KIPAC Computing Resources and Infrastructure

Stuart Marshall

KIPAC Objectives:

- Address science questions arising at the interface between physics and astronomy: DM, DE, CMB, size, shape, age of the universe, growth of first structures, …
- Bridge the campus and SLAC research communities and to facilitate collaboration between the three federal agencies, DOE, NASA and NSF.
- We are building, from the ground up, the infrastructure that will facilitate the success of these research programs. A substantial portion of that infrastructure is computing capability.
  - Our computing capabilities have just served the needs of the first KIPAC members and warrant significant development if we are to take advantage of our potential.

The Science of KIPAC

- Particle Astrophysics
  - Black Holes, Neutron Stars, White Dwarfs…
  - GRBs, magnetars, supernovae…
  - Accretion disks and jets…
  - Relativistic shocks, particle acceleration, UHECR…
  - Solar Physics
- Cosmology
  - Dark energy, dark matter
  - Gravitational lenses
  - Clusters of galaxies and intergalactic medium
  - Microwave background observations
  - First stars, galaxy formation
  - Supernovae

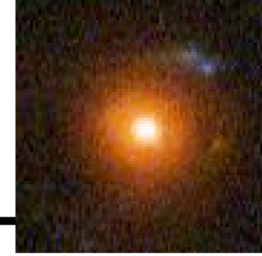
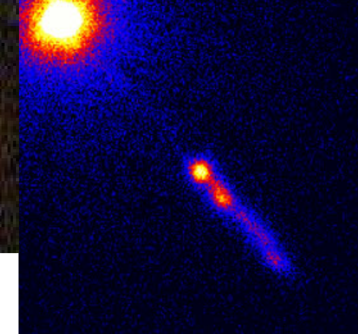
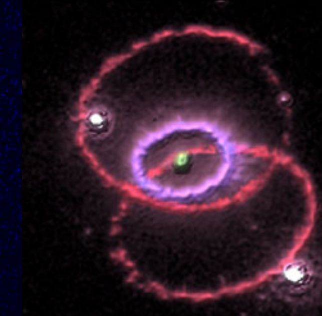

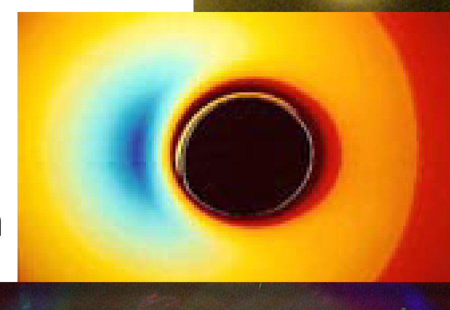
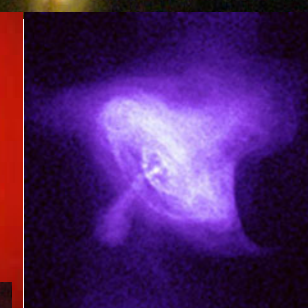
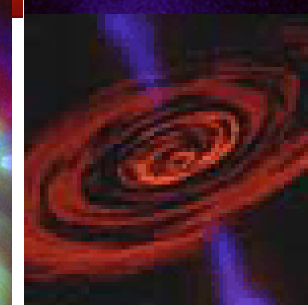


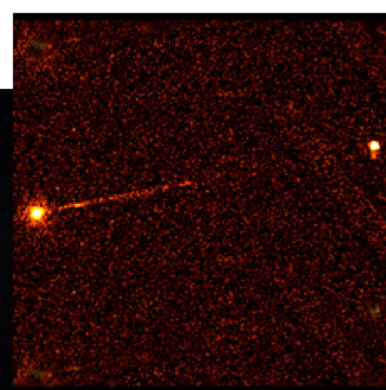
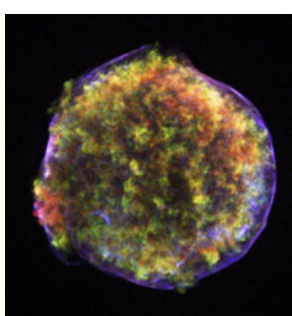
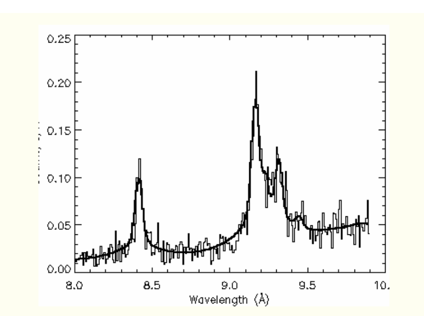

