Atlas Great Lakes Tier 2 Computing Center AGL-T2

UM & MSU

This is a discussion of a proposed computing center at MSU for analysis of data and simulation for the CERN Large Hadron Collider (LHC) experiment, “ATLAS.” It would make direct use–the first research use–of the newly installed high speed network in our state called the Michigan LambdaRail (MiLR).

This Center would be a collaboration of the High Energy Physics groups at MSU and UM.

While we have been collaborators for decades at Fermilab, CERN, BNL, and ANL...this would be the first global cooperative effort involving essentially all of both groups.

We believe it is fitting that the first statewide use of MiLR be a tightly coupled collaboration between our two universities.

MSU and UM have been involved

for many decades in HEP research

continuously supported by NSF and DOE projects

currently 8 MSU

and 15 UM experimental faculty

Every major university in the world has two major subfields in physics research: Condensed Matter Physics (CMP) and High Energy Physics (HEP).

The HEP groups at MSU and UM have been very successful over the last 3 decades in designing and building of highly technical equipment for use in our experiments at international laboratories; for coaxing important scientific results from these large experiments with graduate students and research associates; and in assuming positions of responsibility in the management of these international collaborations and advising the U.S. government on scientific policy.

We are supported in an ongoing basis by three operating grants from the NSF (MSU and UM) and DOE (UM). This support has been uninterrupted for since its beginnings in the 1970's. In addition, we have both been successful in garnishing project funds from the DOE for the construction of large components of our experiments.

High Energy Physicists

(aka "HEP")

We work to unravel the processes

which controlled the evolution of the universe $<10^{-8}$ s after its birth

High Energy Physics (HEP) is the study of the fundamental entities and forces in Nature.

We make use of international facilities called "particle accelerators" in the U.S., Switzerland, Germany, Japan, and China in order to produce and collide protons, antiprotons, electrons, positrons, neutrinos, and other particles...

The energy-densities of these collisions are similar to the first pico seconds or so of the birth of the universe and any objects that would have been present at that time, would be produced by us, artificially at these facilities. Sometimes, this field is also called "Elementary Particle Physics," since we seek to unravel the ways in which energy manifests itself in quantum mechanical bundles which have lifetimes, inertias, and particular ways of interacting with other different bundles. The name "particles" is a hold-over from the colloquial use of that word and the fact that many of the say dynamical laws govern the collisions (albeit, relativistic laws) of these "particles" as for actual matter-particles. "Elementary" refers to our belief that these are the most elementary, indivisible, entities in nature.

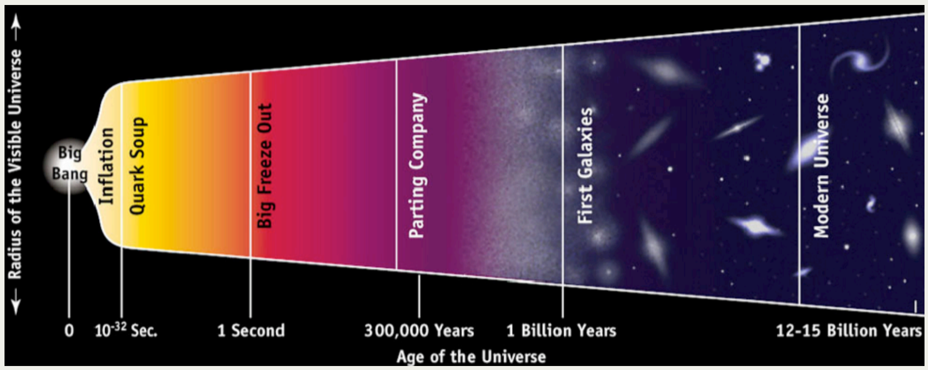

The evolution of the universe can be thought of as a time-evolution, since t = 0, the Big Bang. But, it can also be thought of as an orderly cooling as the universe expands and the initially hot byproducts are dispersed. Modern particle physics laboratories regularly probe the region in the figure around “Orange.”

$$E = mc^2$$

&

beam “luminosity”

the calling cards to fundamental physics

the higher the energy, the heavier the particle masses that can be created

the more particles which collide, the rarer the processes that can be studied

The higher the energies available, the higher are the massive, elementary states that can be excited. This excitation is in the head-to-head collisions of very energetic beams of typically protons or electrons with their antimatter partners, anti-protons and positrons.

