

Executive abstract

Technology advances in both laboratory- and field-data acquisition, as well as computing, present us with an embarrassment of riches of sorts; it is difficult to store and archive, process and analyze, as well as visualize and absorb the large amounts of data required for progress in many areas of research. In many disciplines, such as fluid turbulence, astronomy, global-climate modeling, biology and neuroscience, high-energy physics, and others, the nature of many thrusts requires dealing with very large amounts of data. Often, the capability of acquiring and generating the requisite data sets is at hand. The ability and infrastructure to handle such data sets is, however, lagging. The *Distributed Teravoxel Data System: Acquisition, Networking, Archiving, Analysis, and Visualization*, proposes to address this deficiency through the development of generic acquisition and processing capabilities and infrastructure, to be hosted at the California Institute of Technology.

1. Research activities

Many research activities stand to benefit from the proposed *Teravoxel* data acquisition and infrastructure system. Primary activities will focus on fluid turbulence and ground-based astronomical observations. Other areas also stand to gain, as discussed below.

1.1 Primary research

1.1.1 Flow turbulence

Of the 19th-century science issues bequeathed to the 20th, turbulence is the only one relayed to the 21st. Important over a broad spectrum of science, technology, and environmental applications and problems, and despite significant 20th-century progress, turbulence must be regarded as the last, largely unsolved, classical-physics problem. One of the reasons is that turbulence is a *non-linear, three-dimensional, high-dynamic-range, unsteady*, process. As a consequence, the volume of space-time data required for its description alone, either experimentally or computationally, is very large.

Some preliminary results




Fig. 1 Jet-concentration isosurface.

An illustration of envisaged data in the form of a 3-D (space-time) laser-induced fluorescence image (Fig. 1) is included [Dimotakis & Catrakis 1999]. Computed from data recorded and visualized three years ago, it depicts isosurfaces of jet-fluid concentration in the far field of a turbulent jet (~ 300 diameters from jet nozzle). If this was a chemically reacting jet and the isosurface concentration level was the stoichiometric mixture ratio, the depicted surface would represent the burning surface in a turbulent, non-premixed flame. The complex geometry captured is more complicated than (power-law) fractal. The (2-D space + time) data were acquired with a spatial resolution of 1024^2 -pixel frames, at 1frame/s (time increases into the paper),

using a synchronized, swept argon-ion laser beam. Processing limitations at the time restricted the data that could be analyzed to 1/400 of the set recorded and visualized (depicted image). The sparseness of data of this kind, along with the complexity of the geometry that must be described to quantify turbulent mixing have limited progress in this important area.

A second illustration (Fig. 2) depicts the space-time isosurface of a turbulent jet in a cross flow [Shan, Dimotakis, & Lang 1999]; an important flow for pollutant dispersion, fuel injection in supersonic combustor ramjets, and many other applications. Freestream flow is from left to right with time increasing into the paper. The data were acquired with the Cassini digital camera (see below), as 250 1024²-pixel frames, 12-bits/pixel, 10 frames/s (15MB/s for 25 seconds: ~375MB), using a 10Hz Nd:YAG laser. Angles in this view represent convection velocity. The 3-D space-time structure illustrates the complex entrainment process that mixes freestream fluid with injected jet fluid. Processing limitations force us to subsample the data to 512²-pixel frames, to make the computation of the 3-D flow structure geometry possible.




Fig. 2 Transverse jet space-time isosurface

Very similar data can be generated via direct numerical simulation (DNS). Fig. 3 derives from a simulation by A. Cook, Lawrence Livermore National Laboratory (LLNL), in collaboration with P.




Fig. 3 Rayleigh-Taylor instability. $\rho_2/\rho_1=3$ fluid atop $\rho_1=1$ fluid; downward g . $\rho_2/\rho_1=2$ isosurface evolution.

pressure field (publication submission pending). Data handling, rendering, and visualization