Current Computation: Diverse range of topics

- **Multivariate Monte Carlo Simulations of X-ray Clusters of galaxies (cluster physics, cosmological constraints) (Peterson)**
- **LSST sims: photons from galaxies to detector (Peterson)**
- **Optical and x-ray studies of galaxy clusters: constraining cosmological parameters (Allen)**
- **Neutron stars: pulsar magnetospheres, nuclear burning o rotating NS during type I x-ray bursts (Spitkovsky)**
- **X-ray spectroscopy & atomic processes (Sako, Gu)**
- **Gravitational Lensing (strong, weak) (P. Marshall, Bradac)**
- **Formation of first objects in Universe: (Abel)**
  - **Stars, dwarf galaxies, supermassive black holes**
- **Gamma Ray Burst modeling: (Zhang, Granot)**
  - **Relativistic jet, accretion disk simulations, prediction of observables**
- **Super symmetric particle physics models and observables: comparison with data (Baltz)**

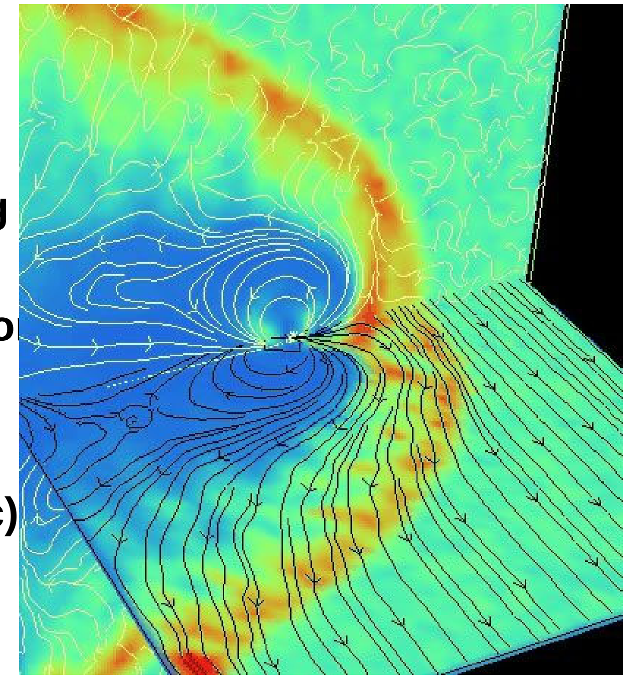

Anatoly Spitkovsky: Relavitistic wind interacting with a rotating pulsar magnetosphere. Used to model recently discovered binary pulsar system J0737.

Current Computation: Scope

- **Range of computing scales:**
  - **Interactive small jobs: 20-30 processors 24/7 on desktop and batch systems**
  - **Medium scale: <100 processors**
    - Global shared memory (SGI) for quick algorithm development
    - Distributed shared memory (MPI on 32, 128p at SLAC)
    - Fully independent (up to ~3000 processors at SLAC)
  - **Large scale: >128p: Big codes (ENZO, RAM, TRISTAN) on big machines @ (NERSC, ORNL, NCSA, SDSC, *SLAC*?)**
  - **DOE Huge Memory architecture prototype (future uses)**
- **Visualization facility: prototype 3D projection display, software**
- **Software: idl, mathematica, iraf, aips, heasoft, pgplot, ciao and many personal/group packages. To date, management is pretty ad-hoc.**