Usually, the probability of producing states of matter becomes smaller as the masses become larger, so in order to study them we must have more and more collisions between the primary beams.

Both of these requirements run into technical limitations as in order to accelerate high energy particles, high-field superconducting magnets are required to direct the beams, but this is a technically difficult job and so it makes these facilities geographically large.

Everything about HEP is big...

Typically, circular or linear particle accelerators are measured in miles around or in length.

Collaborations are now regularly 500-1000 PhD’s.

Future collaborations are 1500 PhD’s.

Costs and commitments are national and negotiated at the level of State Departments. Obligations extend over decades.

As a result

High Energy Physics requires enormous facilities

We are engaged at two international facilities:

The Fermi National Accelerator Laboratory, Fermilab

The European Center for Nuclear Research, CERN

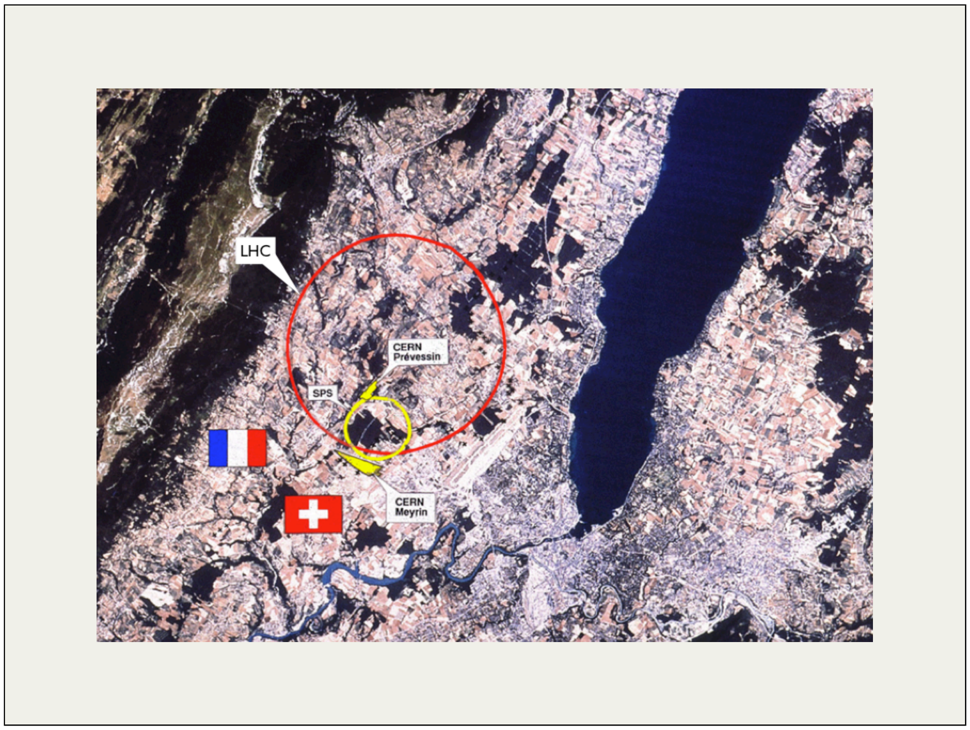

The Large Hadron Collider is the most recent of these increasingly mammoth, international enterprises. It straddles the border between France and Switzerland, just outside of Geneva.