Dimotakis, of the initial stages of flow generated by the Rayleigh-Taylor instability. This instability is encountered when low-density fluid accelerates (pushes) higher-density fluid, as occurs when heavy fluid finds itself atop light fluid in a gravitational field, for example. The phenomenon is important in astrophysics, geophysical flows, mixing in accelerating/rotating flows, in inertial confinement fusion, and many others. Recognized as an important instability about a hundred years ago, the characterization of its latter non-linear and turbulent stages has proven elusive to date. The simulation leading to the included figure was of two fluids with a density ratio, $\rho_2/\rho_1 = 3$. It was performed on a $256^2 \cdot 512$ spatial grid ($3.36 \cdot 10^7$ spatial nodes), on 96 IBM/SP2 processors (LLNL ASCI-Blue computer); approximately one month of wall-clock time. Five (floating-point) fields must be stored for each time step archived (density, pressure, and three velocity components). Approximately, 0.1TB (10^{10} bytes) of data/run must be stored. To date, only scalar-field (density) data have been analyzed and visualized from three such runs, with a preliminary visualization of the

limitations have prevented analysis of velocity-field data. New simulations are in development as part of the Caltech/ASAP program (DOE) to design a Virtual Test Facility (VTF) to simulate the dynamic response of materials (PI: D. Meiron). These simulations do not extend to fully developed turbulence. That would require at least $\cdot 8$ more grid points ($>1024^2 \cdot 2048$), and $\cdot 16$ more data (1.6TB/run; shorter time steps). Such simulations are planned for the next-generation (ASCI-White) machines. The proposed *Teravoxel* system will accommodate and analyze such DNS data sets.

Multi-D turbulence data

Three- and four-dimensional measurements of scalar-field concentration in turbulent flow, of the type illustrated above, permit other information to be extracted. Scalars are convected in a manner described by the scalar-transport equation. A succession of such data can be inverted to yield measurements of the flow velocity field. Dubbed Image Correlation Velocimetry (ICV), this method is akin to DPIV techniques, but can process both particle and/or continuous-scalar images. Other important turbulence quantities can also be extracted from such multi-D data, such as spectra, geometric (fractal) isosurface properties, local mixing rates, velocity-gradient-tensor eigenvalues and eigenvectors, *etc.* Such data, necessary for probing the dynamics of turbulence and model validation, are mostly out of reach of conventional measurement methods; a consequence, in part, of the difficulties associated with the acquisition and handling of the massive amounts of data required.

Turbulence measurement requirements and the KFS imaging system

Turbulence scaling dictates the minimum spatial and temporal scales required to represent multi-D, spatio-temporally resolved (alias-free) data. These specifications led to the design of a kilo-frame/second (KFS) digital-imaging system that could record 1024^2 -pixel frames at rates up to 10^3 frames/s, to capture a 10^9 -voxel (volume-element) data set in 1s. Such a system is able to capture alias-free turbulent-flow data, at Reynolds numbers sufficiently high for fully developed turbulence. The instrumentation development, undertaken in collaboration with the JPL imaging group, with whom a special (32-output-channel) focal-plane CCD was designed and fabricated, is now at the




Fig. 4 First KFS-CCD images of spinning drill chuck. Synchronized strobe-light illumination, 1024^2 -pixel frames, 12-bits/pixel, 50 frames/s.

bench-test and calibration stage. In its present, Phase-I, implementation, the system can capture 4GB/run, or 4-6 seconds of data with (lossless) data compression. Many such runs must be recorded and analyzed to capture the statistical excursions characterizing turbulence. A pair of 1024^2 -pixel, 12-bit/pixel frames (Fig. 4) depicts raw KFS-CCD data (no pixel-to-pixel offset/sensitivity corrections) from a sample 50-frame/s sequence of a spinning drill chuck, illuminated with a synchronized strobe light. Needless to say, the reproductions do not capture the spatial and dynamic range of the data.

For laser-induced fluorescence applications (*e.g.*, Figs. 1, 2), the laser pulse-repetition frequency represents the technology-limiting component of this new powerful new diagnostic. At this writing, commercially available Nd:YAG lasers can be stacked in pairs to attain up to 100Hz at 300mJ/pulse, providing the illumination source for 100-frame/s sequences. Much new and important turbulence information will be within reach at such frame rates. Fully developed turbulence, at Reynolds numbers, $Re \approx 1\text{-}2 \cdot 10^4$, however, will require acquisition at up to the full kHz-frame rates for which the KFS system was designed. We are presently discussing with the Laser Division of the Lawrence Livermore National Laboratory to develop a 350mJ/pulse (at 532nm), 1kHz Nd:YAG laser (350W average power at 532nm), to be used in a new class of 3-D and 4-D turbulence measurements.

