data analysis:

- **POSIX Library** - This is the Portable Operating System Interface, providing a standard low level interface to the operating system.
- **Frame I/O Library** - This library handles the reading and writing of data in the Frame Format.

- **Meta Data I/O Library** - This library handles the reading and writing of meta data and will most likely be implemented with database technology.
- **Event Data I/O Library** - This library handles the reading and writing of event data and will most likely be implemented with database technology.
- **Light-Weight I/O Library** - This library handles the reading and writing of data in the TBD light-weight format. It will include the ability to present data in ASCII format.
- **MPI Library** - This is the Message Passing Interface library used to perform the high performance distributed computing necessary for LIGO data analysis.

4.3 LDAS Interfaces

The LIGO Data Analysis System will contain components used for configuration, data flow, event detection, event characterization, data ingestion, meta data ingestion, data processing and data analysis. To carry out these actions on the data distributed in this system, the following application programming interfaces will be needed:

- **Data Ingestion API** - This is the component of the system responsible for data ingestion. This involves routing new data from DAQS tapes to the archive, and disk cache, as well as reporting new meta data to the meta data database. The actions performed by this component are controlled through the API.
- **Frame Data API** - This is the component of the system responsible for making request of the Frame Library from the Frame data archive in a form appropriate for data analysis. This includes reducing the content of frames, concatenating frames, performing basic data selection of information stored within frames as requested through the API.
- **Meta Data API** - This is the component of the system responsible for updating the meta database, making requests for information within the meta database, building relationships within the meta database using an API equipped with a query language appropriate for LIGO data analysis.
- **Event Data API** - This is the component of the system responsible for updating the event database, making requests for information within the event database, building relationships within the event database using an API appropriate for various event search strategies, including hierarchical searches and event characterization.
- **Light-Weight Data API** - This is the component of the system responsible for translating data from the frames, meta data and event data into components of the light-weight data format.
- **Control and Monitoring API** - This is the component of the system responsible for configuring, controlling, monitoring the filters being carried out by the LDAS distributed computational resource. This includes detection, parameterization and specialized user requested configurations.

- **Data Conditioning API** - This is the component of the system responsible for preparing data for analysis. This includes conditioning of raw data using calibrations, line removal, regression and bandwidth reduction of the original data. The choice of conditioning is managed through the API. Actual conditioning can be performed within this component or using the filter kernels available through the distributed computing environment.
- **Event Management API** - This is the component of the system responsible for managing events from the filtering kernels. This includes collecting event results, collecting event search results and initiating second tier narrowing of searches for course detection events.
- **LDAS Fast Disk Cache State API** - This is the component of the system responsible for managing the data resources cached on fast disk. This is configurable and included identifying users of cached data, duration of data being cached, as well as types of data cached. The API also manages queuing of requests for cached data.

4.4 Distributed Computing Environment

Computer resources for the LIGO Data Analysis System will be distributed over both local and wide area networks. To perform the scale of analysis necessary for LIGO data, this requires a standard for distributed computing. This standard is handled at the lowest level by MPI. Higher level APIs also exist to focus the specialized needs of LIGO data analysis into cleaner, more object oriented components:

- **Message Passing Interface API** - This is the high level message passing interface layered around the MPI library. It handles the distribution of data and commands appropriate to the LIGO Data Analysis System.
- **LIGO Distributed Data Analysis Manager API** - This is the highest level distributed computing component of the LIGO Data Analysis System. It is the switchboard routing data, commands and results of analysis to the appropriate destinations. In essence, it is the command center for the system. It communicates using LDAS commands, MPI, TCP/IP and WWW protocols.

4.4.1 User Interfaces

Authorized users of the LDAS will communicate with the system through three functionally equivalent, though differing in implementation, user interfaces:

- **LDAS Command Language API** - This API provided the primitives for LDAS communications. The actual interface will support command line and script based interfacing to the LDAS. This API sets under the Graphical User Interfaces discussed next.
- **X11 API** - This API is a Graphical User Interface based on X Windows to provide all the interface to the LDAS. It is basically a smart wrapper around the LDAS Command Language API discussed above.

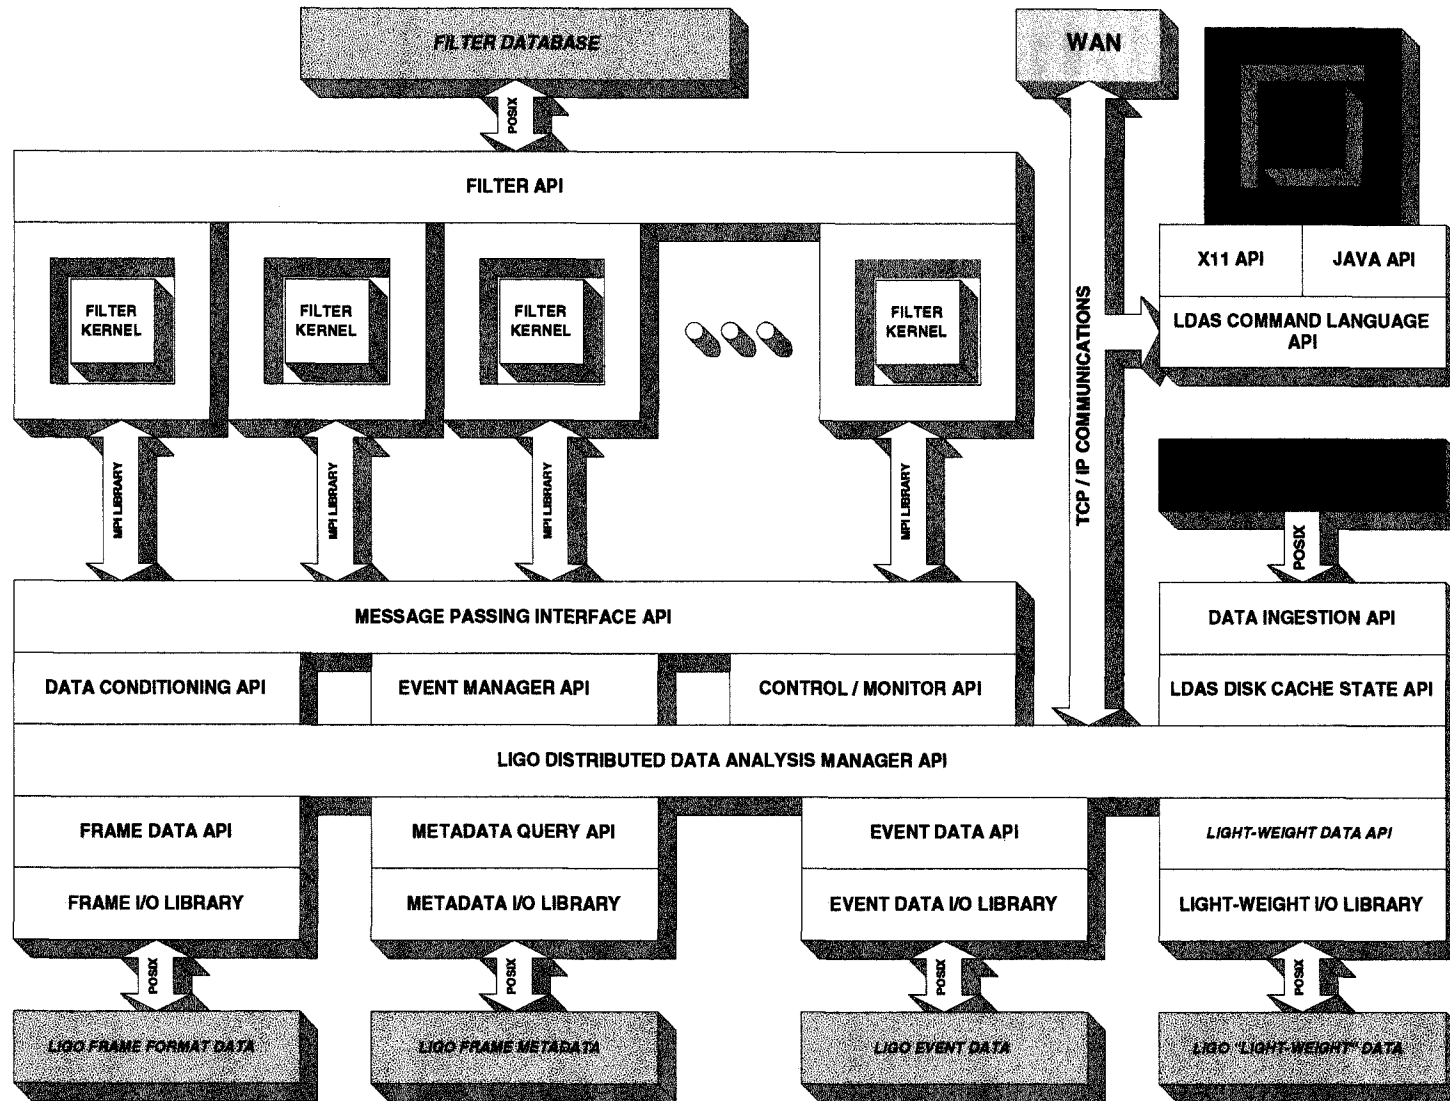
- **JAVA API** - This API is a Graphical User Interface based on JAVA to provide all the interface to the LDAS. It is basically a smart wrapper around the LDAS Command Language API discussed above. Additionally, it can be accessed through the Web services provided by the LIGO Distributed Data Analysis Manager API, providing access from remote machines such as laptops with no more than a JAVA based web browser.

4.5 Computational Data Analysis (Filtering)

Filtering lies at the heart of the LIGO Data Analysis System. Filters of one form or another will be used to carry out virtually all types of searches for gravitational waves in LIGO. Since many of the filters are yet to be envisioned, this component of the system is being conceptualized as modular units with all the interfacing to the LDAS layered away from the actual algorithm so that new filters can easily and quickly be added to the system, possibly even on the fly making the data analysis highly flexible and ultra distributed.

- **Filter Kernel** - This is the compute engine of the analysis filter. It is primarily an algorithm used to by the LDAS to perform a particular type of analysis of the data being distributed from the Data Conditioning API based on configuration directed by the Control and Monitoring API.
- **Kernel API** - This API acts as a wrapper around the Filter Kernel, and will be implemented as a base class in an object oriented environment. It provided the needed integration and communication functionality that allows all filtering algorithms to fit cleanly into the LDAS.
- **Filter API** - This API is responsible for managing various waveforms and parameters used by the distributed network of filter kernels. Among its functions are generating, storing and retrieving waveforms, templates, parameters and other data sets used by various filters that may be active in the LDAS.

Figure 7 Block Diagram of the LDAS Software Components



5 LIGO NETWORKS

5.1 Networking and data access

LIGO is working to develop a wide area network to provide transparent access to data. The wide area network (WAN) will link the four major LIGO Laboratory sites (Caltech, Hanford, Livingston, MIT) as well as institutions in the LIGO Scientific Collaboration. Access to the data archive will be via a web-based technology database query and analysis environment which is being designed and prototyped.

5.2 LIGO Laboratory

- *Universities*

Caltech will comprise of two LANs: the General Computing LAN, inherited from the construction phase and supporting day-to-day activities of LIGO staff at the university, and an LDAS LAN, supporting the data analysis functions described in this document. The LDAS LAN will be distributed on campus between the analysis office area and CACR, where the archive and computational resources reside. Refer to Figure 12 for details on the LIGO GC LAN.

MIT will have a similar topology, although the LDAS component will be of reduced scope, reflecting the fact that there is one common archive and set of computational resources, which will be provided by CACR. Refer to Figure 13 for details about the LIGO GC LAN.

The present baseline is that MIT-Caltech communications will proceed over vBNS.

- *Observatories*

The observatories are in remote locations, and unlike the universities, have no existing networking infrastructure. The plan is to utilize non-LIGO network infrastructure to provide the required connectivity. Each observatory is in proximity to other institutions with whom LIGO is negotiating access to networking resources. Access to the Observatories from the Universities will be provided by a combination of vBNS and ESnet, and the two sites will be linked to the LIGO Laboratory WAN differently.

Refer to Figure 14 and Figure 15 for schematic details of the observatory LAN architecture which complements the LDAS architecture presented in Figure 4 above.

The present baseline is for the LIGO Hanford Observatory to be linked to the LIGO WAN by way of DOE ESnet. The terms for access are being discussed at present. LIGO has requested access to ESnet's ATM cloud system with an initial bandwidth of 3X T1, growing with time as ESnet converts to OC3 and eventually to OC12.

The baseline is that there will be a single crossover of LIGO traffic from ESnet to vBNS. This crossover is TBD, but will either be on the Caltech campus through Caltech/HEP or via SDSC, which already serves as an ESnet/vBNS crossover. MIT communications with Hanford will be routed through this west coast cross-over gateway. It is possible that a similar cross-over gateway may be established at MIT; however discussions on this have not yet to started.

Livingston will be connected to vBNS via a gateway to be provided as an agreement with Louisiana State University (LSU) in Baton Rouge, LA. LSU will be linked to vBNS either directly, or

possibly as part of a consortium comprising of five other southeastern states in a regional network called SEPSCoR. SEPSCoR will be linked, via the University of Kentucky, to the vBNS backbone. LIGO is providing for the establishment of a high bandwidth fiber optic link from its Livingston Parish site to the LSU campus gateway.

Figure 8 and Figure 9 show network maps for the vBNS and ESnet infrastructure. Overlaid on the respective maps are the proposed links for LIGO using PNNL at Hanford and LSU at Livingston as local gateways. Table 12 presents the inter-site connectivity matrix as it is presently planned.

Table 12 WAN/LAN Connectivity among LIGO Laboratory Sites

<i>Site</i>	<i>Livingston, LA</i>	<i>Hanford, WA</i>	<i>MIT</i>	<i>Caltech</i>
<i>Caltech</i>	vBNS/OC3	ESnet (3 X T1) <-> vBNS/OC3	vBNS/OC3	OC3/ATM 100BT
<i>MIT</i>	vBNS/OC3	ESnet (3 X T1) <-> vBNS/OC3	100BT OC3(?)	
<i>Hanford, WA</i>	ESnet (3 X T1) <-> vBNS/OC3	OC3 100BT		
<i>Livingston, LA</i>	OC3 100BT			

5.3 LIGO Scientific Collaboration

The LIGO Scientific Collaboration is composed of the LIGO Laboratory itself and other institutions that have signed memoranda of understandings (MOUs) with the Laboratory. These participating institutions will also require access to LIGO computing and archival resources. Researchers' parent institutions are expected to provide their respective researchers with local network access to LIGO resources through the vBNS network.

6 INSTALLATION APPROACH

At present, several times per year, the computer market is witnessing upward increasing trends in compute power for the same price: CPU clock speeds are expected to double by mid-1998; RAM and disk costs are also expected to decrease; similarly, networking architectures are being introduced which can speed up inter-node communications as well as improve long-distance communications between researchers and computational hardware. These trends continue to evolve following Moore's Law of computer hardware scaling with time. This trend in technology argues for an approach to implementing the the conceptual design described above in a staged manner which can ensure the LIGO achieves the most cost-effective solution available at the time a commitment to procuring computer hardware is required. In order to do this, LIGO will undertake a phased installation approach.

Figure 8 NSF's vBNS Infrastructure will allow LIGO Livingston to link up with other Laboratory Sites via NCSA and Louisiana's regional access via SEPSCoR to vBNS (LIGO and SEPSCoR have been added to map; NSF map is from mid 1996)

The National Science Foundation Very-High-Speed Backbone Network Service Logical Network Map

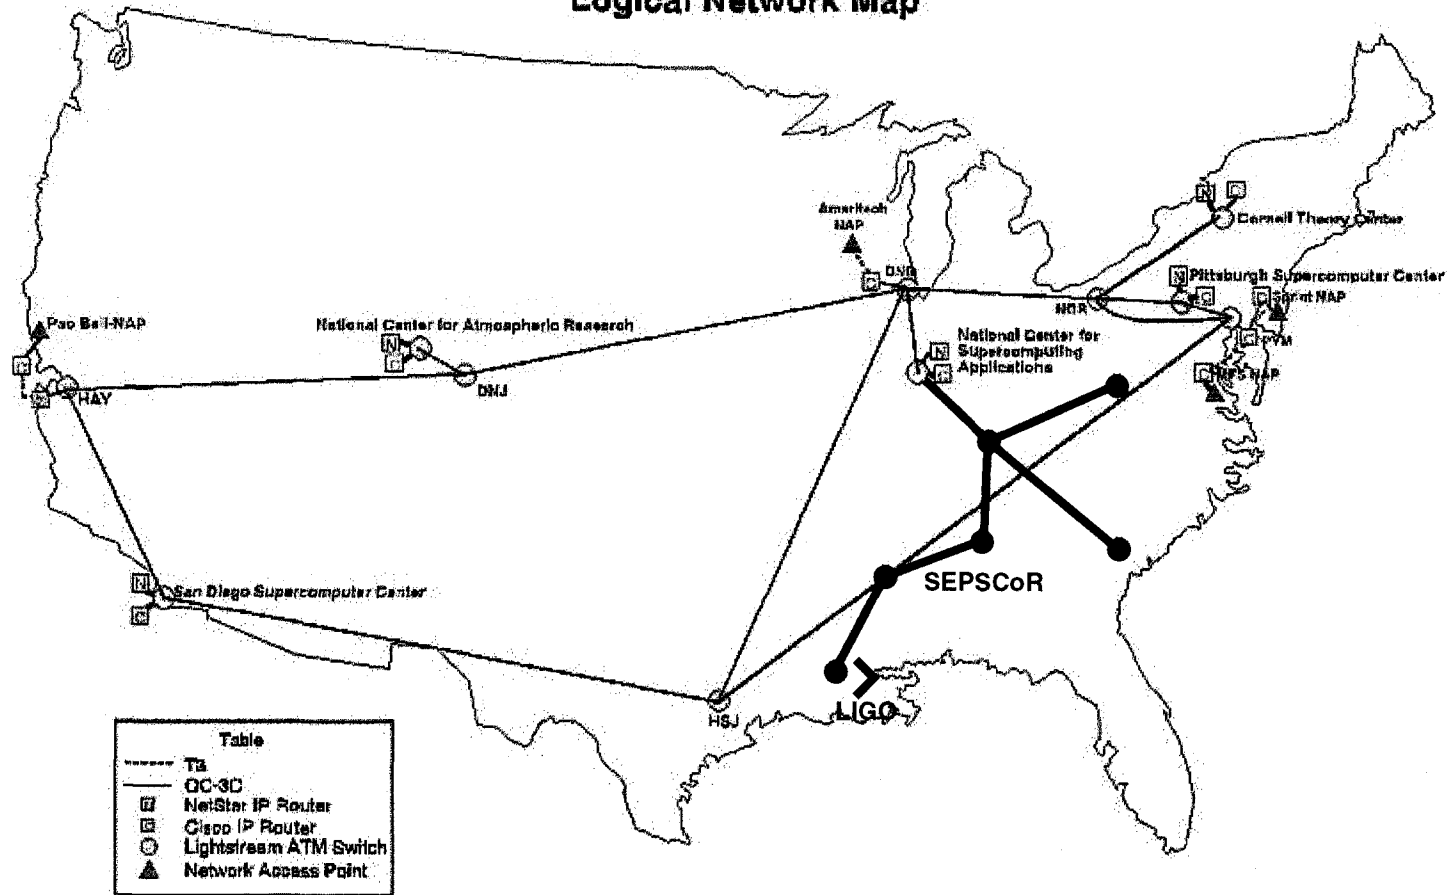
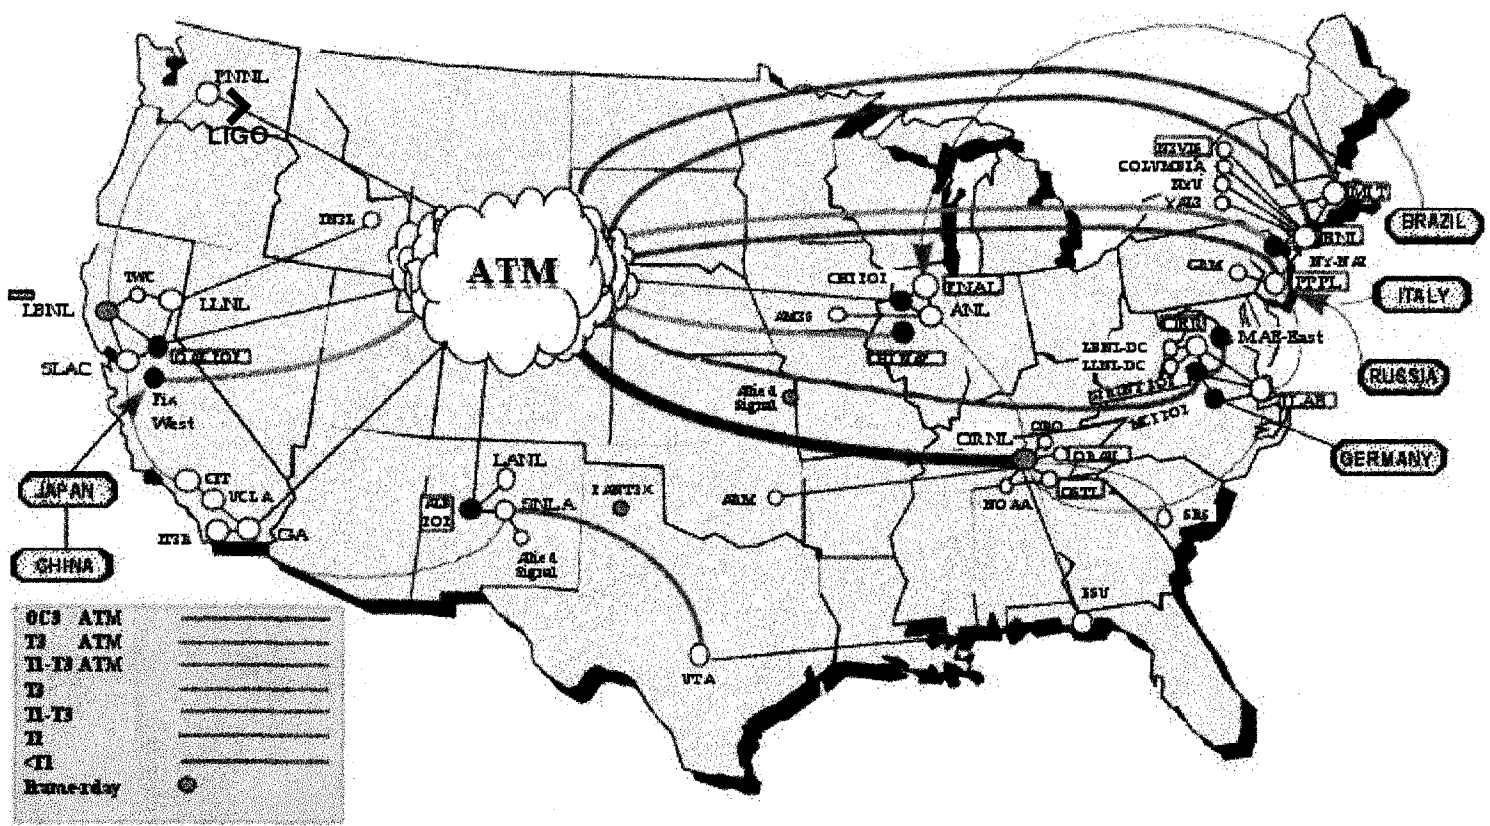


Figure 9 DOE's ESnet Infrastructure will allow LIGO Hanford to link up with other Laboratory Sites via PNNL
 (LIGO has been added to map)

ESnet BACKBONE
Mid 1997



6.1 Benchmarking

LIGO will identify target hardware platforms which can meet LDAS performance requirements by performing a sequence of benchmark tests on candidate systems. This activity is presently being carried as a collaborative effort with the University of Wisconsin at Milwaukee (Prof. Bruce Allen).