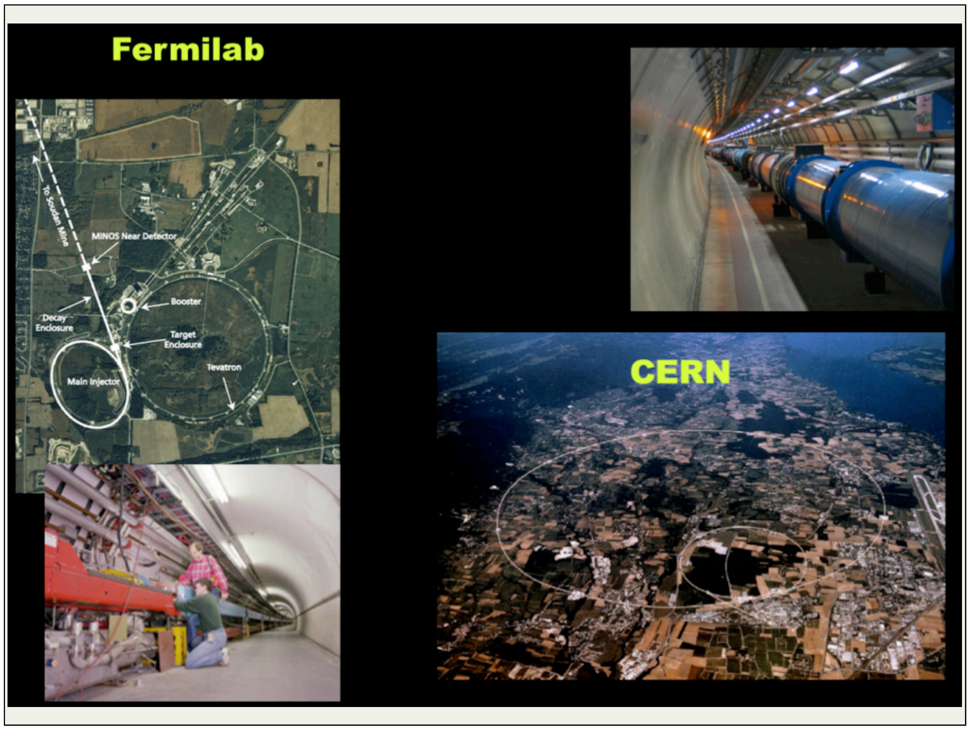

The fermilab complex is in Batavia, Illinois and houses the DØ and CDF collider detectors. This is where we work now, and will continue until about 2008. We currently have 5 graduate students and 4 Research Associates in residence in Illinois.

UM and MSU have been involved in these experiments, fulfilling leadership roles in the physics, engineering, and management of the collaborations.

The CERN accelerator complex is now called the Large Hadron Collider (LHC) in Switzerland. UM and MSU have been involved in experiments at CERN off-and-on for two decades , but most recently for the last nearly 10 years as the slow, complicated ramp-up to a new proton-proton collider facility commences.

The Fermiab ring is 4 miles in circumference, while the CERN LHC ring is 17 miles in circumference: need that for the highest energies with beams at 7 times the momentum of Fermilab's.

The LHC Is still under construction and scheduled for physics quality beams in late 2007 or early 2008.