1.1.2 Ground-based astronomical observations

Ground-based astronomical observations are hampered by atmospheric turbulence responsible for (primarily) phase aberrations of the light beam collected by the telescope aperture. Recently, such effects have been mitigated through adaptive optics. Present implementations of such technologies, however, are limited in the spatial degrees of freedom afforded by the number of phase sensors and actuators, the associated control-system parameter space dimensionality, and bandwidth that can be accommodated in real time. In an alternative approach, we have proposed to record short-time-exposure images aberrated by atmospheric turbulence, concurrently with phase-front data, at frame rates commensurate with characteristic atmospheric-turbulence time scales (2-6ms), for *ex post facto* image reconstruction in the computer from a (long) data sequence. Recording over several minutes at a time would be ideal, at commensurate frame rate (150-500frames/s). At this writing, the 9-10 May 2000 nights have been reserved on the 200" Palomar telescope, for the first such measurements using the Phase-I KFS imaging system (B. Kern, D. Lang, P. Dimotakis, & C. Martin).

1.2 Research support sources

The work on turbulence is supported by the Air Force Office of Scientific Research (AFOSR – PI: P. Dimotakis) and by the Dept. of Energy (Caltech ASCI/ASAP program – PI: D. Meiron). The KFS system was developed as part of a DOD (DURIP) grant and under continuing AFOSR support (PI: P. Dimotakis). Past and pending support on optical beam propagation through turbulence has come from AFOSR (PI: P. Dimotakis). The *ex post facto image* reconstruction scheme for ground-based astronomy is currently supported by the National Science Foundation (PI: C. Martin).

[DR & JCTP please review] Work on high-speed networking is supported by the Lee Center for Advanced Networking (Director: D. Rutledge) and Caltech internal funds. High-performance computing, storage, and visualization development is part of the on-going work at Caltech's Center

for Advanced Computing Research (CACR; Director: J. C. T. Pool). This is supported by NPACI, sponsors of various computational-/computer-science CACR projects, and Caltech. [JCTP: +?] Computer-visualization efforts represent a major part of Computer Science and CACR activities at Caltech, as well as part of various computational and other efforts on campus. The NSF-sponsored (Grant ACI-9982273) focus on Multiresolution Visualization Tools for Interactive Analysis of Large-Scale N -Dimensional Datasets [PI: D. Breen] is addressing many closely related issues.

1.3 Other and future research activities

Some other data-intensive applications, whose output tax presently available data-processing capabilities, are listed below. In many ways, progress in these areas is also limited by the ability to handle high-volume data. They will also benefit from the proposed *Teravoxel* infrastructure.

Multi-Echo MRI of Developing Brains

The Caltech Biological Imaging Center collects magnetic resonance (MR) data sets of mouse, quail and mouse lemur embryos, from gestation Day 7 to neonates, at approximately half-day intervals, to understand the relationship between structure and function in developing brains. The MR data are multi-echo (typically 3 echoes) and in some cases utilize several recycle times, along with gradient echo images. This produces multi-D volume data sets, with 3 to 10 values per voxel. The data sets may be as large as $512^2 \cdot 256$, with values stored as floating point numbers, *i.e.*, over 2GB, each.

Electron Tomography of Neuronal Structures

The National Center for Microscopy and Imaging Research (NCMIR) currently utilizes high-voltage electron microscopy (HVEM) to view fine features of neuronal spiny dendrites and organelles, such as the Golgi apparatus and endoplasmic reticulum. The problems with such data has been analyzing and understanding complex 3-D geometries, and extraction of quantitative information. These were recently addressed through computational methods to derive 3-D data from HVEM data sets, using electron tomography. In electron tomography, as with other tomographic medical imaging methods, different specimen views are used to create a 3-D reconstruction amenable to quantitative analysis. In the past, the production of such tomographic volumes was prohibitively time-consuming. With the increase in computing power and the development of digital cameras, an entire tilt series can be acquired and aligned in 3 hours, or less. Using parallelized reconstruction algorithms, a 400MB volume can be computed in less than 30 minutes. With these advances in electron tomography it is possible to acquire and generate 10GB/week with the NCMIR electron microscope.

Large Arrays for Radio Astronomy (NSF), Deep-Space Communications (NASA), and Search for Extra-Terrestrial Intelligence (SETI Institute).