A sequence of benchmark tests have already been performed using MPI based data analysis flows for the particular problem of the detection of inspiraling binary coalescence events. The tests have compared shared memory architectures with multinode networked architectures and clustered architectures employing large numbers of stand-alone platforms networked together on a computational LAN.

This activity will be continued in order to identify the best solutions as the hardware continues to evolve. By using prototype analysis programs, tests such as these also serve to verify the design approach LIGO is taking that software analysis tools be transposable across platforms. Platforms demonstrating incompatibility with this design approach will not be regarded further.

6.2 Prototyping Activities

As the results from benchmarking studies become verified and are available, LIGO will procure one or more candidate platforms to verify code development, portability, and expected performance can be met. In the case of clustered workstation solutions, the workstations purchased may be recycled, if needed, as individual workstations for researchers at any one of the LIGO Laboratories.

Hardware prototyping tests will include, but not be limited to, the following activities:

- Individual platform performance;
- Scaling of performance with CPU number;
- Platform POSIX compliance;
- Internode networking switching speed and throughput;
- Mass media storage and retrieval throughput;
- System maintainability and reliability;
- Data archival and retrieval throughput;
- etc.

In addition, prototyping will include software development tests, including, but not limited to:

- Data analysis flows;
- Algorithm development and optimal implementation in hardware;
- Display and visualization tools and concepts;
- Alarm handling; event management tools;
- Database management concepts;
- Data retrieval efficiency and speed;
- Data reduction and compression concepts;
- Signal regression and cross-correlation techniques;
- etc.

Software prototyping activities will focus on developing experience with implementing MPI on candidate platforms, verifying the availability and functionality of commercial software packages which are needed, and to improvement in detection algorithms through experience in prototype data analysis tests using data available from prototype (e.g., 40m) interferometers.

6.3 Implementation

6.3.1 Observatories

The transition from initial stages of commissioning to declaration of operational status for the initial LIGO I searches will likely take as long as 12 months. During this time, it is envisioned that LIGO's understanding of interferometer performance will evolve from an initial phase of discovery and removal progressively smaller, but not fundamental, noise sources to a phase of stable, understood performance at the design instrument sensitivity.

In parallel with this operational evolution, the on-line LDAS will evolve and expand from a small complement of hardware initially supporting hardware debugging and diagnostics activities to the configuration described in this document. Because of the technology trends discussed in section 6, LIGO will increment the hardware investment only to the extent justifiable at that time by interferometer performance: computational power will be matched to the tasks at hand at the time. With this strategy, the least risk and best leveraging of the LDAS budget can be expected.

6.3.2 Universities

Installation of hardware in off-line LDAS components will lag the sites because the need for data reduction, archival, and retrieval will be non-existent in the early phases of detector commissioning. The universities will serve as staging and prototyping sites for hardware destined for the on-line components. Also, data analysis activities at the universities will likely support diagnostics and debugging analyses for the observatories.

At Caltech/CACR, the LIGO off-line component, which features an archival and retrieval system, will become available as the observatories begin to produce data in quantity and of such quality that it becomes desirable to archive them.

At MIT, the principal components of LDAS involve a local mass storage system and a small cluster of workstations; these will be of the same type of hardware which is procured for Caltech.

APPENDIX A LIGO ON-LINE DATA ANALYSIS FLOW

The on-line LDAS component will be used for several functions. The most compute intensive analysis appears at present to be that of searching in real time for binary inspiral chirped waveforms. This is used to determine the target specification for the system. Refer to Figure 10

1. Preprocessing the analysis frame data to obtain calibrated, regressed estimates of strain with instrumental narrowband features removed. This involves several steps.
 - **DATA DROPOUT COMPENSATION.** If the datastream is disrupted during a particular epoch of acquisition - due to loss of lock or some other instrumental or operational problem - it will be desirable to provide for data continuity by either inserting zeroes or mean values into the data stream. This must be done before any other transformations are performed. Process: replace missing or invalid data with predetermined values. Computational impact is negligible.
 - **NARROWBAND INSTRUMENTAL LINE FEATURE REMOVAL.** Using either multi-tapered transform techniques or predictive (e.g., Kalman filter) time domain techniques, narrowband features are removed to decrease signal dynamic range and to improve the SNR of the optimal Wiener filtering step. With a multi-taper method, a number $n < 16$ separate Fourier transformations of the data stream are performed, each with a different weighting function ("Slepian taper filters"). A regression analysis and subtraction process follows to remove by subtraction the narrowband features. Process: $n(<16)$ (real) FFTs of the data stream at the full sample rate (16384 S/s) for a suitably short time $T_{\max} < 1 - 10$ s. This is followed by a linear regression analysis to obtain best estimates of line feature amplitudes and phase offsets for the epoch. Dimensions of the regression analysis is determined by the number of lines to be removed. For m features, the regression involves an $m \times n \times 2$ matrix operations resulting in m pairs of numbers, $\{A_i, \phi_i\}$. The sum of these line feature contributions is then subtracted in the time domain from the raw data.
 - **FOURIER TRANSFORMATION of DATA STREAM.** The data series is converted to frequency domain using FFT techniques. This performed over epochs of duration $T_{\max} < 512$ s.
 - **CALIBRATION AND UNWHITENING.** The data in the spectral domain are multiplied by calibration complex transfer functions to remove end-to-end effects of interferometer response to mechanical strain and data acquisition dynamic range compression. The calibration involves the strain-to-demodulated output transfer function of the interferometer; the unwhitening restores signal dynamic range which was compressed during the acquisition process. Process: frequency bin-by-frequency bin multiplication of the complex calibration function by the line-removed dataset.
 - **DATA REDUCTION.** The previous steps will be carried out at the full temporal and frequency resolution of the acquired datastream. At this point, for inspiral chirp detection, it is possible to compress the volume of data by retaining the frequency range within which the GW waveforms being detected contain most of their signal power. The high-frequency cutoff of the GW channel will be reduced from a Nyquist frequency of 8192 to 512 Hz. This corresponds to a 16:1 reduction in the volume of data to be processed downstream and a corresponding reduction in computation required to process the datastream through the template filter bank.

- **REGRESSION WITH ANCILLARY CHANNELS.** The GW channel will be regressed against a number of other channels acquired simultaneously and which may exhibit instrumental cross-coupling into the GW channel. This regression step, if performed with adequate SNR, will improve the quality of the estimate of GW strain. Process: up to $k (< 16)$ ancillary channels will be regressed over the reduced frequency domain. Each of these k channels will require calibration for instrumental effects across f frequency bins. Then multiplication and accumulation of k complex products over each of f frequency bins.
2. Processing the analysis frame data to obtain candidate events. This involves several steps.
- **DATA PREPARATION STEPS** - assuming one is in the middle of a template analysis cycle, there are a number of steps taken for each data epoch and which is common to all template filtering steps:
 - Obtain next epoch of data from the on-line server
 - Additional preprocessing. For example, additional filtering and sample decimation to reduce the sample rate (if this was not already performed earlier)
 - Time windowing for template sizes. Different mass templates have different temporal durations in the LIGO band. The data need to be weighted by a windowing function to minimize finite-sample effects which will occur in transforming a finite sample of data into the frequency domain.
 - Updating noise floor. It is expected that on some (long) time scale the interferometer noise floor will change; maintaining the Wiener filtering at its optimum requires the noise to be stationary. Tracking slight drifts in the measured noise spectrum can be used to keep the filter at optimum.
 - Fast Fourier Transformation. The signal time series is transformed into the frequency domain for optimal filter processing.
 - The signal is divided by the estimated average noise power spectrum in preparation for the filtering process. This prepares the FFT kernel for the next sequence of parallelized steps.
 - **TEMPLATE FILTERING.** The output from the previous sequence of calculations is provided in common to all template filtering processors.
 - Weight kernel by template waveforms. The kernel calculated above is multiplied by the template frequency spectrum.
 - Inverse Fourier transformation. The function is transformed back into the time domain. The output is a vector whose magnitude for each time bin is a measure of signal to noise. The time bin corresponds to a possible time-of-coalescence.
 - Peak detection. The (complex) time-domain vector is scanned for amplitude exceedence which may be indicative of a coalescence event. The output is one or more pairs $\{\xi_i, t_i\}$ corresponding to the possible event. The index "i" corresponds to template number for which the SNR was observed to exceed a threshold.
 - Events are broadcast. The parallelized implementation of this multi-filter process results in a list of candidate events within the epoch of time being analyzed. These are outputs of the data flow.
 - The cycle continues until all templates have been analyzed.
 - Data are provided for visualization, summarization, and further correlation.

The data flow represented in Figure 10 has been modeled using a spread sheet analysis which tracks in detail the computations required to implement the indicated data flow. Where possible, empirical results derived from benchmark tests using LIGO 40m prototype datasets. These tests have been implemented on a variety of computational resources available at Caltech's CACR.

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
CPU/NODE Engine Model		Calculated 22.09.97
Rated CPU Performance (MFLOPs):	300	400 MFLOPs currently available, 200 MFLOPs PentiumPro
Percent CPU Utilization by FFTS:	70.0%	Based best empirical values measured at CACR
Floating point Operations per Add:	4	Note: adding/subtracting requires a normalization
Floating point Operations per Sub:	4	Note: adding/subtracting requires a normalization
Floating point Operations per Mult:	2	Need to get values in registers
Floating point Operations per Div:	2	
Operations to Cast Integer to Float:	3	
Operations in Memory to Memory Copy:	3	Future hardware will probably off-load this from the CPU
Size of Integer (Bytes):	4	
Size of Float (Bytes):	4	
Size of Double (Bytes):	8	
Size of Complex (Bytes):	8	
Size of Double Complex (Bytes):	16	
Disk I/O Performance (MB/sec):	20	Can do better than this with high end SCSI options
Message Passing Interface (MB/sec):	12	Comes from 100 Mb/second ethernet = 12MB/sec
Number of Nodes in System:	32	Model assumes one CPU per Node
RAM per NODE (MB):	128	
Disk Storage per NODE (MB):	9216	
Memory Resident Operating System (MB):	32	Based on the IBM SP2 which uses 30MB
Disk Resident Oper. Sys. & SWAP (MB):	256	Stripped down version of Unix with approx 100MB of swap

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
System CPU Performance (GFLOPs):	9.6	
System CPU FFT Performance (GFLOPs):	6.72	
System RAM Installed (MB):	4096	
System Disk Storage Installed (GB):	288	
Data Acquisition System Model		
Samples per Second:	16384	Based on CDS DAQS design
Integer Sample Size (Bytes):	4	The option for 24 bit ADCs exists for some DAQS channels
Float Sample Size (Bytes):	4	
Complex Sample Size (Bytes):	8	
Number of Channels used in LDAS:	32	A guess of the number of channels from Frames used by LDAS
Number of IFOs handled by LDAS:	1	
Physical Constants		
Nyquist Frequency of DAQS (Hz):	8192	
Total Data Supplied by DAQS (MB):	2	
Noise Model		
Seismic Cut Off Frequency (Hz):	30	Based on KB's Noise Model from Summer 97
Minimum in Noise PSD (Hz):	145	Based on KB's Noise Model from Summer 97
Physical Constants		
Gravitational Constant:	6.67E-11	MKS
Speed of Light:	3.00E+08	MKS
Mass of the Sun:	1.98E+30	MKS

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
π	3.14E+00	MKS
Binary Inspiral Source Model		
Each Component Mass (Msun):	1.2	
Percent Loss of Events Desired:	10.0%	One out of Ten events fail to fall close enough to a template
Maximum Frequency of Interest (Hz):	512	Based on Document from Sam Finn of Sept. 97
Inspiral Time for Binary (Seconds):	70.038	Based on 1PN
Number of Inspiral Templates per IFO:	34323	Based on Ben Owen's latest 2PN restricted waveform model
Total Number of Inspiral Templates:	34323	Factor up by the number of IFOs serviced by this system
Nearest 2 ^N Nyquist Frequency (Hz):	512	Closest power of two frequency to the Max Freq above
Black Hole Ring-Down Source Model		
Percent Loss of Events Desired:	10.0%	One out of Ten events fail to fall close enough to a template
Maximum Quality Factor:	89.5	Values up to 100 astrophysically expected, limited by CPU
Minimum Frequency of Interest (Hz):	30	Basically the Seismic cutoff
Maximum Frequency of Interest (Hz):	1024	Near the upper cutoff from astrophysics, higher expensive
Decibels Spanned by Ring-Down Template:	20	Two orders of magnitude range in the signal amplitude
Ring-Down Time (Seconds):	4.373	Based on decay of exponential through the decibels above
Number of Ring-Down Templates per IFO:	1852	Based on model from GRASP manual and J. Creighton
Total Number of Ring-Down Templates:	1852	Factor up by the number of IFOs serviced by this system
Nearest 2 ^N Nyquist Frequency (Hz):	1024	Closest power of two frequency to the Max Freq above

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Pre-Analysis Data Conditioning		
<u>Data Drop-Out Corrections</u>		
Percent of DAQS Frames with Drop-Outs:	0.10%	Only have to correct Frames with Drop-Outs
Data Samples in One Second Frame:	16384	
Number of Frames Resident in Memory:	3	Need before and after Frames to make best fill of data
Frame Data in Memory (MB):	6	
MFLOP to Calc pre RMS post of Channels:	4.2	Need to average all points in before and after Frames
MFLOP to Fill in Drop-Outs:	3.1	Need to fill in with RMS or straight line
Average MFLOPS Needed:	0.007	Accumulated FLOP only in the percentage of bad Frames
<u>Spectral Line Removal</u>		
Seconds of data Windowed per Pass:	4	Reasonable value for Multi-Taper Methods
Number of Lines Removed per Pass:	64	40 lines isolated and removed from 40-meter...slight increase
Number of Slepian's Used:	7	Must less than 2x number of windows in pass above
Samples in Each Data Segment:	8192	Only need Nyquest Frequency for Bandwidth used in Searches
MFLOP to Window Each Segment:	0.1	Each Slepian is multiplied by data to window
MFLOP to FFT each Windowed Segment:	37.3	Each windowed data set is FFT'ed
MFLOP to Set Phase of Cosine Vector:	0.0	The cosine vectors stored in circular buffer so phase is indexed
MFLOP to Set Amp of Cosine Vector:	0.5	Each cosine vector has a multiplication for the amplitude
MFLOP to Remove Lines from Segment:	2.1	Each cosine vector must be subtracted from the data
Total MFLOP Needed to Do Line Removal:	40.0	
Average MFLOPS Needed:	10.00	Average over the number of seconds of data
MB Memory Needed by Line Removal:	11.3	Memory needed to store data, cosine vectors and slepian's
<u>Calibration of GW Channel</u>		
Samples per second to Calibrate:	2048	Only calibrate the bandwidth relevant to searches

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Window Data (Time Domain):	0.004	Apply a time domain window to the GW data
MFLOP to FFT One Second of Data:	0.1	FFT the windowed GW data from time to freq domain
MFLOP to Mult by Complex Calibration:	0.016	Perform complex multiply with calibration transfer function
MFLOP to Inverse FFT Data:	0.1	Out of place IFFT to get time domain calibrated data
MFLOPS to Calibrate Data:	0.2	Total MFLOPS for this calibration: MFLOPS / one second
<u>Perform Data Reduction</u>		
Reduction for Inspiral:	16	Data sampling rate reduction for inspiral
MFLOP to Low Pass Inspiral Data:	0.786	Use 8(?) coefficient FIR for low pass filter
MFLOP to Average Down Inspiral Data:	0.068	Take average of filtered samples to reduce sample rate
Reduction for Ring-Down:	8	Data sampling rate reduction for ringdown
MFLOP to Low Pass Ring-Down Data:	0.786	Use 8(?) coefficient FIR for low pass filter
MFLOP to Average Down Ring-Down Data:	0.070	Take average of filtered samples to reduce sample rate
MFLOPS to Reduce Data Rates:	1.7	Total MFLOPS for reducing both inspiral and ringdown
MB Memory Needed by Data Reduction:	0.074	Total memory needed for performing the data reduction
<u>Perform Linear Regression</u>		
Samples per second to Regress:	2048	Only regress the bandwidth that will be used in searches
MFLOP to FFT Non-GW Channels	3.49	Perform FFT on other channels that are part of the regression
MFLOP to Calibrate Non-GW Channels:	0.016384	Perform complex multiplication of these with freq calibration
MFLOP to Regress GW Channel:	0.524288	Perform the single row complex multiply to regress GW channel
MFLOP to Inverse FFT Regressed Data:	0.11264	Inverse FFT to get time domain regressed data
MFLOPS to Regress GW Channel:	4.15	Total MFLOPS needed to perform this regression
MB Memory Needed by Linear Regression:	0.2578125	Estimated memory needed to perform this regression
Total MFLOPs Used in Data Conditioning:	16.10	Estimate of total CPU performance needed for data conditioning
Total MB Used in Data Conditioning:	17.626	Estimate of total Memory needed for data conditioning

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Binary Inspiral Data Analysis Model		
<u>Inspiral Data Model</u>		
Overlap Factor (power of 2):	8	Overlap is to apply template to more data per convolution
Max Template Size (power of 2):	1048576	The overall template size, (signal + 0 padding) should be 2^N
Zero Padding Added to Template:	976857	This is the about of zeros added to make 2^N and \sim overlap
Convolution Size (samples):	1048576	Number of samples (in time domain) for template/data
Size of Template (MB):	8	Size of Template in the frequency domain - templates not real!
Size of GW Data (MB):	8	Size of data in freq domain - made equal to template
Size of PSD (MB):	2	Size of Power Spectral Density of the Noise
New Data Analyzed Each Pass (sec):	953.962	Seconds of new data introduced into the convolution each pass
Percent of Data Analyzed Each Pass:	93.16%	Percentage of the convolution that is new each pass
Total Template Data Bank Size (GB):	268.1	Number of templates times Template Size
<u>Inspiral Node Model</u>		
Analysis Code Core Size:	2	Assume Analysis Code resides in 2 MB of core(RAM)
Amount of Temp Space (xGWdata):	7	Rough estimate of the about of temporary storage needed
RAM used by Data + Temp Space (MB):	64	Amount of RAM used by data and associated temp storage
RAM Available for Templates (MB):	32	Remaining RAM available on the system for holding templates
Number of Templates on Node Disk:	1073	Number of templates stored on disk on each node
Number of Templates in Node RAM:	4	Number of templates that fit in RAM at one time on each node
Template Disk Storage Needed: (GB):	8.38	Total Disk Storage needed on each node to hold its templates
Number of Template Swap-Outs Needed:	268.25	Number of times the templates in RAM must be swapped
<u>Once per Inspiral Pass Data I/O</u>		
Percent of Time CPU involved with I/O:	5.00%	Estimate of time percentage CPU used for I/O

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Load PS weighted Data on Node (sec):	0.667	By MPI rate, time spent loading preconditioned data on node
Template Swap-Out Time per Node (sec):	429.038	By SCSI rate, time spent loading templates on node
Total Time Moving Data (sec):	429.704	Total seconds of I/O each pass per node
Effective MFLOP lost to I/O:	6445.6	Lost CPU cycles to this I/O
<u>Computation per Inspiral Template</u>		
Number of values stored in Result Element:	3	The result is array of length number of samples by this number
Mult Template & weighted Data (MFLOP):	16.0	MFLOP used in complex multiplication of template with data
Convolve with Inverse FFT (MFLOP):	142.9	MFLOP for FFT based on $5N\log_2(N)$
Square Convolution (MFLOP):	8.0	Multiply complex convolution by its complex conjugate
Compare to Previous Convolution (MFLOP):	4.0	Do compare of each square in array with MAX stored for pass
Replace MAX info in Result (MFLOP):	3	Replace 3 elements in Max array when Max occurs
Computation for One Template (MFLOP):	173.9	Total Mega FLOP used per template on node
<u>All Inspiral Template Computations</u>		
Total MFLOP for All Templates on Node:	186548.7	Total MFLOP used for all template calculations
Total MFLOP when I/O included on Node:	192994.3	Add in the effective MFLOP count lost to I/O
Total Time Available for Computations (sec):	953.962	Amount of new data in seconds analyzed in this pass
Node CPU Performance Needed (MFLOPs):	202.31	Needed node performance of CPU in MFLOP per second
Ring-Down Data Analysis Model:		
<u>Ringdown Data Model</u>		
Overlap Factor (power of 2):	8	Overlap is to apply template to more data per convolution
Max Template Size (power of 2):	131072	The overall template size, (signal + 0 padding) should be 2^N
Zero Padding Added to Template:	122116	This is the about of zeros added to make 2^N and \sim overlap
Convolution Size (samples):	131072	Number of samples (in time domain) for template/data

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Size of Template (MB):	1.0	Size of Template in the frequency domain - templates not real!
Size of GW Data (MB):	1.0	Size of data in freq domain - made equal to template
Size of PSD (MB):	0.5	Size of Power Spectral Density of the Noise
New Data Analyzed Each Pass (sec):	59.627	Seconds of new data introduced into the convolution each pass
Percent of Data Analyzed Each Pass:	93.17%	Percentage of the convolution that is new each pass
Total Template Data Bank Size (GB):	1.809	Number of templates times Template Size
<u>Ringdown Node Model</u>		
Number of Nodes Reserved for Ringdown:	2	Number of Node of identical Class used for Ringdown
Analysis Code Core Size:	2	Assume Analysis Code resides in 2 MB of core(RAM)
Amount of Temp Space (xGWdata):	7	Rough estimate of the about of temporary storage needed
RAM used by Data + Temp Space (MB):	8	Amount of RAM used by data and associated temp storage
RAM Available for Templates (MB):	88	Remaining RAM available on the system for holding templates
Number of Templates on Node Disk:	926	Number of templates stored on disk on each node
Number of Templates in Node RAM:	88	Number of templates that fit in RAM at one time on each node
Template Disk Storage Needed: (GB):	0.90	Total Disk Storage needed on each node to hold its templates
Number of Template Swap-Outs Needed:	10.52	Number of times the templates in RAM must be swapped
<u>Once per Ringdown Pass Data I/O</u>		
Percent of Time CPU involved with I/O:	5.00%	Estimate of time percentage CPU used for I/O
Load PS weighted Data on Node (sec):	0.083	By MPI rate, time spent loading preconditioned data on node
Template Swap-Out Time per Node (sec):	46.3	By SCSI rate, time spent loading templates on node
Total Time Moving Data (sec):	46.383	Total seconds of I/O each pass per node
Effective MFLOP lost to I/O:	695.8	Lost CPU cycles to this I/O
<u>Computation per Ringdown Template</u>		
Number of values stored in Result Element:	3	The result is array of length number of samples by this number
Mult Template & weighted Data (MFLOP):	2.0	MFLOP used in complex multiplication of template with data