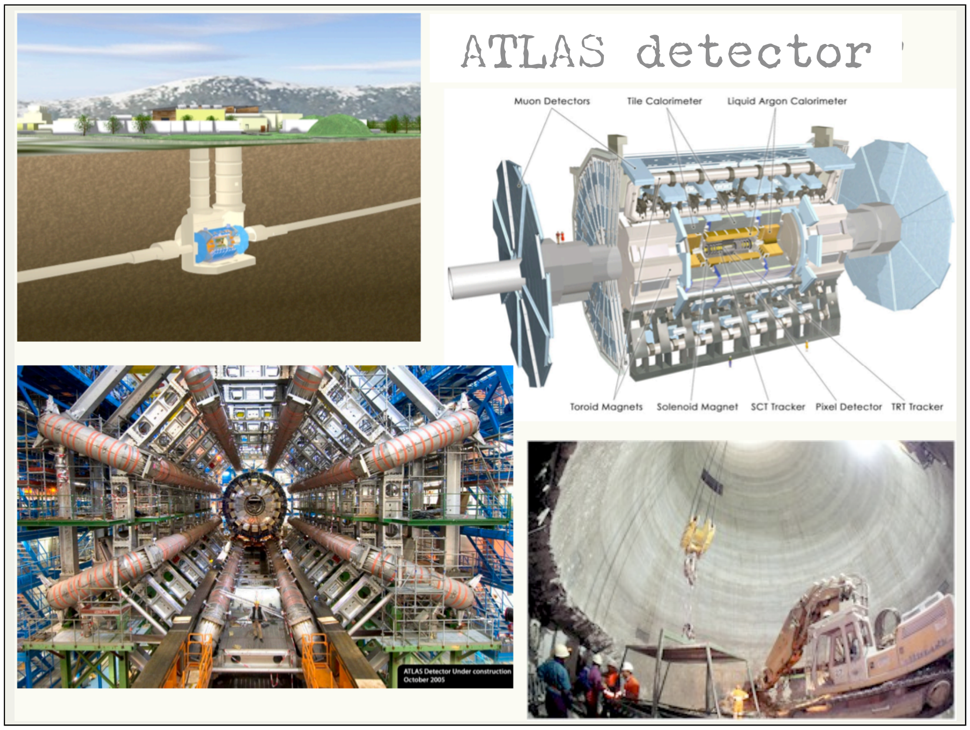

There are a number of detectors built around the proton-ring at the LHC. There are two which are traditional (except for the size) HEP experiments: the ATLAS (A Toroidal LHC ApparatuS) Experiment and the CMS (Compact Muon Solenoid) Experiment. They are both undergoing the final year of installation deep underground at two opposite points in the ring....300 feet underground.

ATLAS involves 2,000 scientists and engineers at 151 institutions in 34 countries. MSU and UM have been members of ATLAS since its inception, following the cancellation in the U.S. of the Superconducting Super Collider in Texas.

units

network bandwidth

Mb/s - millions of bits per second ~ your office is probably 100Mb/s

we'll deal in billion (Giga) bits per second, typically 10Gb/s

computer processing power

"SpecInt2000" (SI2k) standards - your PC ~ 800-900 SI2k units

we'll deal in thousands of -SI2k to millions-SI2k quantities

storage capacity

giga Bytes (GB) - billion-byte - typical consumer hard drive ~100GB

we'll deal in terra or peta Bytes (TB, PB) - trillion and …whatever comes after trillion!

In what follows, we'' need to refer to units of network bandwidth, computer processing power, and data storage. In general, the amounts of these quantities are 2-5 orders of magnitude more than standard quantities.

Remember,…big.

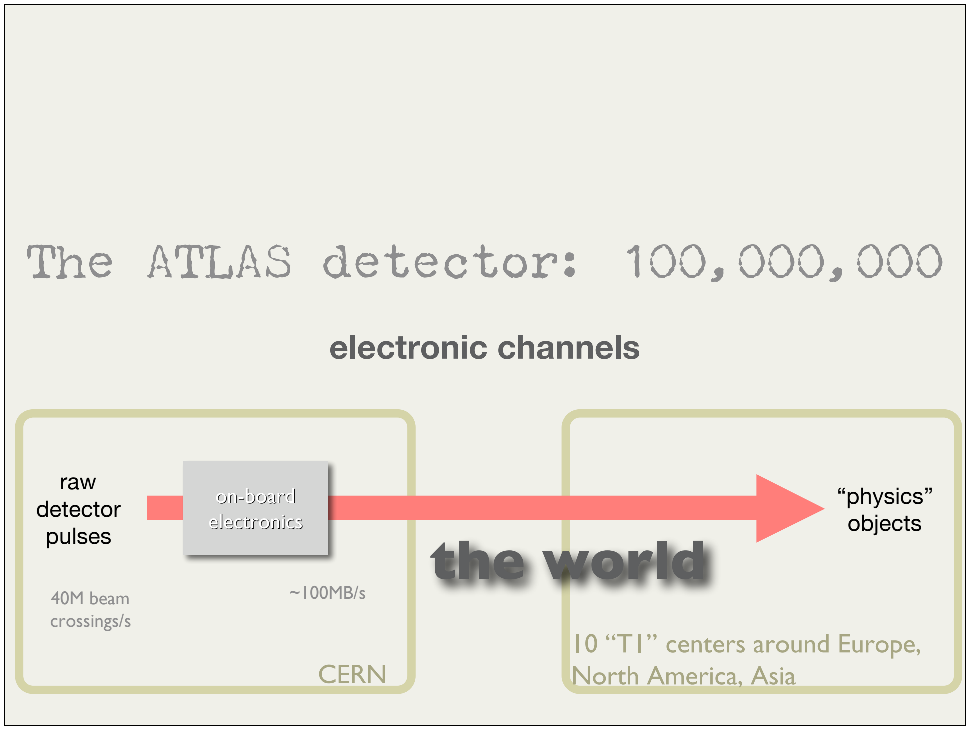

The amazing thing about these detectors is not only their sheer physical size, but their electronic “size.” The ATLAS detector has a tenth of a billion individual electronic channels, and the means of reading them out and forming their raw information into representations of “physics objects” such as electrons, protons, neutrons, muons, and so on, is an enormous near-real-time processing task.