Revolutionary developments in radio astronomy and deep-space communications are spurred by continuing receiver/antenna price reductions, courtesy the satellite television industry. Traditional radio astronomy relies on large, 50-m diameter, or so, parabolic dishes, coupled to hand-assembled low-noise receivers. 50-m dish installations cost in the range of \$50M, each. 5-m diameter installations, however, cost \$20K/ea and can be arrayed to form large synthetic apertures. 1000 5-m dishes would have 10 times the collecting area of a 50-m dish, at comparable cost. However, an array produces large data storage, processing, and communications requirements; 1000 antennas

might produce data rates of order 1Gb/s, each. Image formation requires pair-wise correlations, or, $5 \cdot 10^5$ correlations for a 10^3 -dish array.

2. Proposed research instrumentation

The discussion below assumes a 3-year, two-phase effort. Phases I & II are scheduled as 18-month periods. The proposed *Teravoxel* system is depicted in Fig. 5. It is generic in its architecture and designed to serve a host of applications and needs. What makes it unique is the volume of data that can be handled and rates they can be transported. Its various components are discussed below. Briefly, the envisaged lab/field data acquisition subsystem will be a Phase-II implementation of the KFS data-acquisition system. Similar data will be generated via numerical simulation, as discussed above. Following calibrations that must be applied to laboratory/field data, both experimental and simulation data will form the input stream for common subsequent analysis, and visualization.

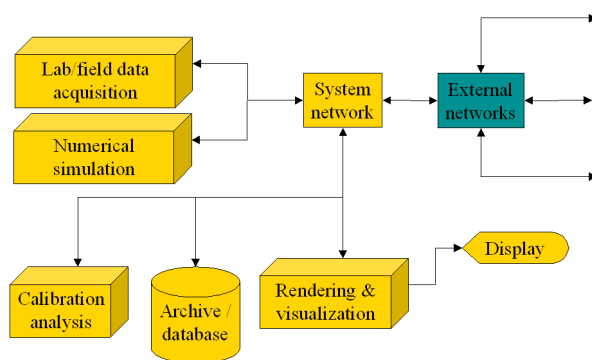


Fig. 5 Proposed *Teravoxel* system.

The network forms the system backbone, whose data-transport needs tax available technologies (transporting 1TB would require 10hrs at 30MB/s). The network interconnects I/O data to system storage, transports data for calibration processing, analysis, geometry and other feature extraction, theoretical and numerical data validation, model testing, *etc.* It also transports data to/from the rendering and visualization system components. Combining adequate speed/size access for storage *and* visualization is a crucial element of the proposed system, even though slightly beyond off-the-shelf implementations at present at the required

transport rates. Part of the proposed development effort is to perform this networking/access integration. The alternative would be represent a costly redundancy.

2.1 Front-end description/specifications

2.1.1 Phase-I KFS system specifications

Much data analysis, such as isosurface extraction (*e.g.*, Figs. 1 & 2), is tantamount to differentiation. Especially after further bit losses to calibration and other corrections, typical 8-bit/pixel image data do not possess adequate dynamic range. In addition to framing rate, this requirement has precluded use of commercial systems and the development of advanced digital-imaging technologies at Caltech (both the in-house Cassini and KFS systems yield 12-bits/pixel).

The Phase-I KFS evolved from the Cassini imaging system, in turn based on the selected second-best CCD built for the NASA/JPL Cassini mission (Fig. 1 data were acquired with it). Development of the KFS system began 5 years ago (DOD funding), following the Cassini imaging system lab/field use release. Presently nearing completion, the KFS system relies on the custom-designed, 1024^2 -pixel, low-noise KFS-CCD, with 32 high-bandwidth output channels for up to 1000-frame/s operation (1Gpixel/s), keeping the 12-bit conversion/channel rate relatively low (≈ 40 MHz). This

mitigates electronic/readout signal-to-noise ratio (SNR) degradation from high signal bandwidths. Data for Fig. 2 were acquired using the Cassini camera and read by the first completed KFS A/D board. Data for Fig. 4 comprises the first integrated KFS system bench test (December 1999).