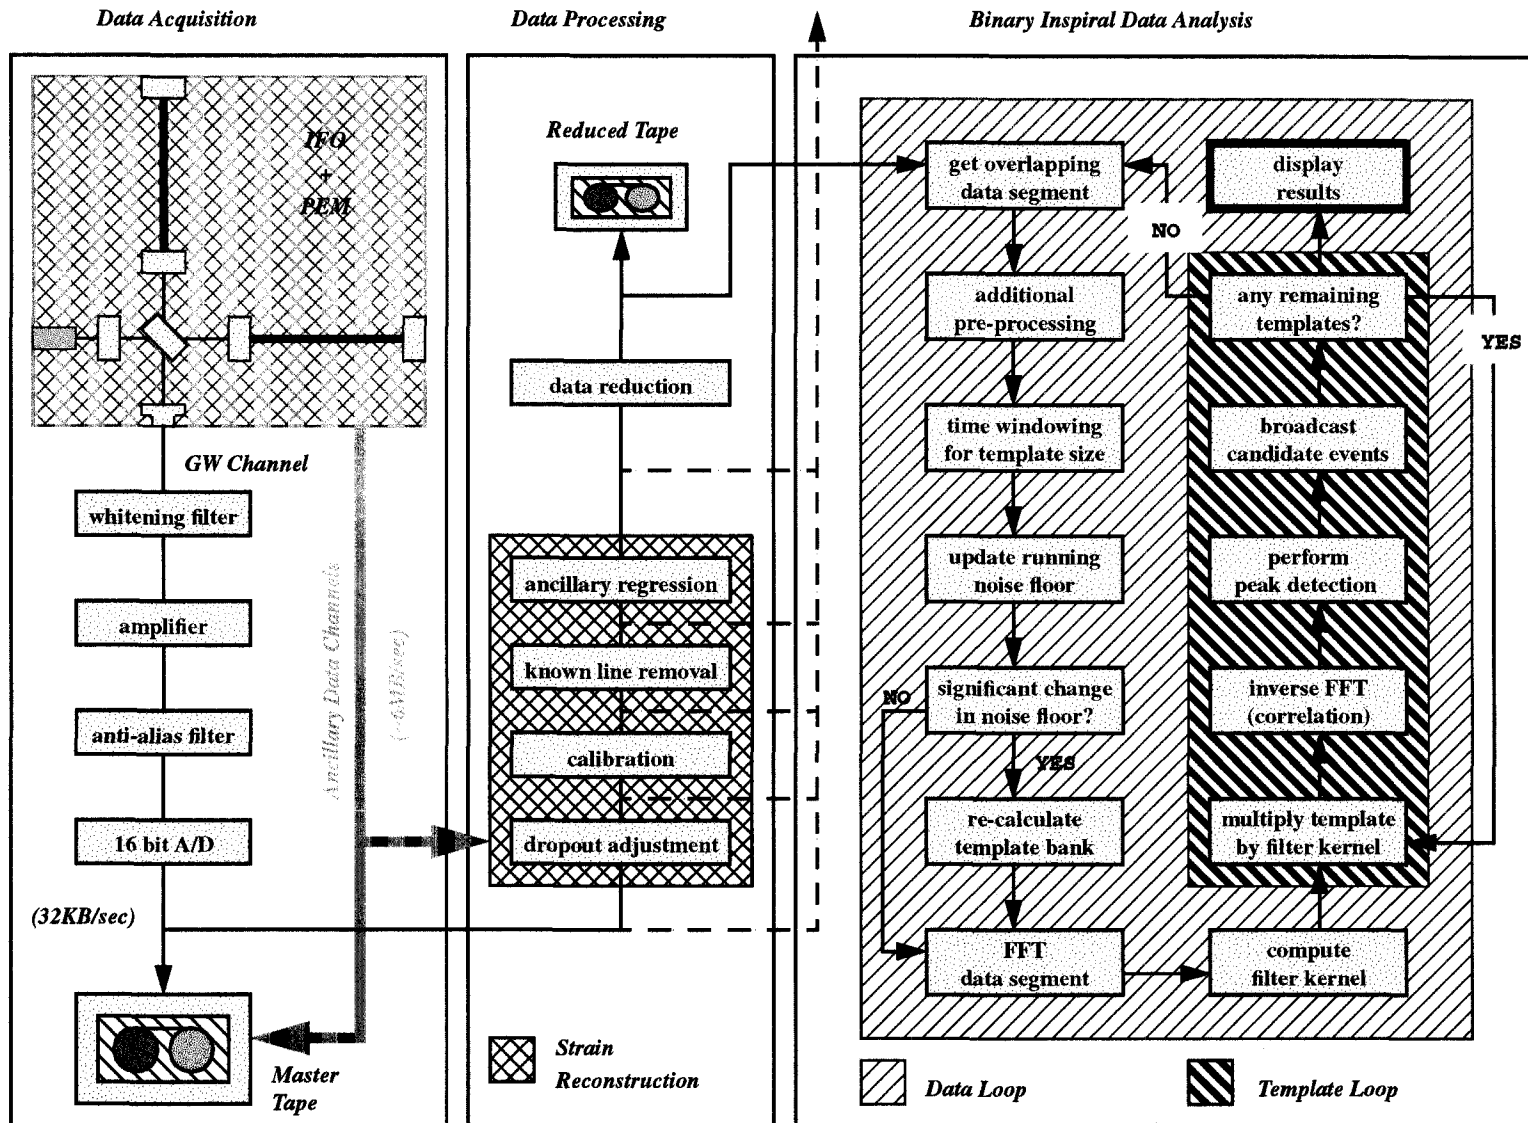
Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Convolve with Inverse FFT (MFLOP):	15.2	MFLOP for FFT based on $5N\log_2(N)$
Square Convolution (MFLOP):	0.5	Multiply complex convolution by its complex conjugate
Compare to Previous Convolution (MFLOP):	0.5	Do compare of each square in array with MAX stored for pass
Replace MAX info in Result (MFLOP):	0.375	Replace 3 elements in Max array when Max occurs
Computation for One Template (MFLOP):	18.6	Total Mega FLOP used per template on node
<u>All Ringdown Template Computations</u>		
Total MFLOP for All Templates on Node:	17180.6	Total MFLOP used for all template calculations
Total MFLOP when I/O included on Node:	17876.4	Add in the effective MFLOP count lost to I/O
Total Time Available for Computations (sec):	59.627	Amount of new data in seconds analyzed in this pass
Node CPU Performance Needed (MFLOPs):	299.8	Needed node performance of CPU in MFLOP per second
ROM Cost Estimates in 1997 US Dollars		
Dollars per MB of RAM:	\$20.00	Estimate of component cost (mid-range values)
Dollars per GB of Disk Storage:	\$200.00	Estimate of component cost (mid-range values)
Dollars per Node:	\$5,000.00	Estimate of component cost (mid-range values)
Dollars per 24 port 100bT Switch:	\$10,000.00	Estimate based on ATM/100bT Switch technology
Dollars per Node 100bT Interface PCI Card:	\$200.00	Estimate of component cost (mid-range values)
Total Number of Nodes:	35	
Total Number of 100bT PCI Cards:	35	
Total Number of Switches:	2	
Total Megabytes of Ram:	4480	
Total GB of Disk Storage:	315	
Cost of RAM:	\$89,600.00	
Cost of Disk Storage:	\$63,000.00	

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Cost of CPUs:	\$175,000.00	
Cost of Misc.:	\$27,000.00	
Total Cost of System:	\$354,600.00	This is per System; 3 will be needed
SUMMARY INFORMATION		
Minimal Mass Achieved in Inspiral:	1.2	
with a cut-off frequency in Hz of:	512	
leaving the following MB of node storage:	379.25	
leaving the following MB of node RAM:	0	Note: algorithm designed to use close to all memory
with the following node-core templates:	4	
utilizing the following percentage of CPU:	67.44%	
of which FFTs are:	79.43%	
Minimal Q Achieved in Ringdown:	89.5	This may be on the high side
with a lower frequency in Hz of:	30	
with a higher frequency in Hz of:	1024	This may be on the low side
leaving the following MB of node storage:	8034	Ringdown templates are short, leaving lots of disk space
leaving the following MB of node RAM:	0	Note: algorithm designed to use close to all memory
with the following node-core templates:	88	
utilizing the following percentage of CPU:	99.93%	
of which FFTs are:	78.63%	

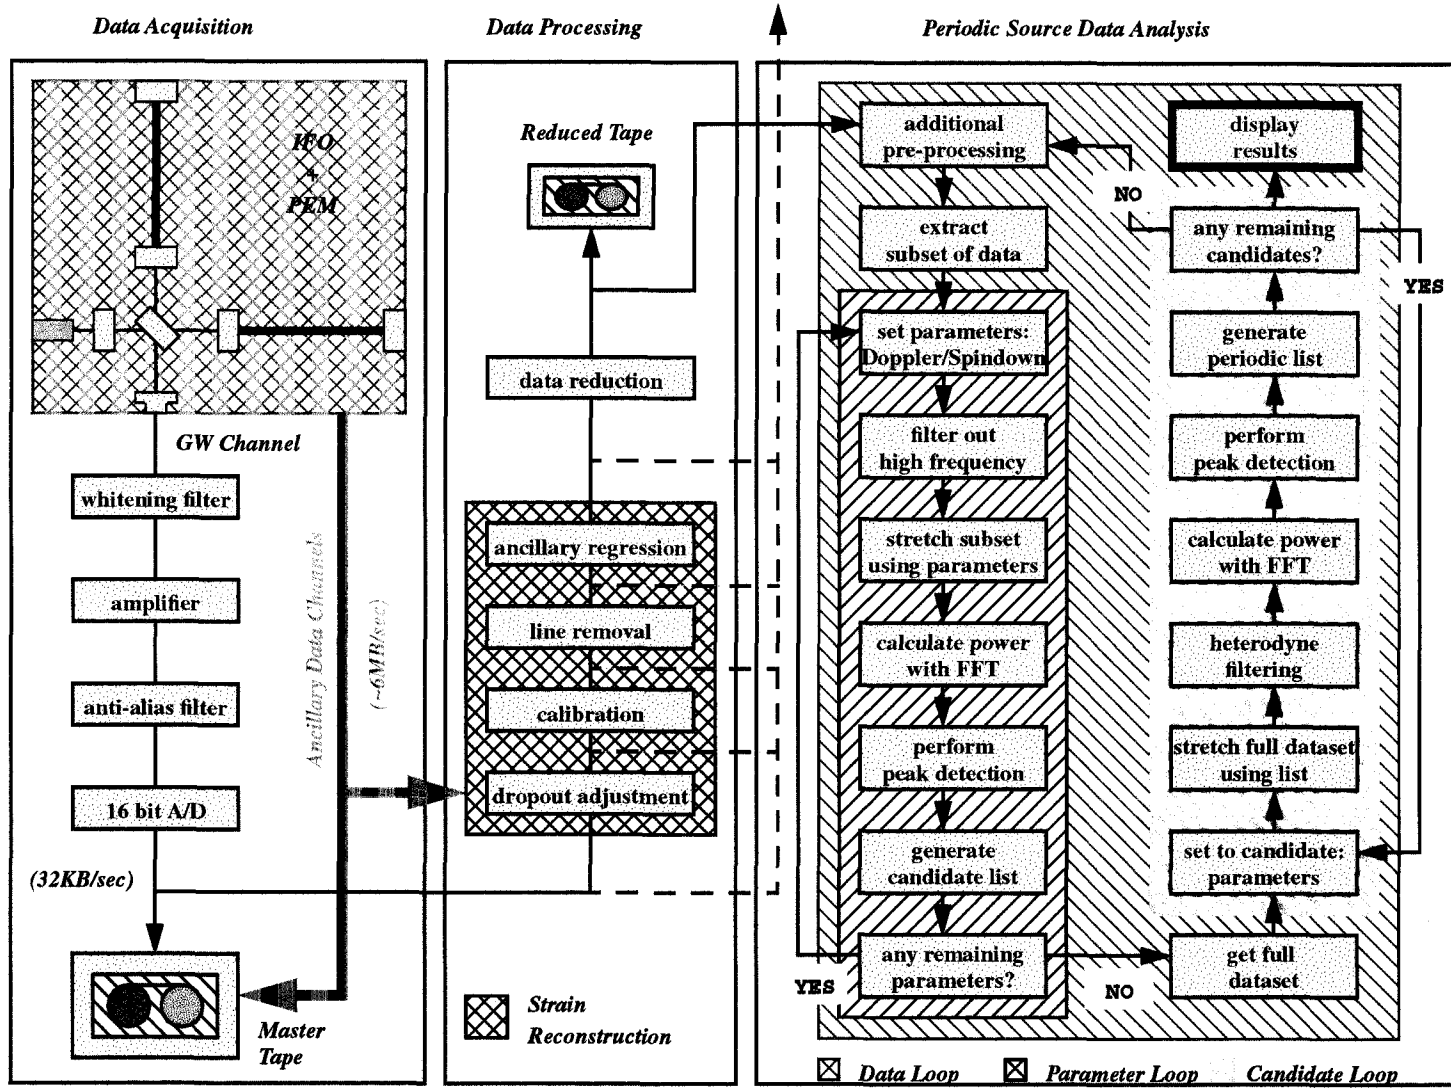
Figure 10 Data analysis flow for binary inspiral detection



APPENDIX B

**LIGO OFF-LINE DATA ANALYSIS
FLOW**

Figure 11 Data analysis flow for periodic source detection



APPENDIX C LIGO LAN ARCHITECTURE DETAILS FOR UNIVERSITIES AND OBSERVATORIES

Figure 12 Planned LIGO General Computing LAN at Caltech showing the link to the Caltech LDAS LAN and also to MIT and Sites

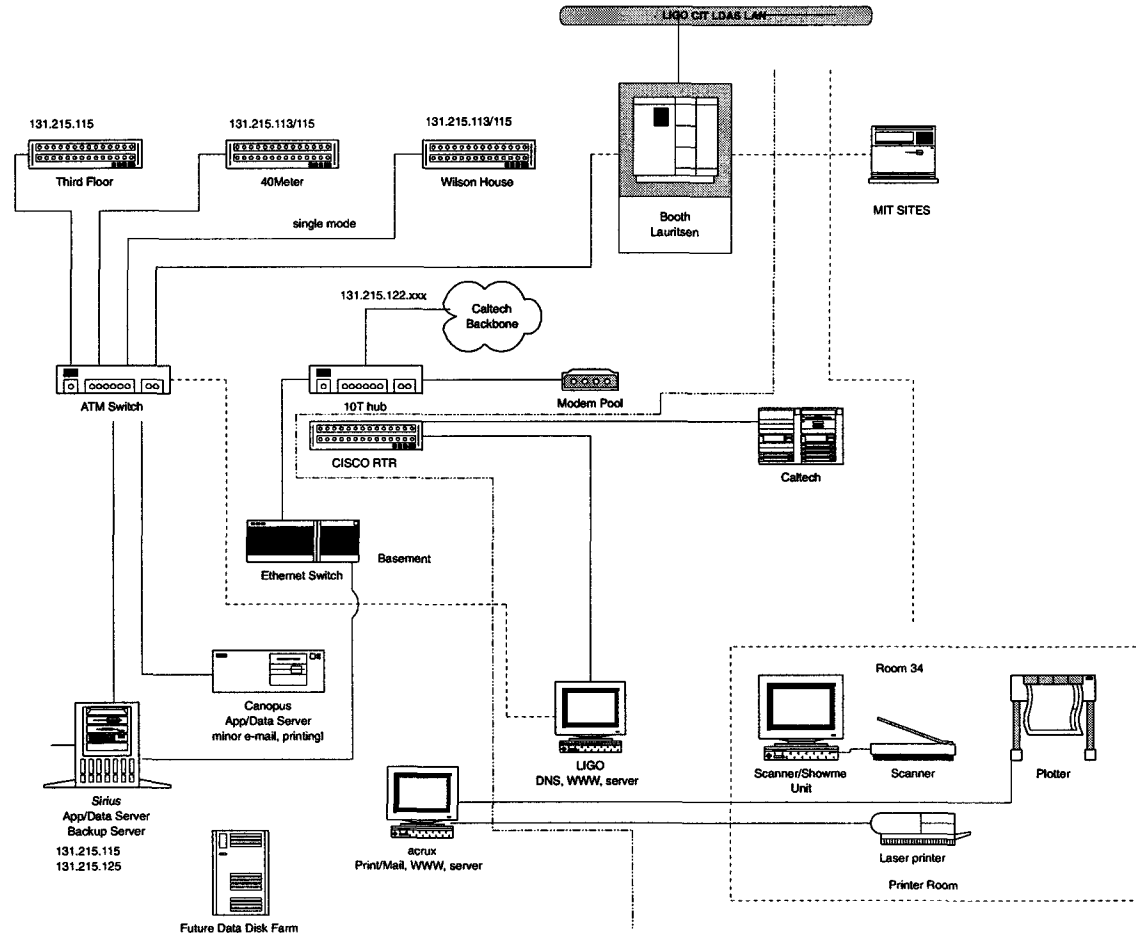


Figure 13 Planned LIGO General Computing LAN at MIT showing the link to Caltech and to the Sites (TBD)

Figure 14 Planned LIGO General Computing LAN at the Observatories showing the link to the Universities and local LDAS LAN

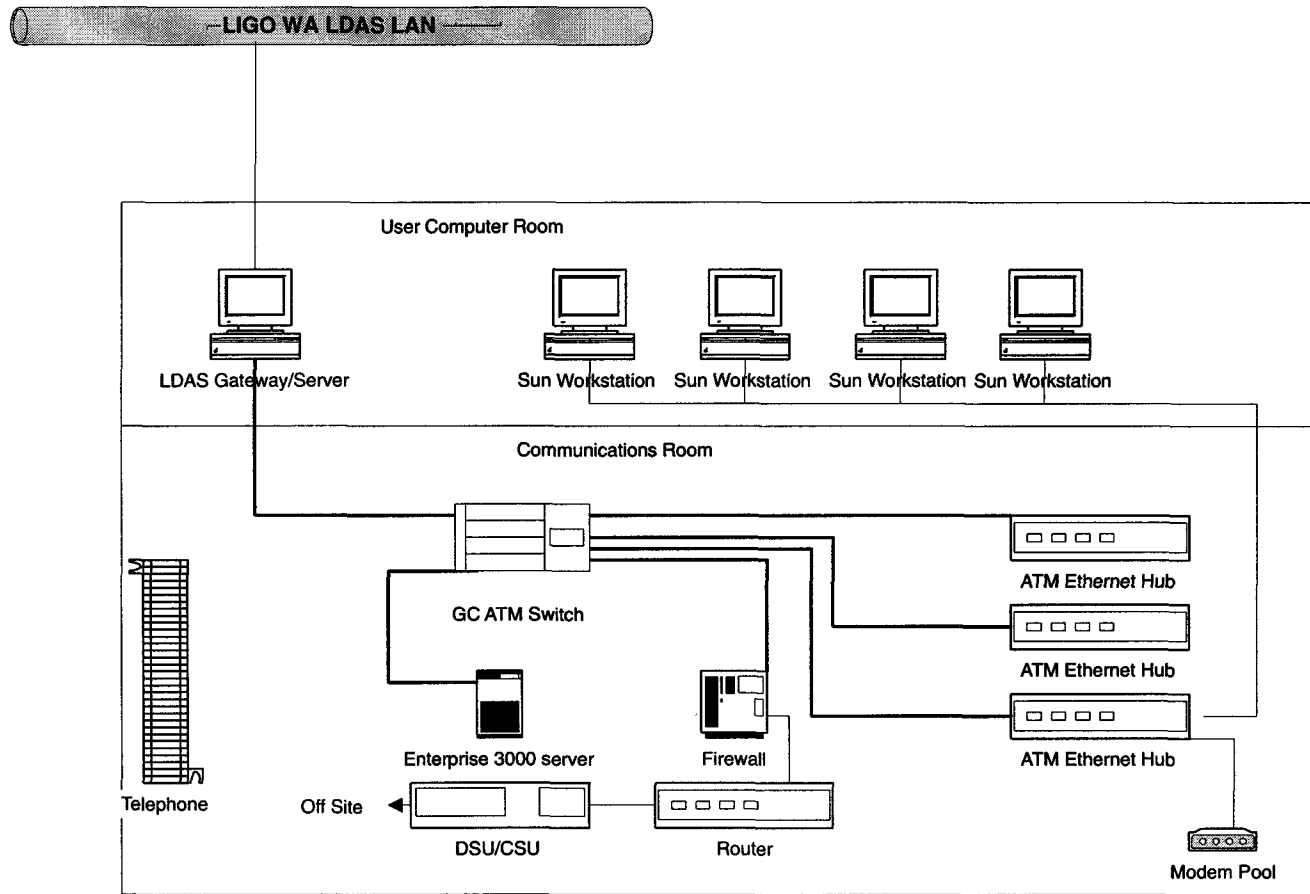
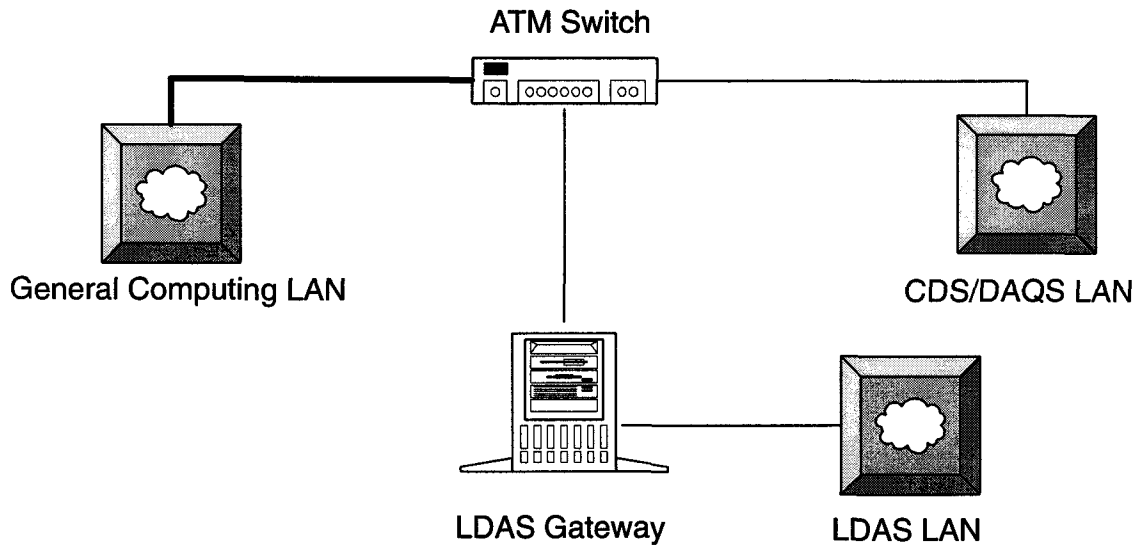