Physics analysis must be done on highly filtered and pre-processed information. CERN starts that task, but the whole world - 3 continents - is requried to complete the job so that individual university scientists (like us here) can extract the useful information.

The data sets which will be analyzed by thousands of physicists must be prepared at 10 different laboratories and available for everyone over the grid that has been running in “data-challenge” mode for a number of years now. The data processing and data storage requirements are staggering.

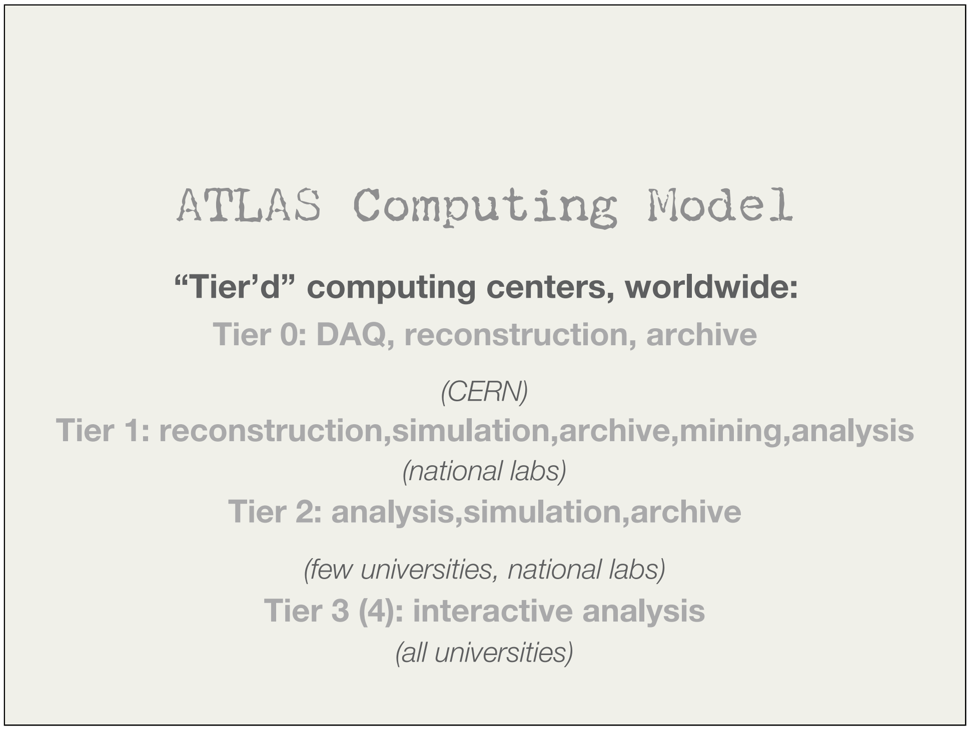

The “computing model” for ATLAS, then, involves successive stages of increasingly recognizable information:

CERN Tier 0: Exactly 1 exists: Record/Archive Raw Data, Prompt Reconstruction, Distribute Raw, ESD, etc.

Worldwide Tier 1 Centers: ~10 (capacities are defined in international MoUs): Archive, perform post calibration and any re-reconstruction of a share of the raw and simulated raw data.

Worldwide Tier 2 Centers: ~30 (5 in US, aggregate capacity defined by MoU). Bulk of simulation, analysis support for ~20 active users/site.

Worldwide Tier 3 Centers: 100’s Institutional Facilities & Individual Users. Primarily physics analysis.