KFS-CCD data are acquired *via* 8 VXI 4-channel A/D converter boards. The specifications of each VXI A/D board, Model 2 (VXIADC2), are as follows (see Fig. 6):

1. 4 analog-input channels/board. Programmable gain, offset, and low-pass filters, reset clamp logic for correlated double sampling (CDS) per channel (CDS doubles the effective sampling rate to 80MHz but removes per-pixel CCD reset noise for a substantial net SNR gain).
2. 4 12-bit, 40MHz A/D converters (one per channel). At lower data rates, programmable variable A/D-reference voltages, with (on-board) multiple A/D-conversion averaging, yield greater precision (16 conversions for 14-bit accuracy, *etc.*).
3. 4 high-speed digital signal processors (one per channel) capable of full-speed, lossless data compression. Compression ratios of 1.5-3:1, depending on signal level, noise, and number of bits used (8, 12, or 13-16 with multiple conversions).
4. 0.5GB high-speed DRAM per A/D board. Memory subsystem aggregate bandwidth is 320MB/s (64 bits at 40MHz). This allows full-speed recording of uncompressed data (48 bits at 40MHz), or full-speed double buffering of compressed data with a compression ratio of 1.5:1, or better (96 bits/1.5 = 64 bits at 40MHz).
5. High-speed 64-bit 40MHz (320MB/s) auxiliary I/O port. Daughter board connector provided for various parallel-interface options.
6. High-speed interface to VXI CCD-timing controller allows tight synchronization of KFS-CCD and camera head, A/D boards, lasers, and other laboratory/field equipment.
7. High-speed VXI interface allows 64-bit block data transfers to controlling computer.

In aggregate, 8 A/D boards provide a total of 4GB of DRAM, allowing for full-speed recording of 2700 frames, w/o data compression. With compression, 4000 to 8000 frames can be recorded, depending on compression ratio (1.5-3:1). Data are then offloaded to the local control/ storage computer through the VXI interface at 33MB/s, for storage to local disk, tape, or CD-R's, and/or transported over the network to a larger computer system for back-up storage, processing, analysis, and visualization.

2.1.2 Proposed Phase-II KFS development

We propose to extend the capabilities of the nearly-completed Phase-I KFS system by exercising the auxiliary I/O port feature (#5 above) to permit high-speed parallel transfer using a FibreChannel interface to a local disk array for data storage (Fig. 6). This would allow data acquisition/run beyond the present 4GB limit to ~1TB (10^{12} bytes) by providing local high-speed mass-storage devices. The proposed extension would address many of the additional data-acquisition requirements, for turbulence research (Sec. 1.1.1), ground-based astronomical observations (Sec. 1.1.2), and much other research listed under Sec. 1.2 that would benefit from the unique Phase-II KFS digital-imaging capabilities. It would also provide a design/implementation template for adaptation to high data-volume, non-imaging applications, some of which were mentioned above (Sec. 1.2).

High-speed local mass-storage disk arrays would be too heavy/bulky to be placed on a telescope, or other confined places. The envisaged design would rely on fiber-optic links between the A/D

converter boards and the mass-storage array. A single length of multiple fiber (16- or 24-fiber) cable that is relatively small ($< 0.5\text{m}$ diameter) and up to 1000 feet long can be employed. In previous experiments at the 200m Palomar telescope, such a fiber link was used to control of the telescope-mounted Cassini system from the observation room.

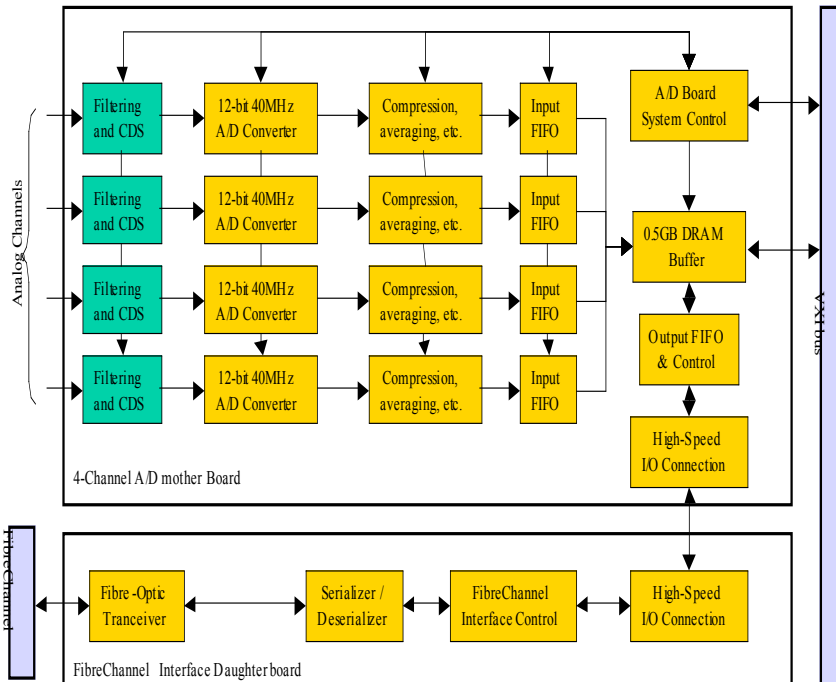


Fig. 6 Proposed Phase-II KFS data system (1 of 8).