Figure 15 LAN Interfaces for the Observatories



BATCH
START

STAPLE
OR
DIVIDER

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -

CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type LIGO-T970160-04- E 10/14/97
LIGO Data Analysis System Conceptual Design.
LIGO Integration Group

Distribution of this draft:

**ELECTRONIC
COPY**

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project - MS 18-34
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: <http://www.ligo.caltech.edu/>

TABLE OF CONTENTS

1 Introduction	4
1.1. Purpose	4
1.2. Scope	4
1.3. Definitions	5
1.4. Acronyms	6
1.5. Applicable Documents	7
2 General description	8
2.1. Product Perspective	8
2.2. General Requirements	9
2.3. Functional Overview	9
3 Design.	12
3.1. System Architecture	12
3.2. On-line LDAS Data Analysis Flow	15
3.2.1. Data Distribution Unit (DDU)	15
3.2.2. Signal Conditioning and Preprocessing Unit (SCU)	16
3.2.3. Data Visualization, Control & Monitoring Server Unit (DVU)	16
3.2.4. Compute Server Unit (CSU)	17
3.2.5. LDAS Layout at the Observatories	17
3.3. Off-line LDAS Data Analysis Flow	18
3.3.1. Data Archival Unit (DAU)	19
3.3.2. Data Distribution Unit (DDU)	20
3.3.3. Data Ingestion Unit (DIU)	20
3.3.4. Signal Conditioning and Preprocessing Unit (SCU)	21
3.3.5. Control & Monitoring Server Unit (DVU)	21
3.3.6. Compute Server Unit (CSU)	22
5.1. Networking and data access	29
5.2. LIGO Laboratory	29
5.3. LIGO Scientific Collaboration	30
Appendix A LIGO On-line data analysis flow	33
Appendix B LIGO Off-line data analysis flow	46
Appendix C LIGO LAN Architecture Details for Universities and Observatories	48

LIST OF TABLES

Table 1	Applicable Documents	7
Table 2	Data Distribution Unit (DDU) -- On-line	15
Table 3	Signal Conditioning Unit (SCU) -- On-line	16
Table 4	Control & Monitoring Unit (DVU) -- On-line	16
Table 5	Compute Server Unit (CSU) -- On-line.....	17
Table 6	Data Archival Unit (DAU) -- Off-line.....	19
Table 7	Data Distribution Unit (DDU) -- Off-line	20
Table 8	Data Ingestion Unit (DIU) -- Off-line	21
Table 9	Signal Conditioning Unit (SCU)-- Off-line.....	21
Table 10	Control & Monitoring Unit (DVU)-- Off-line.....	22
Table 11	Compute Server Unit (CSU) -- Off-line.....	22
Table 12	WAN Connectivity among LIGO Laboratory Sites	30

LIST OF FIGURES

Figure 1	LDAS Functional Units. LAN or WAN refers to the local area networks either within a Laboratory site or between sites.....	8
Figure 2	LDAS Functional Block Diagram - On-line Analysis.....	11
Figure 3	LDAS Functional Block Diagram: Off-line Analysis	11
Figure 4	LDAS architecture for the on-line component.....	13
Figure 5	LDAS architecture for the off-line LIGO Laboratory University Sites (CACR configuration will change on a regular basis)1	4
Figure 6	Layout of on-line LDAS components within the OSB at the observatories.....	18
Figure 7	Block Diagram of the LDAS Software Components.....	28
Figure 8	NSF's vBNS Infrastructure will allow LIGO Livingston to link up with other Laboratory Sites via NCSA and Louisiana's regional access via SEPSCoR to vBNS (LIGO and SEPSCoR have been added to map; NSF map is from mid 1996)	31
Figure 9	DOE's ESnet Infrastructure will allow LIGO Hanford to link up with other Laboratory Sites via PNNL (LIGO has been added to map).....	32
Figure 10	Data analysis flow for binary inspiral detection	45
Figure 11	Data analysis flow for periodic source detection.....	47
Figure 12	Planned LIGO General Computing LAN at Caltech showing the link to the Caltech LDAS LAN and also to MIT and Sites.....	48
Figure 13	Planned LIGO General Computing LAN at MIT showing the link to Caltech and to the Sites (TBD)	48
Figure 14	Planned LIGO General Computing LAN at the Observatories showing the link to the Universities and local LDAS LAN	49
Figure 15	LAN Interfaces for the Observatories.....	50

1 INTRODUCTION

The LIGO observatory systems will generate large amounts of continuous data (approximately 12MB/sec for the Hanford site). At the observatories, these data must be processed in such a way that they may be distributed to various on-line analysis and diagnostic systems. At Caltech, the recorded data will be stored in the LIGO data repository or archive. From here, the data will be reduced for long term archival storage; they will also be available for distribution across the members of the LIGO I scientific collaboration for the purpose of supporting a variety of off-line research and searches. The off-line component will provide a level of parallel processing capacity provided through a collaborative agreement with Caltech's CACR. Data analysis and distribution will be possible via the LIGO wide area network (WAN) that will be available to interconnect LIGO Laboratory sites and sites of institutions who are members of the LIGO I Scientific Collaboration. These functions will be provided by the LIGO Data Analysis System (LDAS).

1.1. Purpose

This technical note presents the conceptual design for the LIGO Data Analysis System (LDAS). This design is in accord with the LIGO LDAS Design Requirements Document (DRD), LIGO T970159. Since it is early in the prototyping phase of data analysis, it is expected that new insights or experience will lead to an evolution of the requirements over time. This is particularly true for those components of LDAS that provide data reduction methods. It is the intent of this document to present a baseline design, based on the present requirements, which is flexible enough to incorporate new features in the future.

All designs presented in this document will identify technology available at the time of writing or known to become available within the design and development phase of the system. However, since the field of data analysis and signal processing is rapidly expanding as new technologies become available, this will open the design to better, faster, cheaper options in the future. Therefore, where possible, flexibility has been included in the design to allow future developments to be incorporated wherever they are applicable. While specific hardware components are shown in the baseline design presented in this document, these will not necessarily be the ones implemented.

1.2. Scope

The LDAS is to be developed to meet the requirements set forth in this document. The LDAS shall have the means to:

[A] Provide on-line data analysis capability to each of the LIGO Observatories. This capability includes the following:

- A means to extract physical strain from the interferometer output(s) and to utilize relevant ancillary channels (e.g., PEM) to remove instrumental or environmental signatures.
- A means to process strain data through real-time detection algorithms for both performance monitoring and scientific purposes. Sufficient computing power to allow processes to keep up with the incoming datastream shall be provided. Sufficient margin shall also be provided to accommodate maintenance down times and other system inefficiencies.
- A means to cross-correlate data (either time series or event lists) from multiple interferom-

eters.

- A means to store data frames and analysis results (local to the observatory) to short term storage media. This functionality will be provided by the LIGO DAQS resources, with possible augmentation by LDAS.
- A means to access both “live” and short term archived data via the local area and wide area networks. Access shall be subject to available bandwidth and demand.
- Means to retrieve, concatenate and extract specific channels of recent data from the on-line storage system.
- Sufficient automation to run continuously and autonomously during periods of normal operation.
- A means to display and visualize results of analyses.

[B] Provide off-line data analysis capability. This capability is likely to be concentrated at one LIGO Laboratory site but shall be available “seamlessly” throughout the Laboratory. This capability includes:

- A means to reduce the raw data to science data representing calibrated GW strain data and a reduced subset of ancillary data and a data quality descriptor.
- A means to archive, retrieve and distribute reduced datasets acquired over a period of time at least 5 years in duration.
- A means for duplicating reduced datasets either for backup or for distribution.
- A means to access the data archive via LAN and WAN by the LIGO Laboratory and LIGO Scientific Collaboration with sufficient bandwidth to support database manipulation at the off-line site by remote users.
- A standardized interface for visualization tools to allow experimenters to see and interpret results from various analyses.
- Sufficient computing margin to enable multiple analyses to be conducted in parallel.

Specifically not considered to be within the scope of the LDAS are:

- Data analysis functions performed at centers other than the LIGO Laboratory Facilities.
- The on-line diagnostics system used for stimulus-response characterization, transfer function determination, and calibration functions. However, it is expected that software developed for the LDAS will find utility within the diagnostics system and vice-versa.
- Simulation capability shall be provided separately from, but coordinated with, the LDAS. The simulation environment is being developed also within LIGO.

1.3. Definitions

On-Line: Data are considered to be “on-line” if they are readily available to clients via the LDAS networks from short term storage.

On-Line Analysis System: The system which processes the data stream in real-time for performance monitoring and for the detection of astrophysical events which are time-critical. There shall be two similarly configured systems (although not necessarily performing identical operations at the same time or having the same capacities) at both LIGO Observatories.

Off-line Analysis System: The system which is used to reduce, archive, retrieve, analyze and duplicate datasets after these have been collected and transmitted to the data repository site. The system shall also provide sufficient computational capacity to members of the LIGO Scientific Collaboration.

Real-time data: The LIGO datastream written to disk cache by the CDS DAQS. The actual latency in the data is determined by the length of frames and access time to read these off the media.

Real-time analysis or processing: Analysis of data so that the information contained in them can be extracted on a time scale sufficiently short so that it is possible to influence either the improvement or maintenance of detector operational performance or the collateral detection of a potential astrophysical event by other detector systems (outside of LIGO).

Time-critical:

[i]Scientific: having the potential to influence the operation of other (non-LIGO) astrophysical or astronomical detection systems so that these instruments may be employed to observe the same phenomenon detected by LIGO (or vice-versa). The time scale for *time criticality* will range from fractions of an hour to fractions of a day or even longer.

ii]Operations: having the potential to allow recovery of a LIGO interferometer from off-nominal operation by virtue of the information extracted from the LIGO datastream.

1.4. Acronyms

- CACR Center for Advanced Computing Research (Caltech)
- CDS Control and Data System
- CSU Compute Server Unit
- DAQS Data AcQuisition System
- DDU Diagnostic Distribution Unit
- DIU Data ingestion Unit
- DRU Data Reduction Unit
- DVU Data Visualization Unit
- GUI Graphical User Interface
- IFO Interferometer
- LAN Local Area Network
- LDAS Data Analysis System
- LIGO Laser Interferometer Gravitational Wave Observatory
- MPI Message Passing Interface
- MTB(C)F Mean Time Before (Critical) Failure
- MTTR Mean Time To Repair
- NFS Network File Services
- OSB Operations & Support Building (at Observatories)
- PEM Physical Environment Monitoring (System)
- RH Relative Humidity
- SCU Signal Conditioning Unit

- SNR Signal to Noise Ratio
- TBD To Be Determined
- WAN Wide Area Network

1.5. Applicable Documents

The documents listed in Table 1 contain relevant design and specification data for the LDAS

Table 1 Applicable Documents

<i>Document Identifier</i>	<i>Description</i>	<i>Comments</i>
M970065	White Paper Outlining the LDAS for LIGO I	
T970140	LIGO LDAS Software Specification and Design Requirements	In process
T970100	LIGO System Software Design Issues	
	Paraflo User's Guide	Author: R. Williams, CACR
T970078	Interferometer Diagnostics Tests and Tools	Draft
T960009	CDS Data Acquisition System Design Requirements Document	
T950054	CDS Control and Monitoring Design Requirements Document	
T960010	CDS Data Acquisition System Conceptual Design	
T960108	Interferometer Diagnostics Conceptual Design	
T960107	LIGO Interferometer Diagnostics System Design Requirements	
VIRGO-SPE-LAP-5400-103	Frame Library Users Manual	
T970130	Specification of a Common Data Frame Format for Interferometric Gravitational Wave Detectors (IGWD)	
T970159	LIGO Data Analysis System (LDAS) Design Requirements	

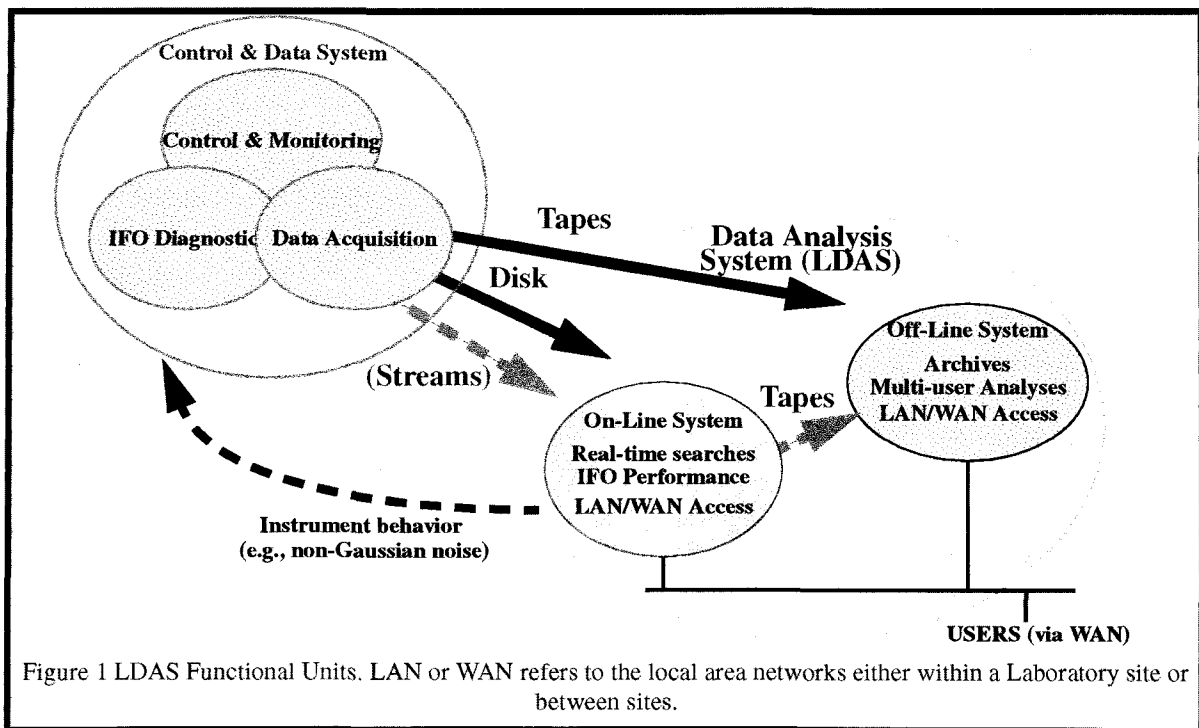
2 GENERAL DESCRIPTION

2.1. Product Perspective

The LDAS is divided into two primary functional units as shown in Figure 1.

These are:

- **On-line LDAS:** The on-line segment consists of two functionally identical but independent units located at the LIGO Observatories (Hanford, WA and Livingston, LA). Each provides the capability to run real-time detection algorithm, and also the ability to provide end-to-end insight into interferometer behavior for specific signal types. The system interfaces to the LIGO DAQS for accessing the real-time data. It also has a limited one-way interface (data displays for operators) for interferometer diagnostics to provide performance metrics (e.g., based on non-Gaussian noise characteristics). The site LAN may be used to access the data cache by scientific analysis workstations at the site. There will also be the ability to access the real-time data remotely by other LIGO Laboratory sites.
- **Off-line LDAS:** The off-line component will likely reside at Caltech's CACR and provides several functions: data reduction and compression for long term archival; data retrieval; and multi-user (independent) refined analysis of LIGO data. Access to the archive will be via a wide area network (WAN) capable of providing high throughput access to the archive. The baseline provides for regular transfer of data tapes from the LIGO Observatories to the off-line analysis center. Once data tapes are received, they will be processed in a yet to be defined manner to extract/compress/refine the science data for the permanent archive.



The on-line and off-line systems shall be independent; however, where it is required, those critical databases or components will be “mirrored” between them. The systems will be highly integrated so that, for example, a remote user will be able to access either the on-line or off-line system with nearly identical user interfaces and commands.

2.2. General Requirements

The specific requirements which this system must meet are given in the DRD, LIGO-T970159. The primary requirement on the system, which drives the design which follows, is to be able to:

- Provide on-line analysis at the observatories.
- Process and reduce the raw LIGO datasets at the off-line center to prepare the data for archival storage and retrieval.
- Provide computational and storage resources for off-line analysis using the archived data
- Provide a flexible design which can be reconfigured to reflect new analysis or computational requirements as they evolve.
- Provide access to LIGO data from all LIGO Laboratory sites and also from member institutions of the LIGO Scientific Collaboration for the LIGO I search.

2.3. Functional Overview

Figure 2 depicts the functionality provided by the on-line LDAS. The on-line LDAS processes the data stream provided by the CDS/DAQS at a rate sufficient to keep up with the acquisition rates. Processing involves data distribution to remote users qualified for access to the datastream. It also involves performing signal conditioning operations on the GW strain channel to calibrate it and to improve the SNR, if feasible, by regressing out instrumental or environmental signatures.

As a minimum, the on-line LDAS will also process the GW (strain) channel to look for one or more astrophysical source types which, if detected, may be correlated with other (non-LIGO) astrophysical detectors. These detection algorithms will be run by one or more processes running at real-time rates to filter the datastream for source signatures. A complementary product of such analyses will be data characterizing the instrument performance in terms of false alarms rates, noise floor sensitivity and spectral variations, etc. Users will be able to access such event and health/status information using X-windows and web technology based interfaces into the LDAS.

The on-line LDAS will have its own local area network (LAN) which is interfaced with the observatory general computing LAN and is also interfaced to the LIGO wide area network (WAN) interconnecting LIGO Laboratory Sites.

The LDAS also has an off-line component resident at Caltech (most likely CACR) which services the entire Laboratory and also the collaboration for the LIGO I searches. Figure 3 shows this component's functional elements. This component serves several functions. It provides access to additional computational resources not feasible to be provided at the observatories. These resources may be utilized and configured to perform more refined analyses than may be possible or warranted at the two sites. It is also here that the full correlated use of data from multiple interferometers will likely be exploited. The computational resources will include dedicated LIGO resources and, on a much larger scale, institutionally provided resources which are available a significant fraction of the time.

The off-line LDAS also provides for the management and creation of the LIGO data archive. Data reduction and compression will be utilized to reduce the data to a volume suitable for commitment to long-term storage and retrieval. Data distribution is provided by the archive both for the (local) types of processing discussed above and also for direct distribution to other Laboratory or Collaboration sites where additional data analysis make take place.

Data ingestion provides for incorporating into the archive new data delivered on media from the observatories. Tape media from the observatories are not expected to be directly compatible with the archive media, hence a translation will need to be performed.

Note that the figures provide functional views generated by the DRD, and they do not translate into distinct and separate hardware (or software) systems. The LDAS is designed such that, to the greatest extent possible, the coupling between the functional units is weak. Functional units will couple over single, well defined interfaces. This will provide a flexible system which may be reconfigured with few problems (e.g. porting a function onto new hardware, moving a functional block from one computer to another).

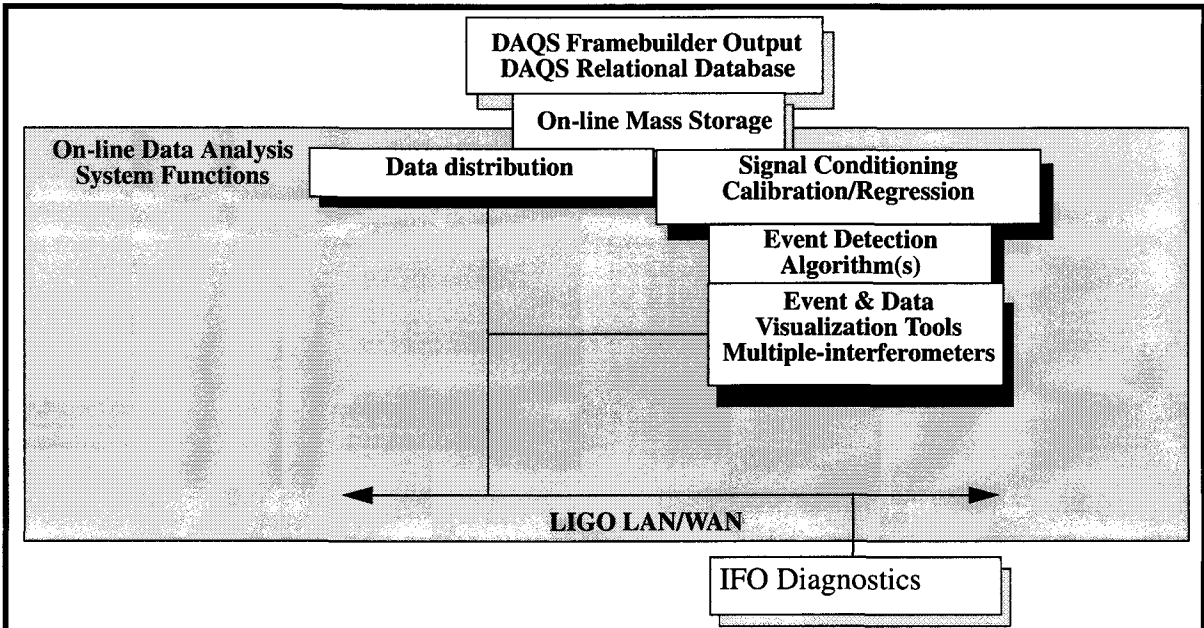


Figure 2 LDAS Functional Block Diagram - On-line Analysis

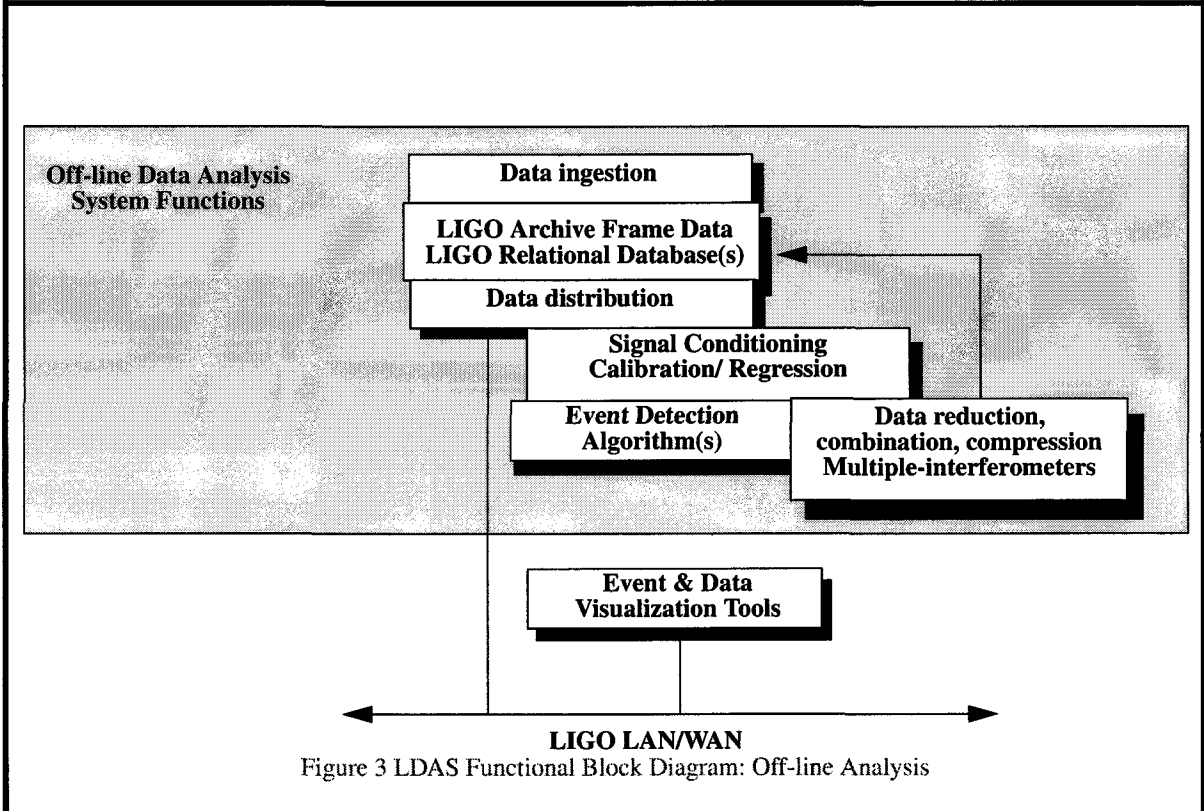


Figure 3 LDAS Functional Block Diagram: Off-line Analysis

3 DESIGN.

3.1. System Architecture

Figure 4 depicts an overview of the LDAS architecture for the on-line component at both sites. The Livingston site will require half the processing and storage capacity because it accommodates only one interferometer.