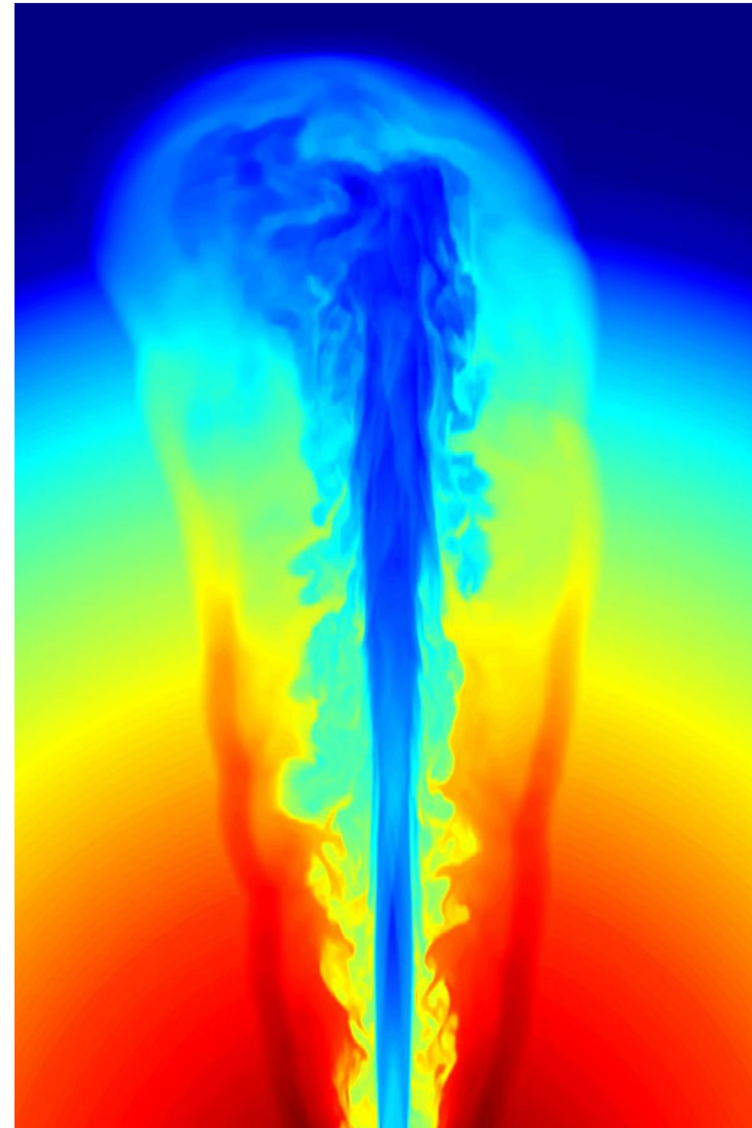

Current Computation: Resources

- **14 Apple XServe G5s (2 FS, 2 interactive, 10 batch, 9.6 TB Storage**
- **SGI Altix w/72p 1.5GHz, 220 GB memory, 8 TB Storage**
- **21 *node dual cpu opteron for XOC group***
- ***15 TB NFS Storage on Sun systems***

HPC on SGI Altix

- Large,parallel runs on Red with heavy Disk I/O as well as message passing usage,using the code ENZO (75k lines,written by Greg Bryan)
- Timesteps,of which there can be tens or hundreds,are typically between 2 and 8 gigabytes,and consisting of thousands to tens of thousands of files.
- Small runs use between 4-16 processors, larger runs (cosmology) use 32 typically. More than 32 is considered inconsiderate to fellow users. Users manage processor allocation.
- Run times can be between one day and several weeks,depending on start and stop times and complexity of the run.
- Simulation complexity can be very difficult to estimate in advance,and until the actual collapse of primordial objects in the simulation,it can be impossible to estimate when simulations will become complex.