A high-speed, FibreChannel (FC) link, was recently developed for mass storage with disk arrays using this technology now available. The current FC implementation has a 1Gb/s ($\sim 80\text{MB/s}$ sustained) transport rate. Each KFS A/D board is capable of 160MB/s (double-buffered mode). One FC link per A/D board will accommodate a maximum of 500frames/s, assuming 12-bit data and a conservative compression ratio of 1.5:1 (higher frame rates at higher compression ratios).

At this time, the fastest drives are the Seagate ST318203FCV (18.2GB) and ST136403FCV (36.4GB), that transfer 18MB/s

to inner tracks and 28MB/s to outer tracks. Six drives are required for the 1Gb/s capacity (current FC implementation) for 500frames/s. For 8 A/D boards, 8 1Gb/s FC disk arrays are the needed (one/board), with 6 drives/array (500frames/s). We propose to employ 48 18GB drives for 0.864TB of local storage, *i.e.*, $\sim 0.85\text{Mframes}$. A 2Gb/s FC implementation is presently under development. In two years, disks should also see a $\cdot 2$ increase in transfer rate. In combination, these will allow the full, 1000frame/s, KFS potential to be realized, with 8 2Gb/s FC links and disk arrays, with 6 drives/array.

The present 33MB/s limit of the VXI bus is acceptable when downloading data from the A/D-board DRAM buffers (4GB total) into local disk/tape storage, but not suitable for terabyte (TB) data transport. Fortunately, FC disk arrays can be dual-ported with one array port to/from an A/D board (8 total) and the other accessed through a dedicated fiber link to the proposed *Teravoxel* system for data downloading, storage, processing, and visualization.

2.2 Networking

Phase-I networking will rely on 100BaseT for transport from lab to the storage and processing center at CACR, as well as ATM links to map drives at lab locations into CACR drives. At 4GB/run, the corresponding transport bandwidths do not represent a serious bottleneck. Phase-II networking needs represent substantial challenges, however. At $\sim 1\text{TB/run}$, data volume is such that reliance on

Phase-I networking technology would represent the bottleneck for data viewing and experiment turnaround. We propose to address these issues by a combination of Gigabit Ethernet as well as (parallel) FibreChannel (FC) technology. The latter holds the most promise as it relies on fiber-optic links, which can span the distances between data-acquisition sites and CACR. Specifically, FC links can either access dual-port local drives (cf. Fig. 6), or connect to the remote (CACR) disk array, as described below. [DR, JP, +?]

2.3 Storage system requirements & development

The Phase-I KFS instrument, at 4GB/run, will produce 3-D/4-D scalar data of order 8GB/run when expanded (byte-unpacked) and uncompressed. Derived data sets (calibrated, 4-D Cartesian-grid-interpolated data, extracted vector-field data, *etc.*) increase data storage requirements to as much as 64GB/run (68GB/run including raw data). Including up to 10 minutes of animation data can add another 90GB. Thus, Phase-I data storage requirements can be expressed as $(68n+90)$ GB, where n is the number of runs. For security and computational efficiency, raw and output data (after the costly calibration/grid-interpolation calculations) will have duplicate tape copies, increasing the requirement by $12n$ GB, *i.e.*, $(80n+90)$ GB total.

Assuming a (max) 20% daily duty cycle over the 18 months of Phase I, we anticipate ~100 runs, leading to an estimated Phase-I total of ~8TB. StorageTek 9840 drives (\$40K media storage costs for 8TB) are proposed for Phase-I. Two drives to be attached to the High Performance Storage System (HPSS), with a third to be attached to the instrument to offload data when the Caltech network is not accessible (off-site use, *e.g.*, Palomar or Mauna Kea telescopes, *etc.*). The tape drives cost \$30K/ea, *i.e.*, \$90K for Phase-I tape drive costs, total. [Need to get official quotes on this, \$27K came from a JPL quote, *i.e.*, \$30K w. tax.] Disks on the Visualization Cluster, as well as the current HPSS disk cache, have sufficient capacity for Phase-I operations.