At each site there will be a dedicated OC3 or 155 Mb/s LAN (nominally ATM) to support on-line data analysis. This LAN is linked via a router to the Observatory General Computing (GC) LAN which is similar in design. Both the GC LAN and LDAS LAN are linked via a gateway to the LIGO WAN. It is planned that the LIGO WAN at Hanford will be able to use DOE's ESnet services. There will be a gateway/router configured to transfer LIGO packets off ESnet and onto vBNS either at SDSC or at Caltech; it is expected that at a future date access to LIGO Hanford via ESnet will also be possible directly at MIT. LIGO Livingston is planned to be accessible via NSF's vBNS using LSU as a gateway.

There will be a limited interconnection from the LDAS LAN to the CDS LAN. A gateway will provide limited TBD services. The intent of the isolation is to limit the loads introduced by data distribution and analysis so that CDS is not affected and so that the GC LAN is also not affected.

LDAS units are interconnected either through Fast Wide SCSI (F/W SCSI) ports to the on-line mass storage system or via LAN to servers and processors in the LDAS.

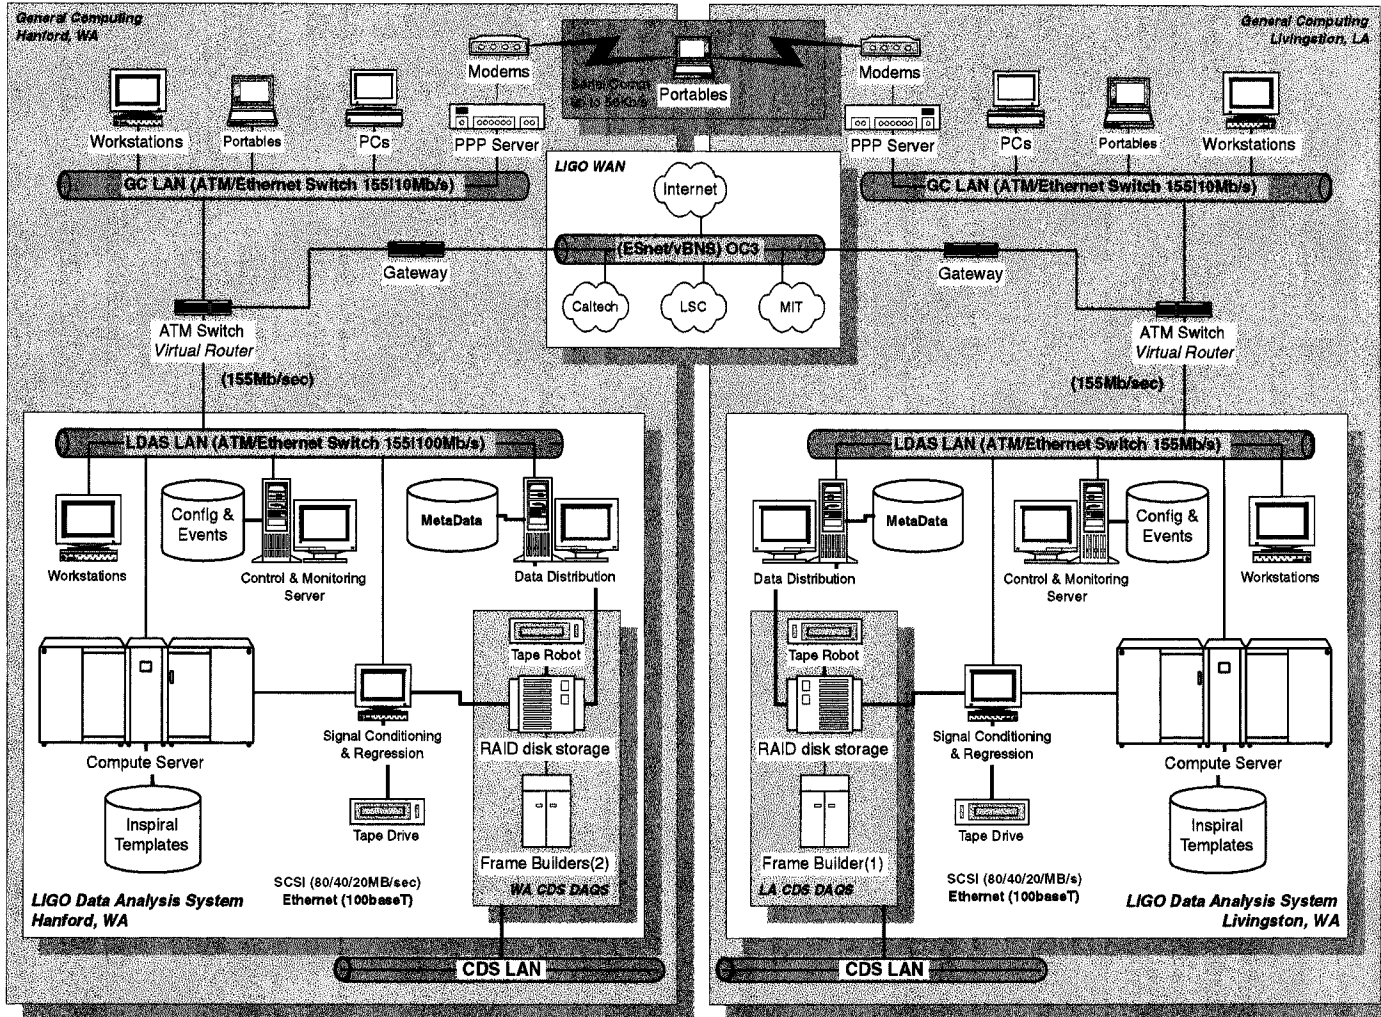
Figure 5 shows an overview of the LDAS architecture for the off-line component. This component is principally implemented at Caltech/CACR but has an independent component at MIT.

At CACR the network dedicated to intensive data transfer and computation will be the extant CACR HiPPI networking infrastructure. The dedicated LIGO resources within CACR will be linked into the HiPPI in an identical manner to the rest of CACR's resources. In this manner, it will be possible to share resources within CACR transparently for parallelized computation based on MPI. As additional resources become available, the LDAS will be able to schedule access to them in order to provide the optimal quality of service consistent with available resources. The CACR LDAS will be accessible on the local LIGO LAN at Caltech using ATM as an interface. much of this interface already exists in a prototyped environment.

The LDAS component at CACR is similar in function to the on-line component at the observatories. The major difference is in the scale of computational resources that are implemented, the hardware dedicated to data ingestion and to data reduction for archival storage, and the archival storage system itself. Moreover, these resources are accessible transparently to all LIGO Laboratory and Collaboration sites via the LIGO WAN and internet in general.

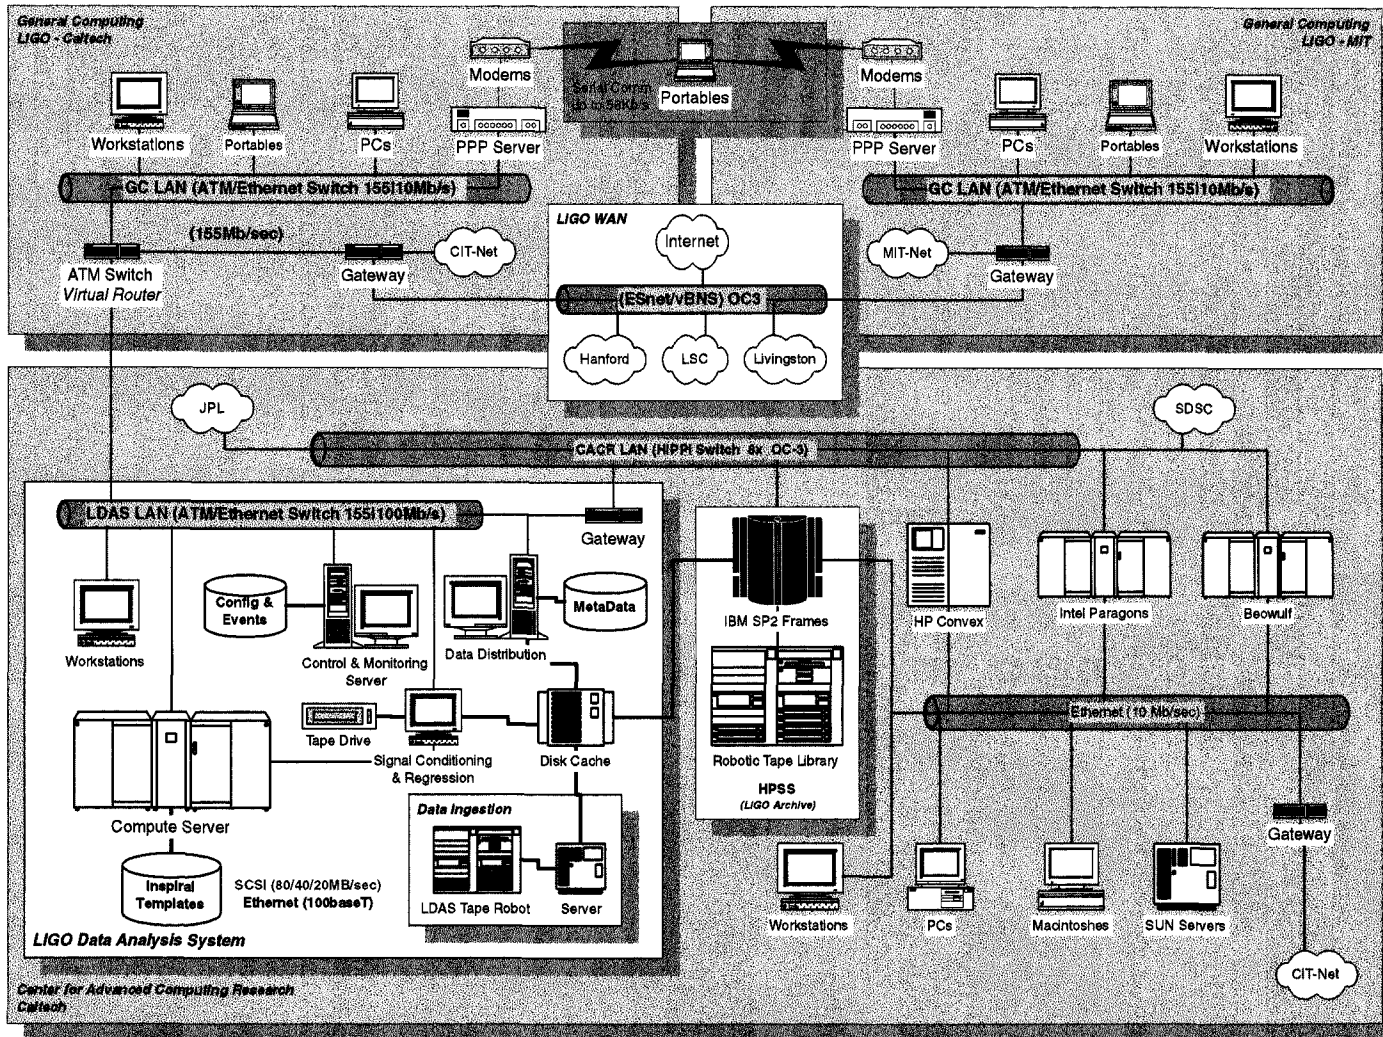
MIT will have a local LDAS component consisting of high end unix workstations for local MPI-based computation (on a more limited scale) and also consisting of a mass storage system sufficient to contain approximately 600 GB of data at any given time. This allows intensive analysis to be done locally when internet-based access to LDAS at CACR proves inadequate.

Figure 4 LDAS architecture for the on-line component



LIGO-T970160-04

Figure 5 LDAS architecture for the off-line LIGO Laboratory University Sites (CACR configuration will change on a regular basis)



3.2. On-line LDAS Data Analysis Flow

Figure 10 and the preceding discussion in Appendix A describe how data will flow through the analysis process for the specific search for binary inspiraling chirped waveforms. Data are served from the on-line data mass storage system by a data distribution system. Because most presently envisioned analysis algorithms for other sources implement similar optimal filtering techniques, these analyses will also follow a similar flow.

A set of preprocessing steps take the raw data and condition the gravitational wave channel by dropout correction, suitable calibration, unwhitening, regression with other channels, and transformation into the Fourier domain in preparation to hand off the particular epoch of data to a massively parallel processing system designed to perform optimal Wiener filtering for the detection algorithm.

The following discussions rely on details of the data analysis flow which are presented in Appendix A and from which follow the hardware specifications presented in the tables which follow.

3.2.1. Data Distribution Unit (DDU)

The DDU the front end component of the on-line LDAS. Refer to Figure 4 in the discussion which follows. The DDU sits at the interface between DAQS and LDAS. The primary function of the DDU is to make data available to remote users at other LIGO Laboratory or Collaboration sites. The DDU is connected to the LDAS ATM LAN and through this LAN to the site LAN via an ATM switch which establishes a virtual network for the Observatory. Access to the GC LAN or CDS LAN is via gateways which serve to limit ethernet traffic on the respective networks. The GC LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The DDU will serve frame data to clients situated either locally, on the CDS or GC LANs). The specifications for the DDU are listed in Table 2.

Note data distribution to the on-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the on-line processes will not be affected by client loads on data distribution and also so that the on-line process may be isolated from inadvertent interference by off-site remote users.

Table 2 Data Distribution Unit (DDU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
RAID System (shared with CDS/DAQS)	-	-	-	>500 GB striped(TBD)
Ultra Wide SCSI Port	-	-	> 40 MB/s	-
Data Server	>200 MFLOPS	>256 MB RAM	-	-
Network	-	-	- LIGO LDAS LAN (100BT/100Mb/s) - LIGO Site LAN (OC3/155Mb/s)	-

3.2.2. Signal Conditioning and Preprocessing Unit (SCU)

Referring once again to Figure 4, the SCU has four interfaces. The input stream is taken from the DDU on a dedicated high speed F/W SCSI port. Refer to Table 3 for the specifications.

Table 3 Signal Conditioning Unit (SCU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS	>128 MB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 2 ports	-
Tape Drive/Robot	-	-	>2.5 MB/s >5 tape storage	>25 GB/tape
Network	-	-	- CSU direct (≥ 100 BT) - LIGO LDAS LAN (OC3/155Mb/s)	-

3.2.3. Data Visualization, Control & Monitoring Server Unit (DVU)

Referring to Figure 4, the DVU provides the user API to configure the on-line LDAS, monitor its health and status, and monitor event lists and other performance metrics derived from on-line processes which are useful for interferometer performance assessment. It consists of one or more high-performance work stations, monitors, and dedicated disk. Access to this system will be provided in the control room in the vicinity of the detector operational monitors so that console operators may have access to scientific information relevant to maintaining acceptable interferometer sensitivity.

The DVU may also be the web server for the LIGO Metadatabase. Refer to Table 4 for the specifications.

Table 4 Control & Monitoring Unit (DVU) -- On-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS /CPU ≥ 2 CPU	512 MB/ CPU	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Disk	-	-	-	> 50 GB (TBD)
Network	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

3.2.4. Compute Server Unit (CSU)

The CSU provides the computational capacity to process in parallel many optimal filter templates. It receives input data directly from the SCU and provides event lists and other reduced data via the LDAS LAN. Results will be post-processed, displayed, or otherwise manipulated separately using DVU components or other local or remote hardware. Refer to Table 5 for its specifications.

Table 5 Compute Server Unit (CSU) -- On-line

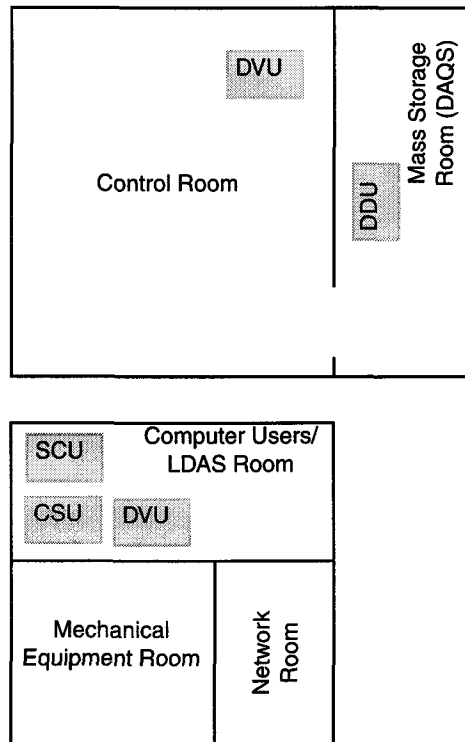
<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	32 nodes, ≥ 10 GFLOPS in aggregate	Per node, ≥ 128 MB RAM	100BT Ethernet	per node, 10GB disk
Fast Wide SCSI Ports	-	-	> 20 MB/s 1 port per node	-
Network	-	-	- LIGO LDAS LAN (OC3/155Mb/s) switched, point-to-point	-

3.2.5. LDAS Layout at the Observatories

Figure 6 depicts a schematic layout of where the principal components of LDAS will be located within the OSB.

The DDU resides in the mass Data Storage Area, in the vicinity of the CDS DAQS because it shares the RAID system with it. The SCU will be located in the Computer Users/LDAS Room. Similarly, the CSU will be located in this area. The DVU, of which there will be several, will be distributed between the Control Room and the Computer Users/LDAS Room. It is anticipated that there will be a need for several DVUs to support operations/diagnostics research and also on-line detection event analysis.

Figure 6 Layout of on-line LDAS components within the OSB at the observatories



3.3. Off-line LDAS Data Analysis Flow

Figure 11 and the accompanying discussion in Appendix B describes schematically how data will flow through the analysis process for the specific search for periodic source waveforms. This is considered at present the most demanding analysis task for which LIGO must prepare to perform off-line. There are several variants on the search task: directed searches looking for expected source waveforms (including spin down); directed searches looking for unknown waveforms; and large-area searches looking for unknown waveforms. The first two problems may also be implemented in part on line for performance monitoring purposes. The last of these problems is presently intractable with computing technology that is projected to be available in the next several years. The limited periodic search will serve as a paradigm for a data analysis flow that must be accommodated by the off-line LDAS component. Of course, the off-line LDAS will also be used to perform refined or deeper (lower chirp mass) searches for inspiraling binary waveforms; this flow parallels that already discussed in Section 3.2.

Data are served from the off-line data archive by a system equivalent to a High Performance Storage System (HPSS), consisting of a tape robot and large volume disk cache system which function as the LIGO Data Archive. The disk caching system will be a partition on a shared disk RAID system also serving the data ingestion and reduction processes.

A set of preprocessing steps take the raw data and condition the gravitational wave channel by dropout correction, suitable calibration, unwhitening, regression with other channels, and trans-

formation into the Fourier domain in preparation to hand off the particular epoch of data to a massively parallel processing system designed to perform optimal Wiener filtering for the detection algorithm.

A significant role of the off-line LDAS will be to reduce and archive the raw data arriving from the Observatories. The data flow for reduction will be similar to the signal conditioning process, except that the end product is written back into frame format, possibly compressed, cataloged, and then transferred to the archive for long term storage and retrieval.

The following discussions rely on details of the data analysis flow which are presented in Appendix B and from which follow the hardware specifications presented in the tables which follow.

3.3.1. Data Archival Unit (DAU)

The DAU is the front end component of the off-line LDAS. Refer to Figure 5 in the discussion which follows. The primary function of the DAU is to store and retrieve data from long-term storage. Data are then served via the Data Distribution Unit (DDU, see below) to users located at any of the LIGO Laboratory or Collaboration sites. The DAU is connected to the LDAS ATM LAN and also to the CACR HiPPI network so that the data may be accessed by all CACR computational resources. Access to the LIGO ATM LAN is via gateways which serve to limit ethernet traffic on the respective networks. The LIGO ATM LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The specifications for the DDU are listed in Table 6.

Note data distribution to the off-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the off-line processes will not be affected by client loads on data distribution and also so that the off-line process may be isolated from inadvertent interference by off-site remote users

Table 6 Data Archival Unit (DAU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Tape Robot & Cabinets			≥ 5 Tape heads 6 MB/s	≥ 100 TB 3 cabinet
Disk System (Shared with DDU)	-	-	HiPPI	>250 GB striped for 5 users =2 * #users * vol/tape (25 GB) Partition on shared resource with DDU/DIU
Data Server/IBM SP2	2 node, >200 MFLOPS per node	≥ 512 MB RAM	HiPPI	-
Network	-	-	- CACR LDAS LAN - CACR HiPPI	-

3.3.2. Data Distribution Unit (DDU)

The LIGO Data Distribution Unit comprises the front end component of the off-line LDAS for users of archived data. Refer to Figure 5 in the discussion which follows. The DDU is comprised of a tape archive residing in cabinets accessible by tape robots; it consists of tape readers which dump frame-based data onto a large volume disk cache; it consists of a data server capable of sustained I/O rates sufficient to support up to 5 users at a time each of whom is capable of accessing data at rates up to 6 MB/s. More users may be supported, but a corresponding degradation of I/O throughput to each user will occur.