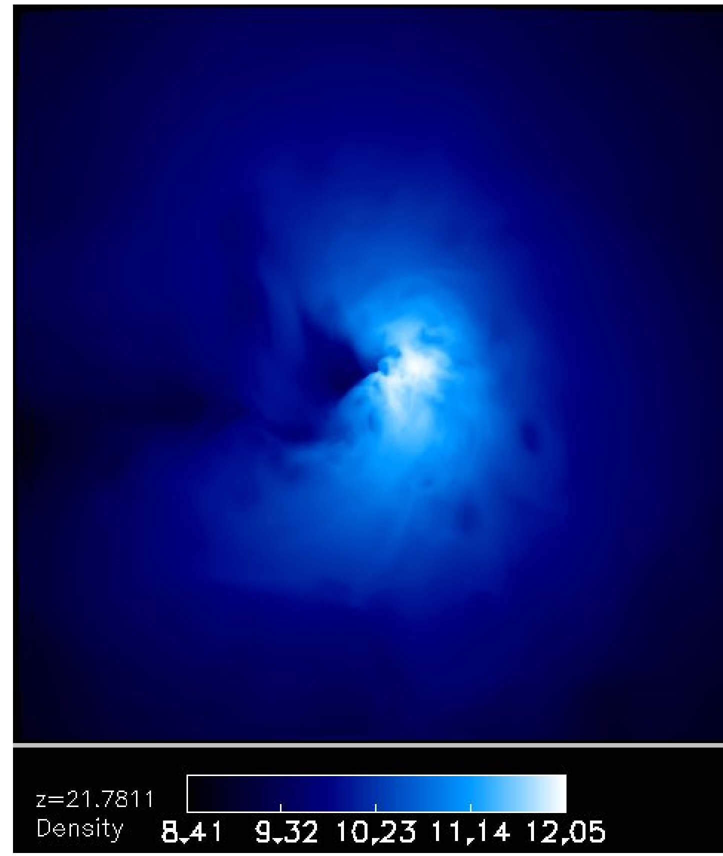

Matt Turk

Visualization: Prototypes

- **Small 3-D projection system**
- **Rendering usingVisIt, developed LLNL (http://www.llnl.gov/visit/)**
- **Based on VTK, with generalized parallel support for analysis and rendering**
- **Can render on Altix or new 4 core Sun visualization host.**
- **Currently forwarding X11,will soon be usi full client/server architecture to forward OpenGL primitives through SSH directly**
- **Stereoscopic rendering, with VisIt as well as particle data in partiview (from NCSA)**
- **Moving analysis tools to VisIt, for generalized parallel architecture support**
- **Apple G5 based multi-screen display**