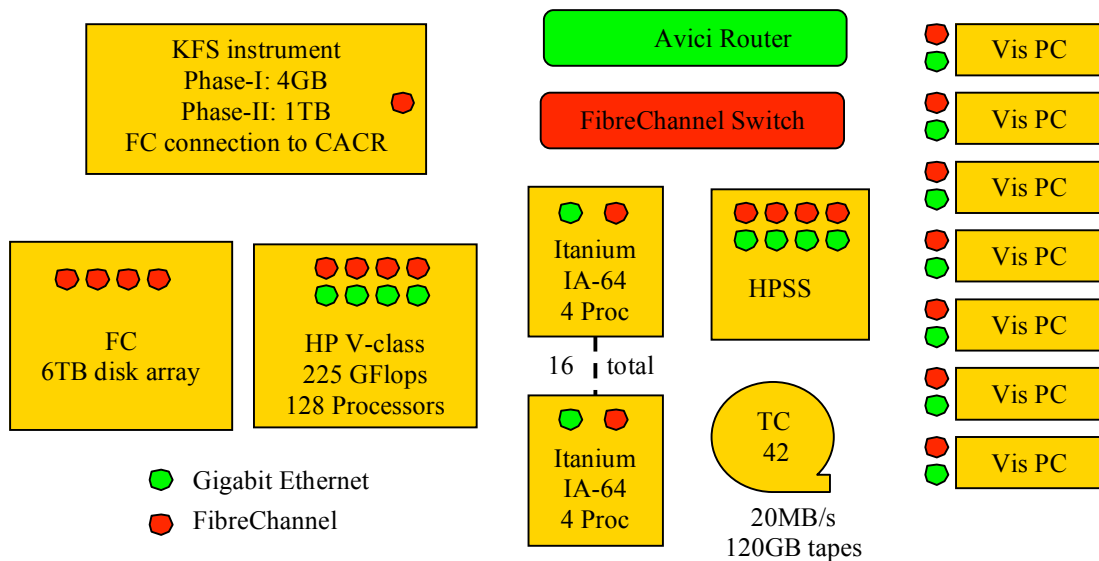


Fig. 7 Storage/analysis/visualization system components/connections.

The proposed Phase-II analysis and visualization systems will be connected to a 6TB FibreChannel (FC) disk array (Fig. 7). A Shared SAN file system (similar to the Mercury Computing SANergy, or the Transoft FibreNetDS) is proposed, so all subsystems are able to read all files. Each subsystem

will be attached to the Storage Area Network *via* a FC interface card and through a FC switch. [Prices for switch equipment] This will provide a 100MB/s transport capability throughout. [Budget-justification pages?] Costs for FC disk arrays at Phase-II start are budgeted at \$60K/TB [consistent with quote to DLang?]. With additional funds for RAID controllers, the disk cache total is estimated to be \$390K.

We are also assuming the StorageTek TC42 tape drive will be available at Phase-II start. Estimated for release 2001Q3, the TC42 will hold 120GB/cartridge (uncompressed) and transfer at 20MB/s. Single-spool cartridges currently cost around \$50/ea, leading to Phase-II storage costs of \$420/TB. Base storage requirements for a maximum-sized run would be the ~1TB of raw instrument data, double copied, plus 2TB for calibrated data, *i.e.*, \$1.68K/run. Storage of computed data sets, such as velocity and vorticity, would add to cost per run. For Phase-II, we project a need for 5 TC42 tape drives; one for the instrument and four for HPSS. The cost of these drives is not known at this time, but assumptions are that StorageTek will keep prices fairly constant. Assume \$30K/drive, this leads to \$150K for Phase-II tape drives.

2.4 Data processing/analysis requirements and system architecture

2.5 Multi-D space/time visualization system development

2.6 User environment

3. Impact of infrastructure projects

4. Project management plans and procedures

4.1 Second KFS camera system

4.2 Storage/use allocation and priorities

4.3 Maintenance, operation, and technical support

5. Budget and funding justification

5.1 FastLane budget forms

5.2 Budget details

5.3 Cost sharing sources

6. Personnel

6.1 Principal investigators and major users

Breen, David
Dimotakis, Paul
Martin, Chris
Meiron, Daniel
Pool, James
Rutledge, David

6.2 System-development

Aivazis, Michael
Brunet, Sharon
Bunn, Julian
Lang, Dan
Lombeyda, Santiago
Patton, James
Sterling, Tom
Williams, Roy