The primary function of the DDU is to provide access to the LIGO archive to users situated at all LIGO Laboratory or Collaboration sites. The HPSS is connected to the LDAS via the HiPPI backbone which forms part of the CACR network infrastructure. The HiPPI interfaces to the LIGO site LAN via an ATM switch which establishes a virtual network for the Laboratory. Access to the GC LAN is via gateways which serve to limit ethernet traffic on the respective networks. The GC LAN gateway also serves to transfer DDU frame data onto the WAN and off-site. The DDU will serve frame data to clients situated either locally or off-site. The specifications for the DDU are listed in Table 7.

Note data distribution to the on-line processing unit will be via a separate dedicated port. The reason for this is to isolate users from the on-line system so that the on-line processes will not be affected by client loads on data distribution and also so that the on-line process may be isolated from inadvertent interference by off-site remote users.

Table 7 Data Distribution Unit (DDU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Data Server	2 node, >200 MFLOPS per node	≥ 1 GB RAM		-
Disk System	-	-		≥ 250 GB striped for 5 users Partition on RAID shared with DAU/DIU
Fast Wide SCSI Port			>40 MB/s	
Network access	-	-	- CACR LDAS LAN (OC3/155 mb/s) - HiPPI	-

3.3.3. Data Ingestion Unit (DIU)

Referring once again to Figure 5, the DIU provides the functionality of reading in raw data tapes arriving from the observatories for processing and transfer to the archive (DDU). Its server dumps the tape data to local disk cache which is then accessible either for data processing and analysis or transfer to the HPSS. The DIU comprises a single interface to the local disk cache based on F/W SCSI technology. Refer to Table 8 for the specifications.

Table 8 Data Ingestion Unit (DIU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
Server/CPU	>200MFLOPS	512 MB		
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Tape Drive/Robot	-	-	>6 MB/s, 3 head 10 tape storage	CDS compatible design
Disk Cache	-	-	40 MB/s, 3 ports	≥ 250 GB Striped for 3 users (3 IFOs) Partition on RAID shared with DAU/DDU

3.3.4. Signal Conditioning and Preprocessing Unit (SCU)

Referring once again to Figure 5, the SCU has four interfaces. The input stream is taken from the disk cache on the DIU on a dedicated high speed F/W SCSI port. Refer to Table 9 for the specifications.

Table 9 Signal Conditioning Unit (SCU)-- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS/CPU ≥ 6 CPUs	1 GB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 2 ports	-
Tape Drive/Robot	-	-	>5 MB/s 5 tape storage	>25 GB/tape
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s) - CSU direct (100BT, 100 Mb/s)	-

3.3.5. Control & Monitoring Server Unit (DVU)

Referring to Figure 5, the DVU provides the user API to configure the off-line LDAS, monitor its health and status, and monitor event lists and other performance metrics derived from on-line processes which are useful for interferometer performance assessment. It consists of one or more high-performance work stations, monitors, and dedicated disk. Access to this system will be provided in the control room in the vicinity of the detector operational monitors so that console oper-

ators may have access to scientific information relevant to maintaining acceptable interferometer sensitivity. Refer to Table 10 for the specifications.

Table 10 Control & Monitoring Unit (DVU)-- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	>200MFLOPS/CPU ≥ 4 CPUs	2 GB RAM	-	-
Ultra Wide SCSI Ports	-	-	> 40 MB/s 1 port	-
Disk	-	-	-	>50 GB
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

3.3.6. Compute Server Unit (CSU)

The CSU provides the computational capacity to process in parallel many optimal filter templates. The CSU will be a dedicated LIGO resource within CACR; however when they are available, all CACR computational resources are accessible for LIGO data processing and analysis. The CSU represents a minimum guaranteed capacity available to the LIGO Laboratory. The EPU is sized to accommodate at least the computational load represented by processing in parallel three interferometer data streams. Thus, as a minimum, it will have 3X the capacity of the individual on-line CSUs at the observatories.

The EPU receives input data directly from the SCU and provides event lists and other reduced data via the LDAS LAN. Results will be post-processed, displayed, or otherwise manipulated separately using DVU components or other local or remote hardware. Refer to Table 11 for the specifications.

Table 11 Compute Server Unit (CSU) -- Off-line

<i>Element</i>	<i>CPU</i>		<i>I/O</i>	<i>STORAGE</i>
CPU	96(TBD) nodes, aggregate 30 IGFLOPS	Per node, 256 MB RAM	100 BT (100 Mb/s)	per node, 10 GB disk
Fast Wide SCSI Ports	-	-	> 20 MB/s 1 port per node	-
Network access	-	-	- LIGO LDAS LAN (OC3/155Mb/s)	-

4 LIGO DATA ANALYSIS SYSTEM SOFTWARE

4.1 General Design Concepts

The LIGO Data Analysis System (LDAS) software will be designed and implemented in a manner that grants the system as much flexibility as possible. This will primarily be achieved by making software portability a high priority in the design. The software will be developed to run under POSIX compliant Unix operating systems with the added flexibility of having all graphical user interfaces (GUIs) supported in both the X11 environment and through a JAVA environment, thus providing a direct web technology interface into the system from remote locations.

4.1.1 Software Standards

Software code developed for the LIGO Data Analysis System will be developed using ANSI standard languages and POSIX compliant system interfaces to the operating system. The primary development language will be C++ which has a mature draft ANSI definition at the time of this writing. C++ is an object oriented language which extends the more traditional procedural approach to programming found in C and FORTRAN with concepts of inheritance, polymorphism and operator overloading which increase code reuse and reduce the cost of code maintenance. C++ also provides direct support for traditional ANSI C code, making binding the two languages straightforward.

C++ is a hybrid language supporting both traditional procedural coding practices and the newer object oriented coding paradigms. LDAS software developed in C++ will be designed and developed in the newer object oriented manner to facilitate code reuse, portability and maintainability.

The standard for GUI design will be based on X11 using Motif to provide a common look and feel to graphical widgets in user interfaces directly connected to the LDAS system. While JAVA will be used to develop functionally identical user interfaces for systems remotely accessing the LDAS system such as through web servers and from portable computers intermittently connected to the system.

To further extend the advantages of software standardizations, an LDAS software style guideline will be integrated into the development program, providing a common approach to classes design, code branching, comments and other aspects of software that offer multiple implementation options.

4.1.2 Software Development

Software development will primarily be carried out within the LIGO Data Analysis System group. The software will be developed with place holders (see FILTERS) where new methods for data analysis can easily be integrated. In addition a Software Developer's Guide will be provided explaining how others can integrate new models easily into the LDAS environment.

4.1.3 Software Code Management

Software configuration management includes the activities of configuration identification, change control, status accounting and audits. Software code development for the LDAS will be under software configuration management using Concurrent Version System (CVS). This provides a full history of the development. CVS is an interface to Revision Control System (RSC) with several additional functions to safely allow concurrent development of individual software components.

4.2 LDAS Software Components

The LIGO Data Analysis System software components are laid out in block form in Figure 7. This figure shows the conceptual functionality of the software as well as the relationship between data, analysis and the end user and the components needed to carry out LIGO data analysis.

4.2.1 Data Types

At the highest level, the data types that will be managed by the LDAS are the following:

- **Frame Format Data** - This is the common data format used in the LIGO Data Acquisition System and in the LIGO Data Archive.
- **Frame Meta Data** - This is the complimentary data associated with the raw data collected by the DAQS, registered during data archival, relational data, results of data analysis and configuration information. This data will reside under a database management system.
- **Event Data** - This is the data inputs and data outputs for the LDAS filters. It includes the LDAS configuration and results of searches by the LDAS. This event data is used as a staging area between event detection and event measurements. As event data matures to the highest level of characterization it becomes permanent Meta Data.
- **Light-Weight Format Data** - This is a simplified TBD data format used to allow importation of data from the LDAS into other software analysis environments such as MATLAB, Mathematica, IDL, etc.
- **DAQS Data Tape** - This is the media used by the DAQS to store the LIGO Frame data at the sites. This data will be transferred to the LDAS for ingestion into the LIGO archive.
- **Filter Data** - This is the data set of waveforms and parameters (including templates) needed by the filters used in gravitational wave analysis of LIGO data.

4.2.2 I/O Libraries

The LDAS will need the following input/output libraries to handle reading and writing of the top level data types used in LIGO data analysis:

- **POSIX Library** - This is the Portable Operating System Interface, providing a standard low level interface to the operating system.
- **Frame I/O Library** - This library handles the reading and writing of data in the Frame Format.

- **Meta Data I/O Library** - This library handles the reading and writing of meta data and will most likely be implemented with database technology.
- **Event Data I/O Library** - This library handles the reading and writing of event data and will most likely be implemented with database technology.
- **Light-Weight I/O Library** - This library handles the reading and writing of data in the TBD light-weight format. It will include the ability to present data in ASCII format.
- **MPI Library** - This is the Message Passing Interface library used to perform the high performance distributed computing necessary for LIGO data analysis.

4.2.3 LDAS Interfaces

The LIGO Data Analysis System will contain components used for configuration, data flow, event detection, event characterization, data ingestion, meta data ingestion, data processing and data analysis. To carry out these actions on the data distributed in this system, the following application programming interfaces will be needed:

- **Data Ingestion API** - This is the component of the system responsible for data ingestion. This involves routing new data from DAQS tapes to the archive, and disk cache, as well as reporting new meta data to the meta data database. The actions performed by this component are controlled through the API.
- **Frame Data API** - This is the component of the system responsible for making request of the Frame Library from the Frame data archive in a form appropriate for data analysis. This includes reducing the content of frames, concatenating frames, performing basic data selection of information stored within frames as requested through the API.
- **Meta Data API** - This is the component of the system responsible for updating the meta database, making requests for information within the meta database, building relationships within the meta database using an API equipped with a query language appropriate for LIGO data analysis.
- **Event Data API** - This is the component of the system responsible for updating the event database, making requests for information within the event database, building relationships within the event database using an API appropriate for various event search strategies, including hierarchical searches and event characterization.
- **Light-Weight Data API** - This is the component of the system responsible for translating data from the frames, meta data and event data into components of the light-weight data format.
- **Control and Monitoring API** - This is the component of the system responsible for configuring, controlling, monitoring the filters being carried out by the LDAS distributed computational resource. This includes detection, parameterization and specialized user requested configurations.
- **Data Conditioning API** - This is the component of the system responsible for preparing data for analysis. This includes conditioning of raw data using calibrations, line removal, regression and bandwidth reduction of the original data. The choice of conditioning is managed through the API. Actual conditioning can be performed within this component or using the filter kernels available through the distributed computing environment.

- **Event Management API** - This is the component of the system responsible for managing events from the filtering kernels. This includes collecting event results, collecting event search results and initiating second tier narrowing of searches for course detection events.
- **LDAS Fast Disk Cache State API** - This is the component of the system responsible for managing the data resources cached on fast disk. This is configurable and included identifying users of cached data, duration of data being cached, as well as types of data cached. The API also manages queuing of requests for cached data.

4.2.4 Distributed Computing Environment

Computer resources for the LIGO Data Analysis System will be distributed over both local and wide area networks. To perform the scale of analysis necessary for LIGO data, this requires a standard for distributed computing. This standard is handled at the lowest level by MPI. Higher level APIs also exist to focus the specialized needs of LIGO data analysis into cleaner, more object oriented components:

- **Message Passing Interface API** - This is the high level message passing interface layered around the MPI library. It handles the distribution of data and commands appropriate to the LIGO Data Analysis System.
- **LIGO Distributed Data Analysis Manager API** - This is the highest level distributed computing component of the LIGO Data Analysis System. It is the switchboard routing data, commands and results of analysis to the appropriate destinations. In essence, it is the command center for the system. It communicates using LDAS commands, MPI, TCP/IP and WWW protocols.

4.2.5 User Interfaces

Authorized users of the LDAS will communicate with the system through three functionally equivalent, though differing in implementation, user interfaces:

- **LDAS Command Language API** - This API provided the primitives for LDAS communications. The actual interface will support command line and script based interfacing to the LDAS. This API sets under the Graphical User Interfaces discussed next.
- **X11 API** - This API is a Graphical User Interface based on X Windows to provide all the interface to the LDAS. It is basically a smart wrapper around the LDAS Command Language API discussed above.
- **JAVA API** - This API is a Graphical User Interface based on JAVA to provide all the interface to the LDAS. It is basically a smart wrapper around the LDAS Command Language API discussed above. Additionally, it can be accessed through the Web services provided by the LIGO Distributed Data Analysis Manager API, providing access from remote machines such as laptops with no more than a JAVA based web browser.


4.2.6 Computational Data Analysis (Filtering)

Filtering lies at the heart of the LIGO Data Analysis System. Filters of one form or another will be used to carry out virtually all types of searches for gravitational waves in LIGO. Since many of

the filters are yet to be envisioned, this component of the system is being conceptualized as modular units with all the interfacing to the LDAS layered away from the actual algorithm so that new filters can easily and quickly be added to the system, possibly even on the fly making the data analysis highly flexible and ultra distributed.

- **Filter Kernel** - This is the compute engine of the analysis filter. It is primarily an algorithm used to by the LDAS to perform a particular type of analysis of the data being distributed from the Data Conditioning API based on configuration directed by the Control and Monitoring API.
- **Kernel API** - This API acts as a wrapper around the Filter Kernel, and will be implemented as a base class in an object oriented environment. It provided the needed integration and communication functionality that allows all filtering algorithms to fit cleanly into the LDAS.
- **Filter API** - This API is responsible for managing various waveforms and parameters used by the distributed network of filter kernels. Among its functions are generating, storing and retrieving waveforms, templates, parameters and other data sets used by various filters that may be active in the LDAS.

Figure 7 Block Diagram of the LDAS Software Components



5 LIGO NETWORKS

5.1. Networking and data access

LIGO is working to develop a wide area network to provide transparent access to data. The wide area network (WAN) will link the four major LIGO Laboratory sites (Caltech, Hanford, Livingston, MIT) as well as institutions in the LIGO Scientific Collaboration. Access to the data archive will be via a web-based technology database query and analysis environment which is being designed and prototyped.

5.2. LIGO Laboratory

- *Universities*

Caltech will comprise of two LANs: the General Computing LAN, inherited from the construction phase and supporting day-to-day activities of LIGO staff at the university, and an LDAS LAN, supporting the data analysis functions described in this document. The LDAS LAN will be distributed on campus between the analysis office area and CACR, where the archive and computational resources reside. Refer to Figure 12 for details on the LIGO GC LAN.

MIT will have a similar topology, although the LDAS component will be of reduced scope, reflecting the fact that there is one common archive and set of computational resources, which will be provided by CACR. Refer to Figure 13 for details about the LIGO GC LAN.

The present baseline is that MIT-Caltech communications will proceed over vBNS.

- *Observatories*

The observatories are in remote locations, and unlike the universities, have no existing networking infrastructure. The plan is to utilize non-LIGO network infrastructure to provide the required connectivity. Each observatory is in proximity to other institutions with whom LIGO is negotiating access to networking resources. Access to the Observatories from the Universities will be provided by a combination of vBNS and ESnet, and the two sites will be linked to the LIGO Laboratory WAN differently.

Refer to Figure 14 and Figure 15 for schematic details of the observatory LAN architecture which complements the LDAS architecture presented in Figure 4 above.

The present baseline is for the LIGO Hanford Observatory to be linked to the LIGO WAN by way of DOE ESnet. The terms for access are being discussed at present. LIGO has requested access to ESnet's ATM cloud system with an initial bandwidth of 3X T1, growing with time as ESnet converts to OC3 and eventually to OC12.

The baseline is that there will be a single crossover of LIGO traffic from ESnet to vBNS. This crossover is TBD, but will either be on the Caltech campus through Caltech/HEP or via SDSC, which already serves as an ESnet/vBNS crossover. MIT communications with Hanford will be routed through this west coast cross-over gateway. It is possible that a similar cross-over gateway may be established at MIT; however discussions on this have not yet started.

Livingston will be connected to vBNS via a gateway to be provided as an agreement with Louisiana State University (LSU) in Baton Rouge, LA. LSU will be linked to vBNS either directly, or

possibly as part of a consortium comprising of five other southeastern states in a regional network called SEPSCoR. SEPSCoR will be linked, via the University of Kentucky, to the vBNS backbone. LIGO is providing for the establishment of a high bandwidth fiber optic link from its Livingston Parish site to the LSU campus gateway.

Figure 8 and Figure 9 show network maps for the vBNS and ESnet infrastructure. Overlaid on the respective maps are the proposed links for LIGO using PNNL at Hanford and LSU at Livingston as local gateways. Table 12 presents the inter-site connectivity matrix as it is presently planned.

Table 12 WAN Connectivity among LIGO Laboratory Sites

<i>Site</i>	<i>Livingston, LA</i>	<i>Hanford, WA</i>	<i>MIT</i>
<i>Caltech</i>	vBNS/OC3	ESnet/3 X T1 <-> vBNS/OC3	vBNS/OC3
<i>MIT</i>	vBNS/OC3	ESnet/3 X T1 <-> vBNS/OC3	
<i>Hanford</i>	ESnet/3 X T1 <-> vBNS/OC3		

5.3. LIGO Scientific Collaboration

The LIGO Scientific Collaboration is composed of the LIGO Laboratory itself and other institutions that have signed memoranda of understandings (MOUs) with the Laboratory. These participating institutions will also require access to LIGO computing and archival resources. Researchers' parent institutions are expected to provide their respective researchers with local network access to LIGO resources through the vBNS network.

Figure 8 NSF's vBNS Infrastructure will allow LIGO Livingston to link up with other Laboratory Sites via NCSA and Louisiana's regional access via SEPSCoR to vBNS (LIGO and SEPSCoR have been added to map; NSF map is from mid 1996)

The National Science Foundation Very-High-Speed Backbone Network Service Logical Network Map

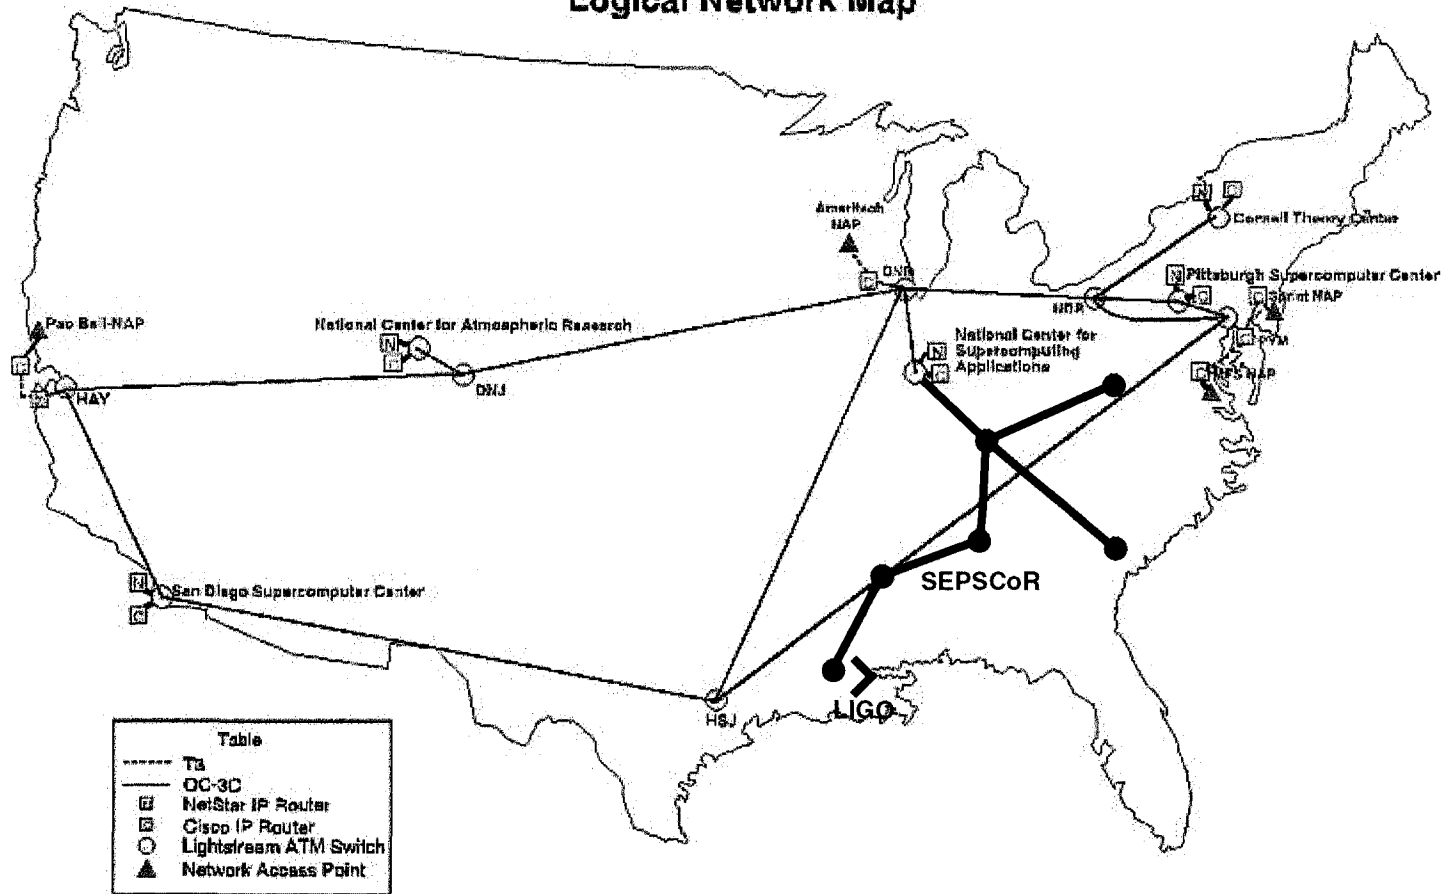
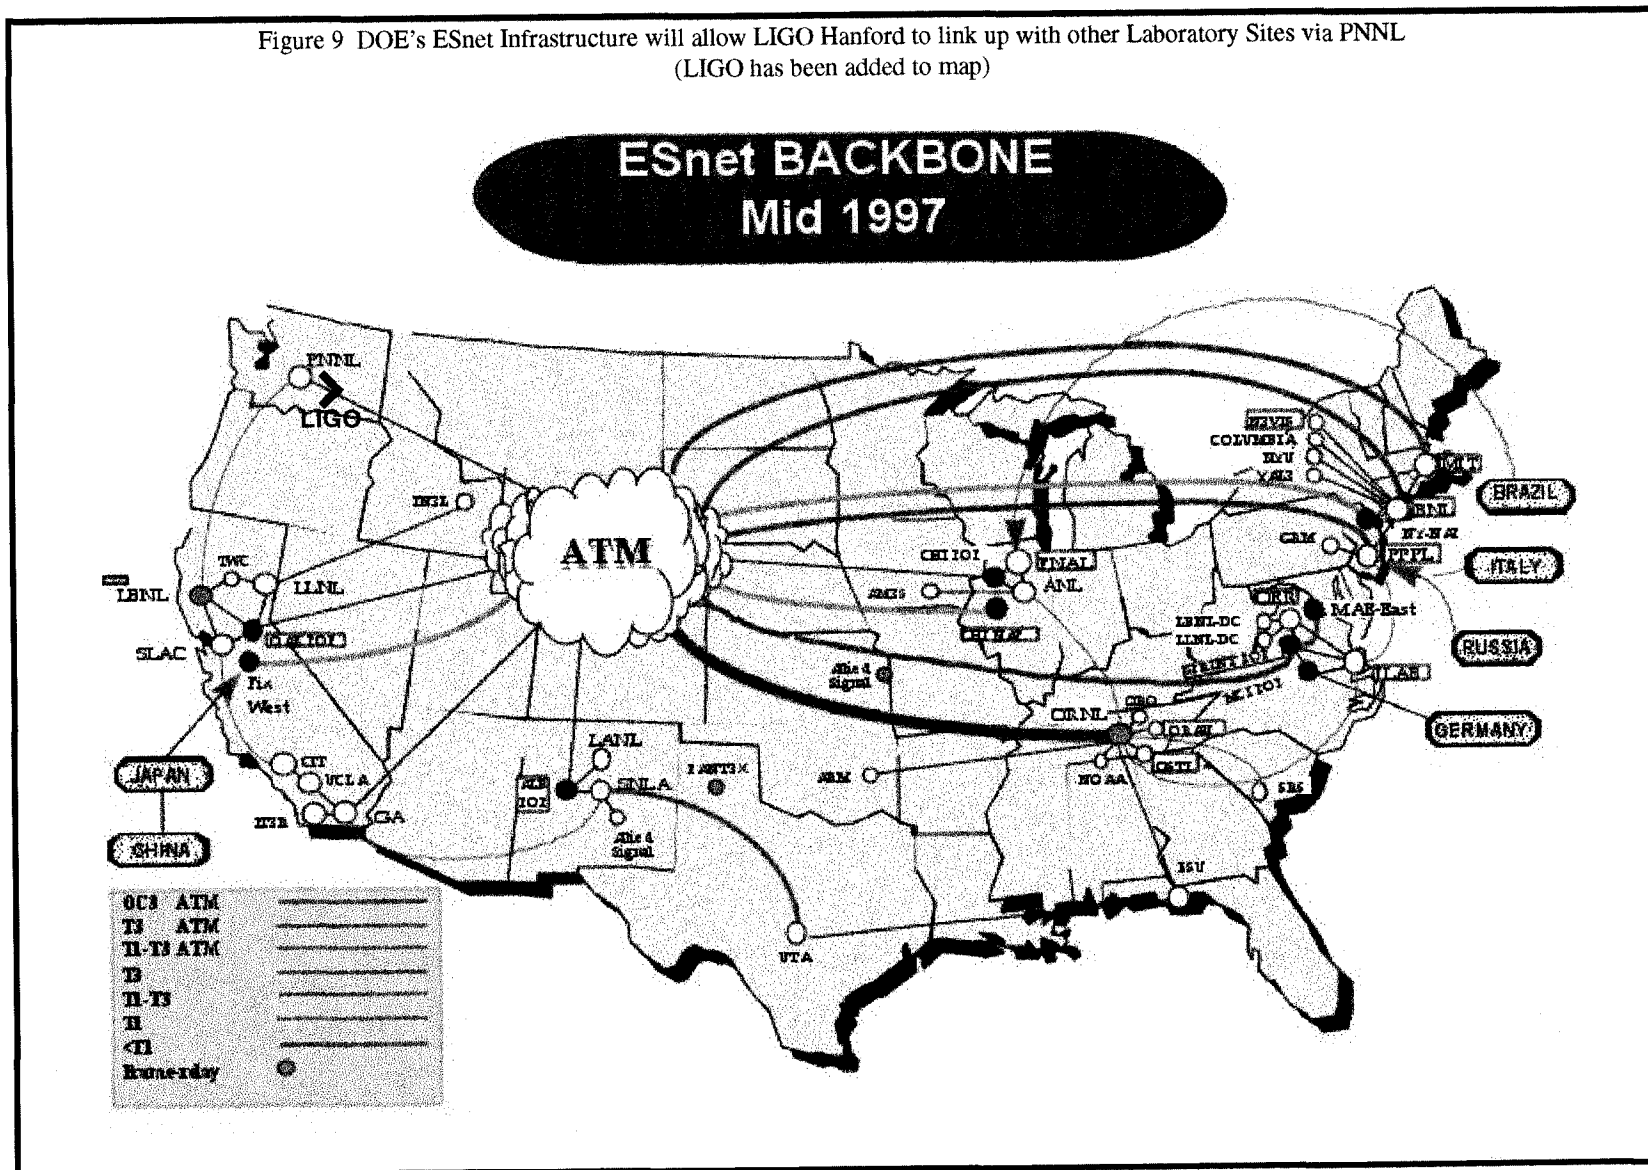