Visualization: Research and Development

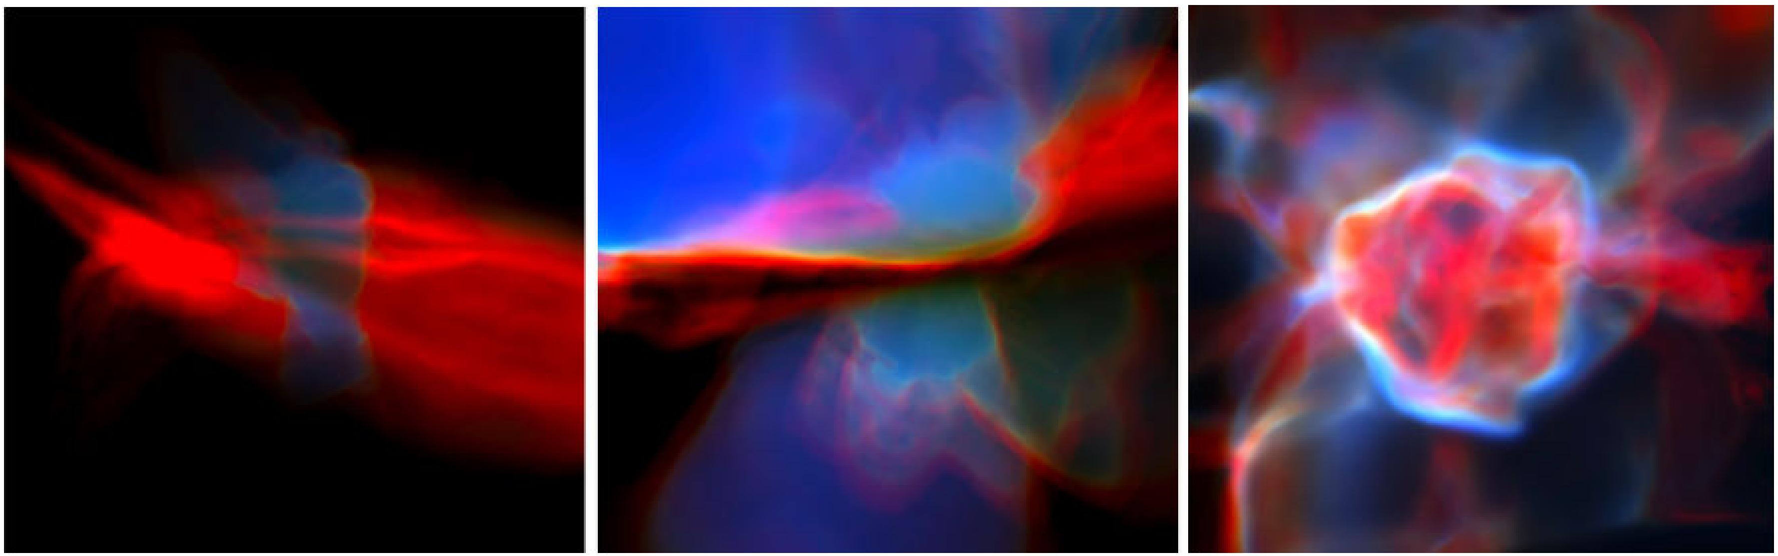

- Three time-steps from a galaxy and star formation simulation rendered with the optical model that maps density and temperature fields to emission and absorption coefficients.
- dark regions in the center result from non-illuminated gas in the front, that absorbs large amounts of light in that region.
- Images were created with a new volume renderer developed by Ralf Kaehler (ZIB), John Wise & Tom Abel (KIPAC).
- Ray traces thousands of grid patches on different spatial resolution and composites them consistently.
- Gives rates of 10 frames per second for a complicated data set which used to take minutes per frame with a software based approach.

What’s to come

- **Black Hole coalescence (Andres Escala*)**
  High resolution SPH simulations of inspiral and merger of MBH binaries. Code is gadget w/MPI.
- **Black Hole Galaxy Connection (Stelios Kazantzidis)**
  Multiscale N-Body+SPH simulations of merging galaxies w/SMBHs at high spatial resolution. Code is “gasoline” SPH optimized for MPI scaling to many processors.
- **Re-ionization of the Universe (Marcelo Alvarez*)**
  Large scale radiative transfer simulations of cosmic reionization. Code C^2-Ray runs on SMP and MPI architectures.
- **Radiation transfer in curved spacetime (Steven Fuerst)**
  General formulation for relativistic effects in photon transport and motion of emitters. Code suitable for larger scale deployment on shared memory machines.