Figure 9 DOE's ESnet Infrastructure will allow LIGO Hanford to link up with other Laboratory Sites via PNNL (LIGO has been added to map)



APPENDIX A LIGO ON-LINE DATA ANALYSIS FLOW

The on-line LDAS component will be used for several functions. The most compute intensive analysis appears at present to be that of searching in real time for binary inspiral chirped waveforms. This is used to determine the target specification for the system. Refer to Figure 10

1. Preprocessing the analysis frame data to obtain calibrated, regressed estimates of strain with instrumental narrowband features removed. This involves several steps.
 - **DATA DROPOUT COMPENSATION.** If the datastream is disrupted during a particular epoch of acquisition - due to loss of lock or some other instrumental or operational problem - it will be desirable to provide for data continuity by either inserting zeroes or mean values into the data stream. This must be done before any other transformations are performed. Process: replace missing or invalid data with predetermined values. Computational impact is negligible.
 - **NARROWBAND INSTRUMENTAL LINE FEATURE REMOVAL.** Using either multi-tapered transform techniques or predictive (e.g., Kalman filter) time domain techniques, narrowband features are removed to decrease signal dynamic range and to improve the SNR of the optimal Wiener filtering step. With a multi-taper method, a number $n < 16$ separate Fourier transformations of the data stream are performed, each with a different weighting function ("Slepian taper filters"). A regression analysis and subtraction process follows to remove by subtraction the narrowband features. Process: $n(<16)$ (real) FFTs of the data stream at the full sample rate (16384 S/s) for a suitably short time $T_{\max} < 1 - 10$ s. This is followed by a linear regression analysis to obtain best estimates of line feature amplitudes and phase offsets for the epoch. Dimensions of the regression analysis is determined by the number of lines to be removed. For m features, the regression involves an $m \times n \times 2$ matrix operations resulting in m pairs of numbers, $\{A_i, \phi_i\}$. The sum of these line feature contributions is then subtracted in the time domain from the raw data.
 - **FOURIER TRANSFORMATION of DATA STREAM.** The data series is converted to frequency domain using FFT techniques. This performed over epochs of duration $T_{\max} < 512$ s.
 - **CALIBRATION AND UNWHITENING.** The data in the spectral domain are multiplied by calibration complex transfer functions to remove end-to-end effects of interferometer response to mechanical strain and data acquisition dynamic range compression. The calibration involves the strain-to-demodulated output transfer function of the interferometer; the unwhitening restores signal dynamic range which was compressed during the acquisition process. Process: frequency bin-by-frequency bin multiplication of the complex calibration function by the line-removed dataset.
 - **DATA REDUCTION.** The previous steps will be carried out at the full temporal and frequency resolution of the acquired datastream. At this point, for inspiral chirp detection, it is possible to compress the volume of data by retaining the frequency range within which the GW waveforms being detected contain most of their signal power. The high-frequency cutoff of the GW channel will be reduced from a Nyquist frequency of 8192 to 512 Hz. This corresponds to a 16:1 reduction in the volume of data to be processed downstream and a corresponding reduction in computation required to process the datastream through the template filter bank.

- **REGRESSION WITH ANCILLARY CHANNELS.** The GW channel will be regressed against a number of other channels acquired simultaneously and which may exhibit instrumental cross-coupling into the GW channel. This regression step, if performed with adequate SNR, will improve the quality of the estimate of GW strain. Process: up to $k(<16)$ ancillary channels will be regressed over the reduced frequency domain. Each of these k channels will require calibration for instrumental effects across f frequency bins. Then multiplication and accumulation of k complex products over each of f frequency bins.
2. Processing the analysis frame data to obtain candidate events. This involves several steps.
- **DATA PREPARATION STEPS** - assuming one is in the middle of a template analysis cycle, there are a number of steps taken for each data epoch and which is common to all template filtering steps:
 - Obtain next epoch of data from the on-line server
 - Additional preprocessing. For example, additional filtering and sample decimation to reduce the sample rate (if this was not already performed earlier)
 - Time windowing for template sizes. Different mass templates have different temporal durations in the LIGO band. The data need to be weighted by a windowing function to minimize finite-sample effects which will occur in transforming a finite sample of data into the frequency domain.
 - Updating noise floor. It is expected that on some (long) time scale the interferometer noise floor will change; maintaining the Wiener filtering at its optimum requires the noise to be stationary. Tracking slight drifts in the measured noise spectrum can be used to keep the filter at optimum.
 - Fast Fourier Transformation. The signal time series is transformed into the frequency domain for optimal filter processing.
 - The signal is divided by the estimated average noise power spectrum in preparation for the filtering process. This prepares the FFT kernel for the next sequence of parallelized steps.
 - **TEMPLATE FILTERING.** The output from the previous sequence of calculations is provided in common to all template filtering processors.
 - Weight kernel by template waveforms. The kernel calculated above is multiplied by the template frequency spectrum.
 - Inverse Fourier transformation. The function is transformed back into the time domain. The output is a vector whose magnitude for each time bin is a measure of signal to noise. The time bin corresponds to a possible time-of-coalescence.
 - Peak detection. The (complex) time-domain vector is scanned for amplitude exceedence which may be indicative of a coalescence event. The output is one or more pairs $\{\xi_i, t_i\}$ corresponding to the possible event. The index “i” corresponds to template number for which the SNR was observed to exceed a threshold.
 - Events are broadcast. The parallelized implementation of this multi-filter process results in a list of candidate events within the epoch of time being analyzed. These are outputs of the data flow.
 - The cycle continues until all templates have been analyzed.
 - Data are provided for visualization, summarization, and further correlation.

The data flow represented in Figure 10 has been modeled using a spread sheet analysis which tracks in detail the computations required to implement the indicated data flow. Where possible,

empirical results derived from benchmark tests using LIGO 40m prototype datasets. These tests have been implemented on a variety of computational resources available at Caltech's CACR.

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
CPU/NODE Engine Model		Calculated 22.09.97
Rated CPU Performance (MFLOPs):	300	400 MFLOPs currently available, 200 MFLOPs PentiumPro
Percent CPU Utilization by FFTS:	70.0%	Based best empirical values measured at CACR
Floating point Operations per Add:	4	Note: adding/subtracting requires a normalization
Floating point Operations per Sub:	4	Note: adding/subtracting requires a normalization
Floating point Operations per Mult:	2	Need to get values in registers
Floating point Operations per Div:	2	
Operations to Cast Integer to Float:	3	
Operations in Memory to Memory Copy:	3	Future hardware will probably off-load this from the CPU
Size of Integer (Bytes):	4	
Size of Float (Bytes):	4	
Size of Double (Bytes):	8	
Size of Complex (Bytes):	8	
Size of Double Complex (Bytes):	16	
Disk I/O Performance (MB/sec):	20	Can do better than this with high end SCSI options
Message Passing Interface (MB/sec):	12	Comes from 100 Mb/second ethernet = 12MB/sec
Number of Nodes in System:	32	Model assumes one CPU per Node
RAM per NODE (MB):	128	
Disk Storage per NODE (MB):	9216	
Memory Resident Operating System (MB):	32	Based on the IBM SP2 which uses 30MB
Disk Resident Oper. Sys. & SWAP (MB):	256	Stripped down version of Unix with approx 100MB of swap

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
System CPU Performance (GFLOPs):	9.6	
System CPU FFT Performance (GFLOPs):	6.72	
System RAM Installed (MB):	4096	
System Disk Storage Installed (GB):	288	
Data Acquisition System Model		
Samples per Second:	16384	Based on CDS DAQS design
Integer Sample Size (Bytes):	4	The option for 24 bit ADCs exists for some DAQS channels
Float Sample Size (Bytes):	4	
Complex Sample Size (Bytes):	8	
Number of Channels used in LDAS:	32	A guess of the number of channels from Frames used by LDAS
Number of IFOs handled by LDAS:	1	
Nyquist Frequency of DAQS (Hz): 8192		
Total Data Supplied by DAQS (MB): 2		
Noise Model		
Seismic Cut Off Frequency (Hz):	30	Based on KB's Noise Model from Summer 97
Minimum in Noise PSD (Hz):	145	Based on KB's Noise Model from Summer 97
Physical Constants		
Gravitational Constant:	6.67E-11	MKS
Speed of Light:	3.00E+08	MKS
Mass of the Sun:	1.98E+30	MKS

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
π	3.14E+00	MKS
Binary Inspiral Source Model		
Each Component Mass (Msun):	1.2	
Percent Loss of Events Desired:	10.0%	One out of Ten events fail to fall close enough to a template
Maximum Frequency of Interest (Hz):	512	Based on Document from Sam Finn of Sept. 97
Inspiral Time for Binary (Seconds):	70.038	Based on 1PN
Number of Inspiral Templates per IFO:	34323	Based on Ben Owen's latest 2PN restricted waveform model
Total Number of Inspiral Templates:	34323	Factor up by the number of IFOs serviced by this system
Nearest 2 ^N Nyquist Frequency (Hz):	512	Closest power of two frequency to the Max Freq above
Black Hole Ring-Down Source Model		
Percent Loss of Events Desired:	10.0%	One out of Ten events fail to fall close enough to a template
Maximum Quality Factor:	89.5	Values up to 100 astrophysically expected, limited by CPU
Minimum Frequency of Interest (Hz):	30	Basically the Seismic cutoff
Maximum Frequency of Interest (Hz):	1024	Near the upper cutoff from astrophysics, higher expensive
Decibels Spanned by Ring-Down Template:	20	Two orders of magnitude range in the signal amplitude
Ring-Down Time (Seconds):	4.373	Based on decay of exponential through the decibels above
Number of Ring-Down Templates per IFO:	1852	Based on model from GRASP manual and J. Creighton
Total Number of Ring-Down Templates:	1852	Factor up by the number of IFOs serviced by this system
Nearest 2 ^N Nyquist Frequency (Hz):	1024	Closest power of two frequency to the Max Freq above

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Pre-Analysis Data Conditioning		
<u>Data Drop-Out Corrections</u>		
Percent of DAQS Frames with Drop-Outs:	0.10%	Only have to correct Frames with Drop-Outs
Data Samples in One Second Frame:	16384	
Number of Frames Resident in Memory:	3	Need before and after Frames to make best fill of data
Frame Data in Memory (MB):	6	
MFLOP to Calc pre RMS post of Channels:	4.2	Need to average all points in before and after Frames
MFLOP to Fill in Drop-Outs:	3.1	Need to fill in with RMS or straight line
Average MFLOPS Needed:	0.007	Accumulated FLOP only in the percentage of bad Frames
<u>Spectral Line Removal</u>		
Seconds of data Windowed per Pass:	4	Reasonable value for Multi-Taper Methods
Number of Lines Removed per Pass:	64	40 lines isolated and removed from 40-meter...slight increase
Number of Slepian's Used:	7	Must less than 2x number of windows in pass above
Samples in Each Data Segment:	8192	Only need Nyquist Frequency for Bandwidth used in Searches
MFLOP to Window Each Segment:	0.1	Each Slepian is multiplied by data to window
MFLOP to FFT each Windowed Segment:	37.3	Each windowed data set is FFT'ed
MFLOP to Set Phase of Cosine Vector:	0.0	The cosine vectors stored in circular buffer so phase is indexed
MFLOP to Set Amp of Cosine Vector:	0.5	Each cosine vector has a multiplication for the amplitude
MFLOP to Remove Lines from Segment:	2.1	Each cosine vector must be subtracted from the data
Total MFLOP Needed to Do Line Removal:	40.0	
Average MFLOPS Needed:	10.00	Average over the number of seconds of data
MB Memory Needed by Line Removal:	11.3	Memory needed to store data, cosine vectors and slepian's
<u>Calibration of GW Channel</u>		
Samples per second to Calibrate:	2048	Only calibrate the bandwidth relevant to searches

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Window Data (Time Domain):	0.004	Apply a time domain window to the GW data
MFLOP to FFT One Second of Data:	0.1	FFT the windowed GW data from time to freq domain
MFLOP to Mult by Complex Calibration:	0.016	Perform complex multiply with calibration transfer function
MFLOP to Inverse FFT Data:	0.1	Out of place IFFT to get time domain calibrated data
MFLOPS to Calibrate Data:	0.2	Total MFLOPS for this calibration: MFLOPS / one second
<u>Perform Data Reduction</u>		
Reduction for Inspiral:	16	Data sampling rate reduction for inspiral
MFLOP to Low Pass Inspiral Data:	0.786	Use 8(?) coefficient FIR for low pass filter
MFLOP to Average Down Inspiral Data:	0.068	Take average of filtered samples to reduce sample rate
Reduction for Ring-Down:	8	Data sampling rate reduction for ringdown
MFLOP to Low Pass Ring-Down Data:	0.786	Use 8(?) coefficient FIR for low pass filter
MFLOP to Average Down Ring-Down Data:	0.070	Take average of filtered samples to reduce sample rate
MFLOPS to Reduce Data Rates:	1.7	Total MFLOPS for reducing both inspiral and ringdown
MB Memory Needed by Data Reduction:	0.074	Total memory needed for performing the data reduction
<u>Perform Linear Regression</u>		
Samples per second to Regress:	2048	Only regress the bandwidth that will be used in searches
MFLOP to FFT Non-GW Channels	3.49	Perform FFT on other channels that are part of the regression
MFLOP to Calibrate Non-GW Channels:	0.016384	Perform complex multiplication of these with freq calibration
MFLOP to Regress GW Channel:	0.524288	Perform the single row complex multiply to regress GW channel
MFLOP to Inverse FFT Regressed Data:	0.11264	Inverse FFT to get time domain regressed data
MFLOPS to Regress GW Channel:	4.15	Total MFLOPS needed to perform this regression
MB Memory Needed by Linear Regression:	0.2578125	Estimated memory needed to perform this regression
<u>Total MFLOPS and Memory Used in Data Conditioning</u>		
Total MFLOPs Used in Data Conditioning:	16.10	Estimate of total CPU performance needed for data conditioning
Total MB Used in Data Conditioning:	17.626	Estimate of total Memory needed for data conditioning

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Binary Inspiral Data Analysis Model		
<u>Inspirial Data Model</u>		
Overlap Factor (power of 2):	8	Overlap is to apply template to more data per convolution
Max Template Size (power of 2):	1048576	The overall template size, (signal + 0 padding) should be 2^N
Zero Padding Added to Template:	976857	This is the about of zeros added to make 2^N and \sim overlap
Convolution Size (samples):	1048576	Number of samples (in time domain) for template/data
Size of Template (MB):	8	Size of Template in the frequency domain - templates not real!
Size of GW Data (MB):	8	Size of data in freq domain - made equal to template
Size of PSD (MB):	2	Size of Power Spectral Density of the Noise
New Data Analyzed Each Pass (sec):	953.962	Seconds of new data introduced into the convolution each pass
Percent of Data Analyzed Each Pass:	93.16%	Percentage of the convolution that is new each pass
Total Template Data Bank Size (GB):	268.1	Number of templates times Template Size
<u>Inspirial Node Model</u>		
Analysis Code Core Size:	2	Assume Analysis Code resides in 2 MB of core(RAM)
Amount of Temp Space (xGWdata):	7	Rough estimate of the about of temporary storage needed
RAM used by Data + Temp Space (MB):	64	Amount of RAM used by data and associated temp storage
RAM Available for Templates (MB):	32	Remaining RAM available on the system for holding templates
Number of Templates on Node Disk:	1073	Number of templates stored on disk on each node
Number of Templates in Node RAM:	4	Number of templates that fit in RAM at one time on each node
Template Disk Storage Needed: (GB):	8.38	Total Disk Storage needed on each node to hold its templates
Number of Template Swap-Outs Needed:	268.25	Number of times the templates in RAM must be swapped
<u>Once per Inspiral Pass Data I/O</u>		
Percent of Time CPU involved with I/O:	5.00%	Estimate of time percentage CPU used for I/O

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Load PS weighted Data on Node (sec):	0.667	By MPI rate, time spent loading preconditioned data on node
Template Swap-Out Time per Node (sec):	429.038	By SCSI rate, time spent loading templates on node
Total Time Moving Data (sec):	429.704	Total seconds of I/O each pass per node
Effective MFLOP lost to I/O:	6445.6	Lost CPU cycles to this I/O
<u>Computation per Inspiral Template</u>		
Number of values stored in Result Element:	3	The result is array of length number of samples by this number
Mult Template & weighted Data (MFLOP):	16.0	MFLOP used in complex multiplication of template with data
Convolve with Inverse FFT (MFLOP):	142.9	MFLOP for FFT based on $5N\log_2(N)$
Square Convolution (MFLOP):	8.0	Multiply complex convolution by its complex conjugate
Compare to Previous Convolution (MFLOP):	4.0	Do compare of each square in array with MAX stored for pass
Replace MAX info in Result (MFLOP):	3	Replace 3 elements in Max array when Max occurs
Computation for One Template (MFLOP):	173.9	Total Mega FLOP used per template on node
<u>All Inspiral Template Computations</u>		
Total MFLOP for All Templates on Node:	186548.7	Total MFLOP used for all template calculations
Total MFLOP when I/O included on Node:	192994.3	Add in the effective MFLOP count lost to I/O
Total Time Available for Computations (sec):	953.962	Amount of new data in seconds analyzed in this pass
Node CPU Performance Needed (MFLOPs):	202.31	Needed node performance of CPU in MFLOP per second
Ring-Down Data Analysis Model:		
<u>Ringdown Data Model</u>		
Overlap Factor (power of 2):	8	Overlap is to apply template to more data per convolution
Max Template Size (power of 2):	131072	The overall template size, (signal + 0 padding) should be 2^N
Zero Padding Added to Template:	122116	This is the about of zeros added to make 2^N and \sim overlap
Convolution Size (samples):	131072	Number of samples (in time domain) for template/data

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Size of Template (MB):	1.0	Size of Template in the frequency domain - templates not real!
Size of GW Data (MB):	1.0	Size of data in freq domain - made equal to template
Size of PSD (MB):	0.5	Size of Power Spectral Density of the Noise
New Data Analyzed Each Pass (sec):	59.627	Seconds of new data introduced into the convolution each pass
Percent of Data Analyzed Each Pass:	93.17%	Percentage of the convolution that is new each pass
Total Template Data Bank Size (GB):	1.809	Number of templates times Template Size
<u>Ringdown Node Model</u>		
Number of Nodes Reserved for Ringdown:	2	Number of Node of identical Class used for Ringdown
Analysis Code Core Size:	2	Assume Analysis Code resides in 2 MB of core(RAM)
Amount of Temp Space (xGWdata):	7	Rough estimate of the about of temporary storage needed
RAM used by Data + Temp Space (MB):	8	Amount of RAM used by data and associated temp storage
RAM Available for Templates (MB):	88	Remaining RAM available on the system for holding templates
Number of Templates on Node Disk:	926	Number of templates stored on disk on each node
Number of Templates in Node RAM:	88	Number of templates that fit in RAM at one time on each node
Template Disk Storage Needed: (GB):	0.90	Total Disk Storage needed on each node to hold its templates
Number of Template Swap-Outs Needed:	10.52	Number of times the templates in RAM must be swapped
<u>Once per Ringdown Pass Data I/O</u>		
Percent of Time CPU involved with I/O:	5.00%	Estimate of time percentage CPU used for I/O
Load PS weighted Data on Node (sec):	0.083	By MPI rate, time spent loading preconditioned data on node
Template Swap-Out Time per Node (sec):	46.3	By SCSI rate, time spent loading templates on node
Total Time Moving Data (sec):	46.383	Total seconds of I/O each pass per node
Effective MFLOP lost to I/O:	695.8	Lost CPU cycles to this I/O
<u>Computation per Ringdown Template</u>		
Number of values stored in Result Element:	3	The result is array of length number of samples by this number
Mult Template & weighted Data (MFLOP):	2.0	MFLOP used in complex multiplication of template with data

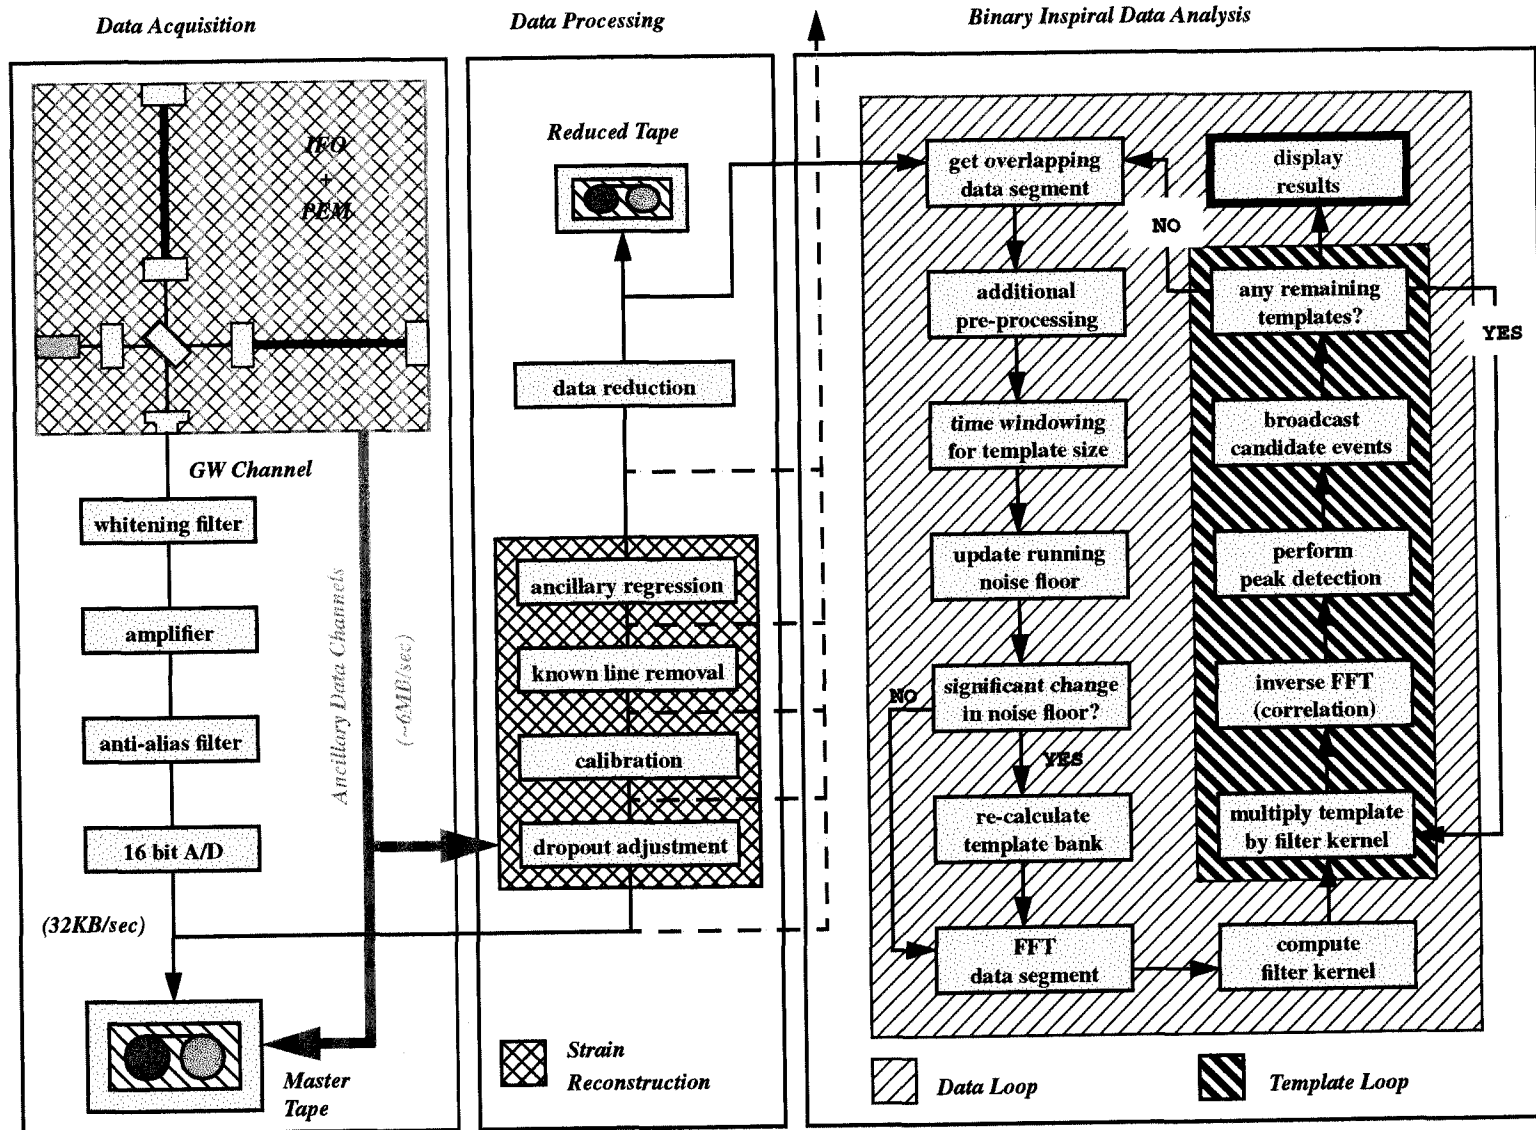
Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

Elements of Data Analysis Flow Model	Value	Comments
Convolve with Inverse FFT (MFLOP):	15.2	MFLOP for FFT based on $5N\log_2(N)$
Square Convolution (MFLOP):	0.5	Multiply complex convolution by its complex conjugate
Compare to Previous Convolution (MFLOP):	0.5	Do compare of each square in array with MAX stored for pass
Replace MAX info in Result (MFLOP):	0.375	Replace 3 elements in Max array when Max occurs
Computation for One Template (MFLOP):	18.6	Total Mega FLOP used per template on node
All Ringdown Template Computations		
Total MFLOP for All Templates on Node:	17180.6	Total MFLOP used for all template calculations
Total MFLOP when I/O included on Node:	17876.4	Add in the effective MFLOP count lost to I/O
Total Time Available for Computations (sec):	59.627	Amount of new data in seconds analyzed in this pass
Node CPU Performance Needed (MFLOPs):	299.8	Needed node performance of CPU in MFLOP per second
ROM Cost Estimates in 1997 US Dollars		
Dollars per MB of RAM:	\$20.00	Estimate of component cost (mid-range values)
Dollars per GB of Disk Storage:	\$200.00	Estimate of component cost (mid-range values)
Dollars per Node:	\$5,000.00	Estimate of component cost (mid-range values)
Dollars per 24 port 100bT Switch:	\$10,000.00	Estimate based on ATM/100bT Switch technology
Dollars per Node 100bT Interface PCI Card:	\$200.00	Estimate of component cost (mid-range values)
Total Number of Nodes:	35	
Total Number of 100bT PCI Cards:	35	
Total Number of Switches:	2	
Total Megabytes of Ram:	4480	
Total GB of Disk Storage:	315	
Cost of RAM:	\$89,600.00	
Cost of Disk Storage:	\$63,000.00	

Table A1: Sample Computational Requirements Analysis For Real-Time Detection of Binary Inspirals and Black Holes Ringdowns

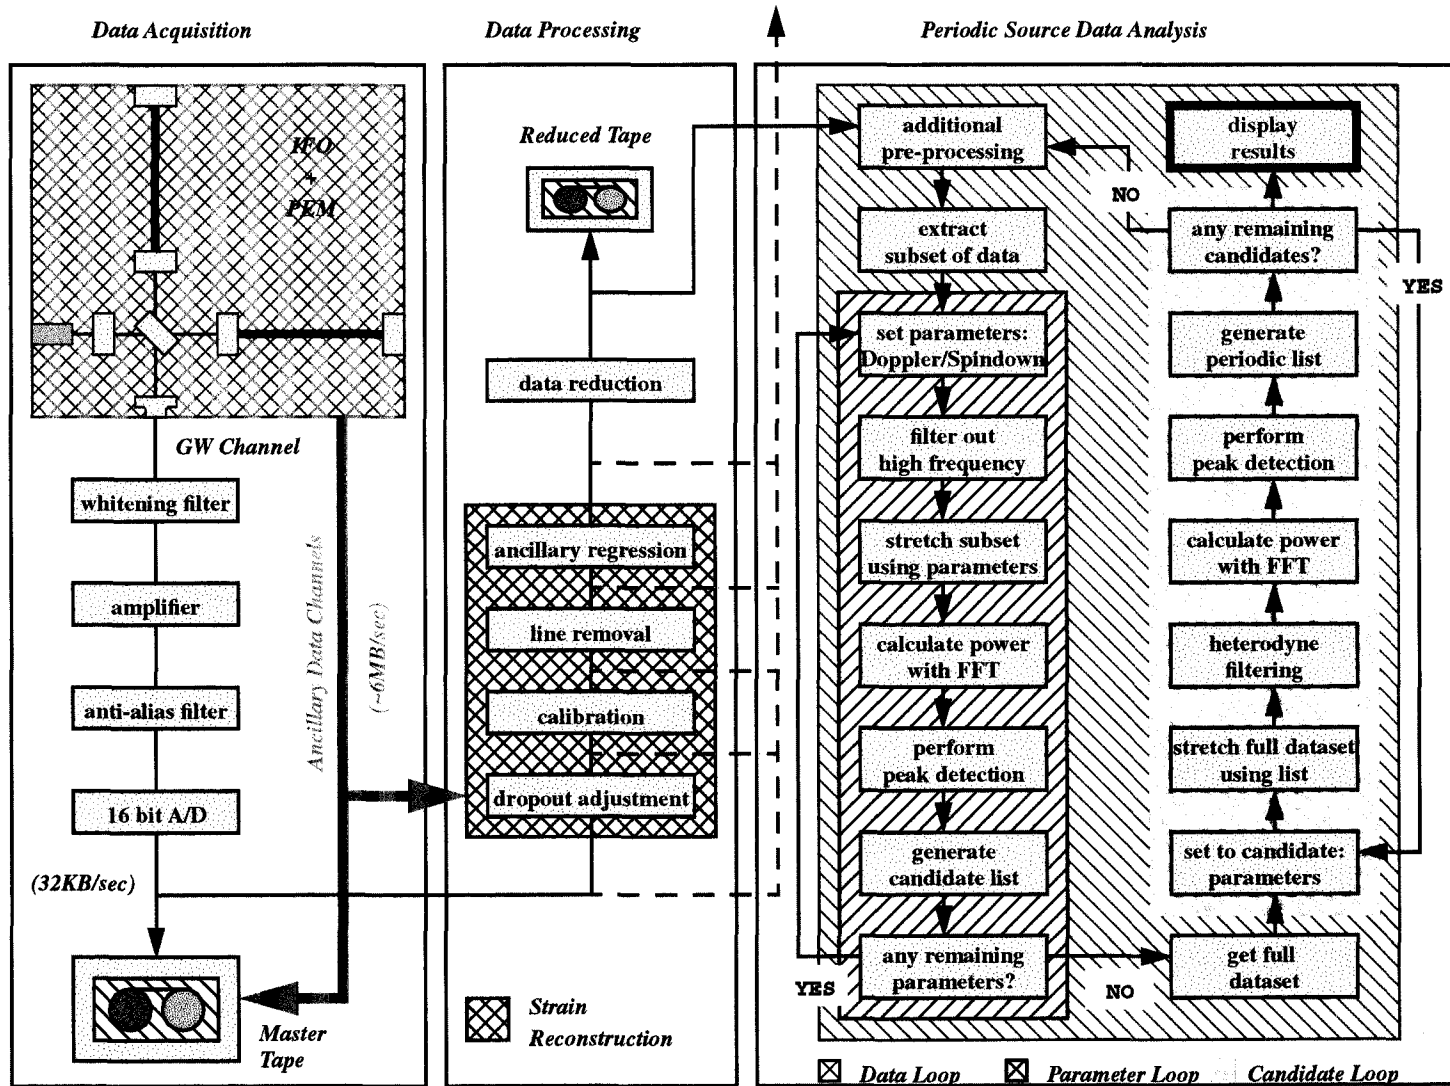
Elements of Data Analysis Flow Model	Value	Comments
Cost of CPUs:	\$175,000.00	
Cost of Misc.:	\$27,000.00	
Total Cost of System:	\$354,600.00	This is per System; 3 will be needed
SUMMARY INFORMATION		
Minimal Mass Achieved in Inspiral:	1.2	
with a cut-off frequency in Hz of:	512	
leaving the following MB of node storage:	379.25	
leaving the following MB of node RAM:	0	Note: algorithm designed to use close to all memory
with the following node-core templates:	4	
utilizing the following percentage of CPU:	67.44%	
of which FFTs are:	79.43%	
Minimal Q Achieved in Ringdown:	89.5	This may be on the high side
with a lower frequency in Hz of:	30	
with a higher frequency in Hz of:	1024	This may be on the low side
leaving the following MB of node storage:	8034	Ringdown templates are short, leaving lots of disk space
leaving the following MB of node RAM:	0	Note: algorithm designed to use close to all memory
with the following node-core templates:	88	
utilizing the following percentage of CPU:	99.93%	
of which FFTs are:	78.63%	

Figure 10 Data analysis flow for binary inspiral detection



**APPENDIX B LIGO OFF-LINE DATA ANALYSIS
FLOW**

Figure 11 Data analysis flow for periodic source detection



APPENDIX C LIGO LAN ARCHITECTURE DETAILS FOR UNIVERSITIES AND OBSERVATORIES

Figure 12 Planned LIGO General Computing LAN at Caltech showing the link to the Caltech LDAS LAN and also to MIT and Sites

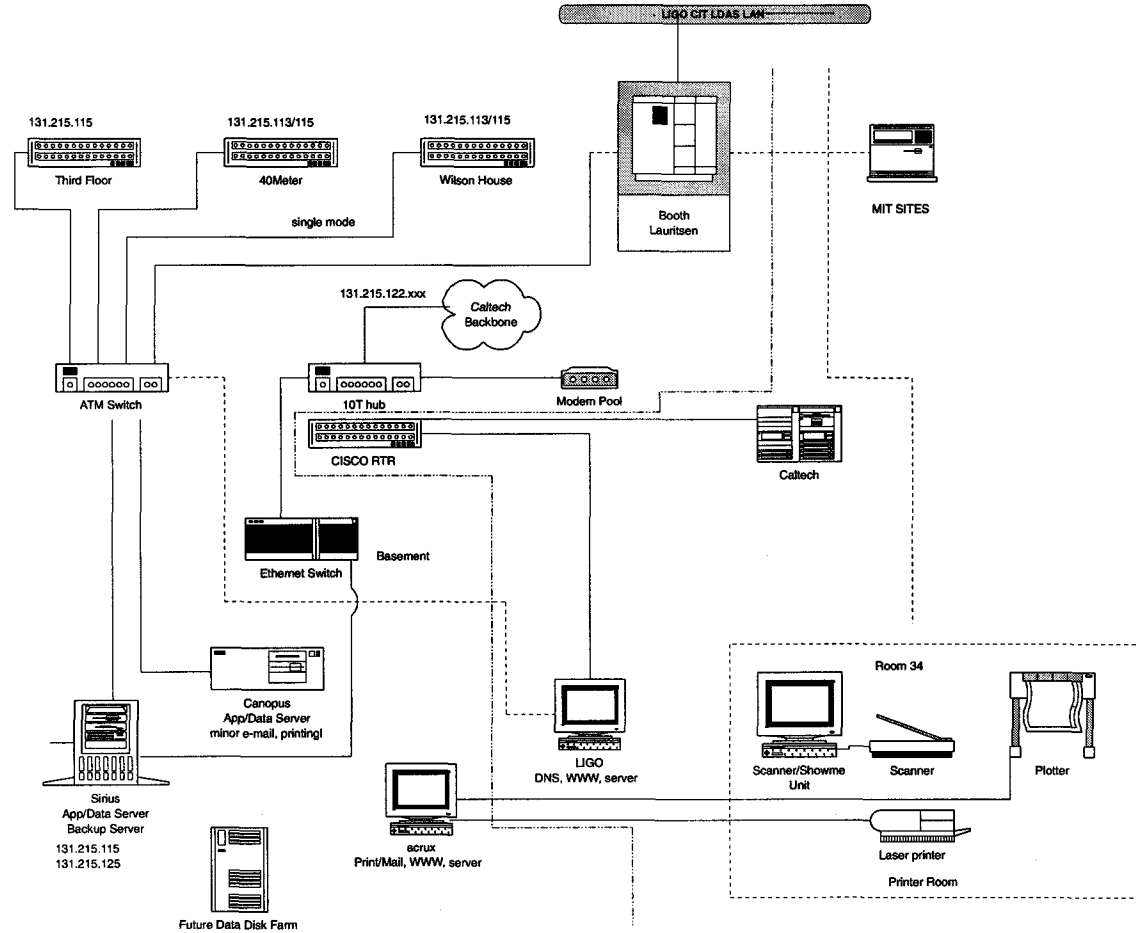


Figure 13 Planned LIGO General Computing LAN at MIT showing the link to Caltech and to the Sites (TBD)

Figure 14 Planned LIGO General Computing LAN at the Observatories showing the link to the Universities and local LDAS LAN

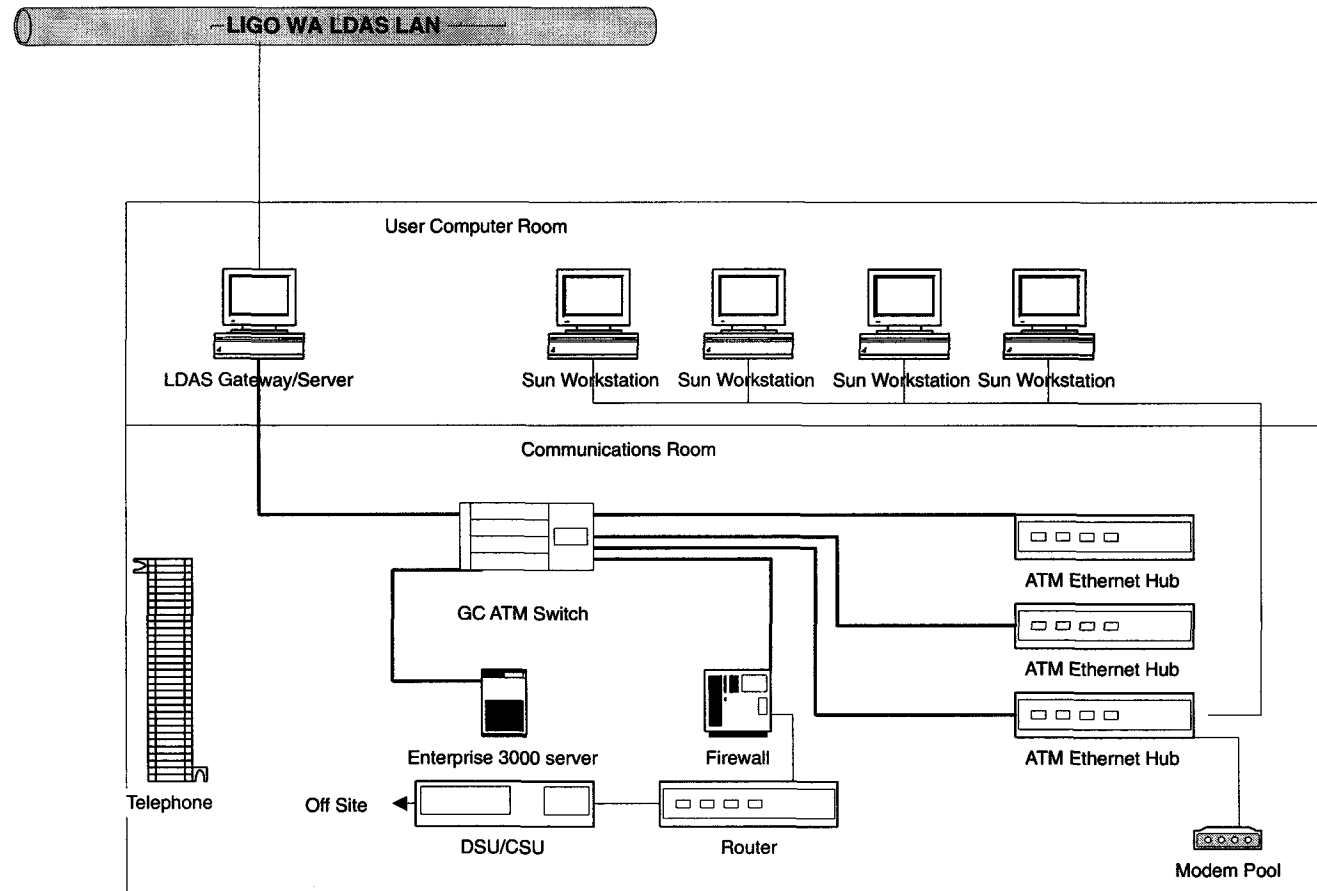


Figure 15 LAN Interfaces for the Observatories