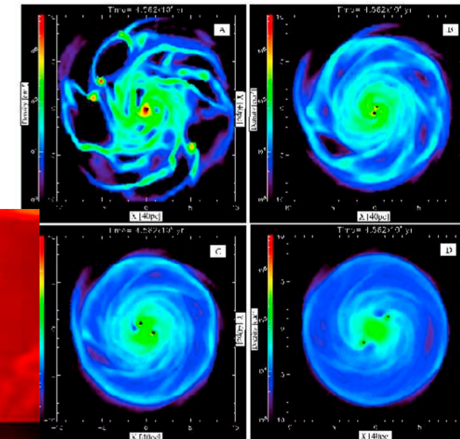

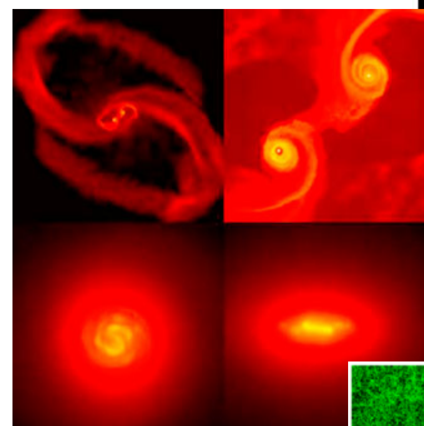

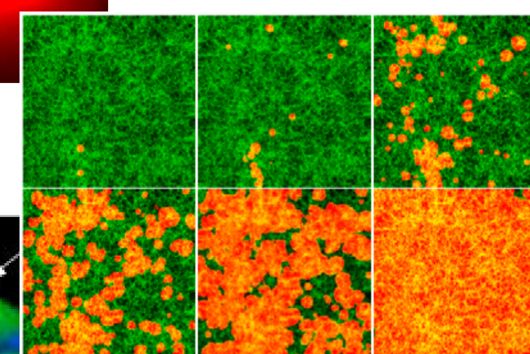

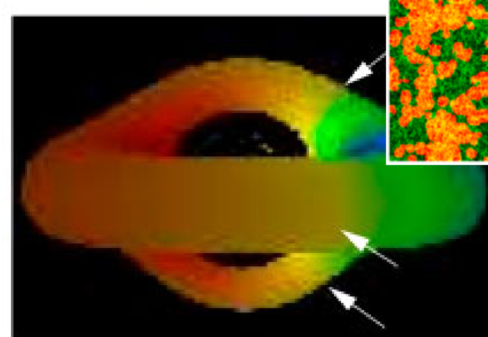

Looking Forward, what is needed

- **Basic services:**
  - **network (FKB, campus)**
  - **Accounts (~100 in KI group)**
- **Individual computing: desktop and laptop**
  - **~85 assigned seats in FKB, over 100 during summer**
- **Communication and collaboration**
- **Software**
  - **Commercial: IDL, Mathematica, Matlab, …**
  - **Community and OSS: IRAF, Heasoft, Imcat, SM, HippoDraw, skycat, starlink, pgplot, partiview, mysql, …**
  - **$30K+ this year, growth dependent on mix of comm/oss**

Communication

KIPAC mailing lists are an important tool for keeping our community informed. Admins manage addition/deletion of visitors/students via web interface.

kipac.stanford.edu

Welcome!

Below is a listing of all the public mailing lists on kipac.stanford.edu. Click on a list name to get more information about th information page for an unadvertised list, open a URL similar to this one, but with a '/' and the list name appended.

List administrators, you can visit the list admin overview page to find the management interface for your list.

If you are having trouble using the lists, please contact mailman@kipac.stanford.edu.

List	Description
admin	KIPAC admin mailing list for announcements/discussion
affiliates	KIPAC Affiliates list for announcements
associates	KIPAC Associates list for announcements
everyone	The union of all KIPAC lists -- for broadcast messages.
faculty	KIPAC faculty mailing list for announcements/discussion
fkb-residents	Fred Kavli Building Resident mailing list
friends	KIPAC Friends list for announcements
full-members	Mailing list for KIPAC full members
lsst	SLAC LSST working group list
nerds	KIPAC mailing list for discussion of software, computing, etc.
postdocs	KIPAC postdoc mailing list for announcements/discussion
scs	SCS/KIPAC discussion mailing list
slackers	KIPAC recreational mailing list
staff	KIPAC scientific staff mailing list for announcements/discussion
students	KIPAC student mailing list for announcements/discussion
visitors	KIPAC visitors list for announcements
xray	KIPAC X-ray group mailing list

Delivered by Mailman
version 2.1.5

Python Powered

Collaboration

Collaboration web space based on open source package Plone.

Prototype set up and tested for several user groups
Next step is to install new instance on dedicated server
Continue discussions with SCCS to similar capabilities hosted at SLAC

Supporting KIPAC Research: Services needed

- **Storage and data management**
  - Not unusual for PD’s and research groups to have multi-TB storage requirements
  - We are purchasing ~30 TB this year with expected constant $$ need for storage
  - “Data Lifecycle Management”
- **Computing cycles: simple to full HPC**
  - Interactive servers
  - Large memory/processor w/few cpu’s
  - MPI cluster 64 nodes w/4 cores
  - Continued SMP development
- **Scientific visualization**
  - Large pixel display w/software
  - 3-D display w/software development

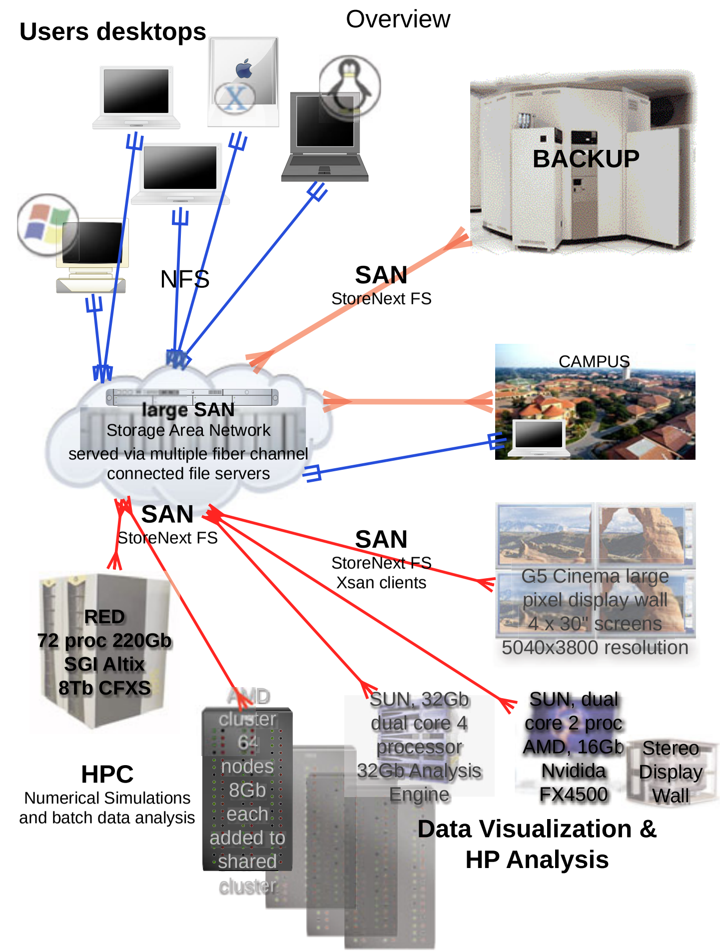

Supporting KIPAC Research: People

- **Staff: Marshall, Kunz, Morris, ..**
  - **With completion of FKB we are just getting together as a team. Kunz brings deep experience at SLAC as physicist and developer. Morris has expertise in modern codes for Cosmological simulations and data analysis.**
  - **New Position opening to support design, implementation, and development of SOA visualization for KIPAC.**
- **KIPAC/SCCS collaboration**
  - **Fortnightly meetings for coordination and development of KIPAC computing resources.**
  - **Major milestones:**
    - Integration of small Apple Xserve cluster and fileservers
    - 72p SGI Altix w/220 GB memory
    - Many publications due to this effort – which was large!
  - **Ongoing work:**
    - Search for suitable networked storage system (Xsan, GPFS, …)
    - Implementation of automated backup via IBM's TSM on SLAC silos
    - Possible use of SLAC batch system for MPI jobs
    - Costing and planning for new hardware
  - **An acknowledgement:**

    SCCS has provided major efforts to support new hardware and operating systems that have in fact enabled much of the successful results KIPAC researchers have published. Continued support for scientific computing by SCCS is critical!

Long Term: Computational Challenges & Development

- **High-performance computing development platform**: Current 72 processor, 220 Gb, SGI Altix SMP is a resource that is useful for development of new numerical techniques, optimization and also some production runs. Maintaining this type of resource is desired.
- **Dedicated batch computing**: More resources with higher interconnect speeds would find much usage.
- **Data and file space management:** Progress in this area would benefit broad range of SCCS users.
- **Scientific Visualization**: Scientific visualization enables human interaction with the vast amount of data generated by numerical simulations and
- **Personal computing environment and training:** Expand current desktop environments to have a well maintained software and development environment tailored to the experimental, observational and theoretical users and resources to train them in its effective use.
- **Information and project management**: Deploy modern tools to document and communicate KIPAC's efforts internally and publicly as well as enable national and international collaborative project management.
- Continuing and expanding our ongoing collaboration efforts with SCCS.
- Collaborating with other groups at SLAC with similar needs and interests must be greatly increased!

People	Year	Desk	Storage	Software	Cycles	Visualization
75	2006	94	10TB FNAS + 30TB NAS	35	128n8Gb	8d + s3Dstereo
80	2007	100	10TB FNAS + 30TB NAS	40	1	8x8 * display wall +install
85	2008	106	60TB BA + 20TB FNAS + 30TB	45	SMP	3 tile 3d stereo
90	2009	113	90TB BA + 20TB FNAS + 30TB	45	3	new GPUS + active stereo system
95	2010	119	120TB BA + 40TB FNAS + 60TB	45	3	visualization workstation
100	2011	125	180TB BA + 40TB FNAS + 90TB	50	5	display wall upgrade