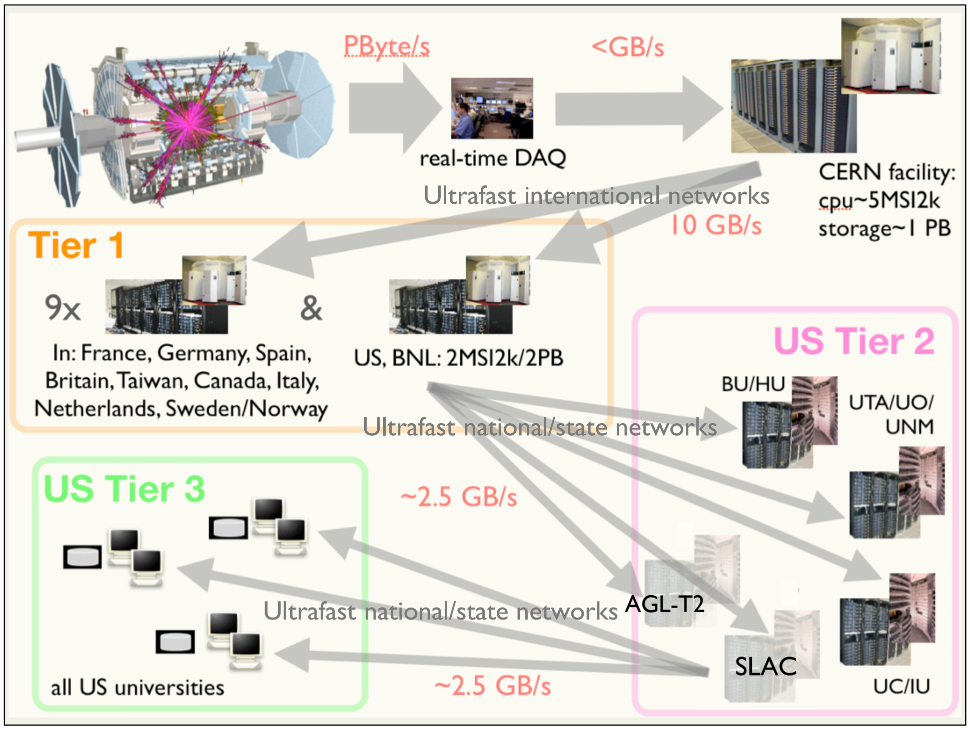

The processing is done in stages from the detector, through to the actual users at their individual institutions and Linux workstations. The processing leads to successively smaller and more manageable data formats and what the user will likely analyze will be event records of ~few 10s of kB or so. There will be many millions of them, depending on the physics reaction of interest. Along with them will be many times the interesting signal process in events which represent the various backgrounds. Along with both of those kinds of data will be *simulations* of each, typically 10's of times the number of events in the real data sets. These simulations are a typical Tier 2 job and will be done many times over and over, as more and more is learned about how best (and fairly) to mimic the apparatus in software.

An individual scientist, in order to analyze an entire data set during the mature years may require 100's GBs to PB's of data access.

This, for 100's of physicists, all at the same time.

Call for proposals

Last spring for last 2 T2 Centers

- **University of Michigan and Michigan State joined together…**

Approved last summer:

Atlas Great Lakes Tier 2 Computing Center

PI and Co-Director of the UM Center: Dr Shawn McKee, University of Michigan

Co-Director of the MSU Center: Prof. Raymond Brock, Michigan State University

ATLAS T2 Status, 2007

Five Tier 2 U.S. sites are required

Three were approved 2005:

Boston University & Harvard University

University of Texas, Arlington, University of Oklahoma, Hampton University, & University of New Mexico

University of Chicago & Indiana University

Two were approved 2006:

Stanford Linear Accelerator Center (SLAC)

University of Michigan and Michigan State University (AGL-T2)

Award: $600k/year for the "life of the research program" ≥ 10 years

A competition was completed about 1.5 years ago and resulted in three T2 sites being selected. MSU was not in that competition, but UM was and was not successful. (At that time, we were not yet involved in high-density Linux computing.) We all realized that in order to be successful that a joint university initiative was required and (quite late) it was decided that a State of Michigan proposal was appropriate and could be quite strong.

The funding agency responsible is ostensibly NSF. However, in the U.S. day-to-day and even year-to-year responsibility for managing the ATLAS experiment has been delegated to a central facility at Brookhaven National Laboratory, "U.S. ATLAS." It has in its budget funds for 5 T2 centers and the RFP came from them and they will make the awards, with NSF consent.

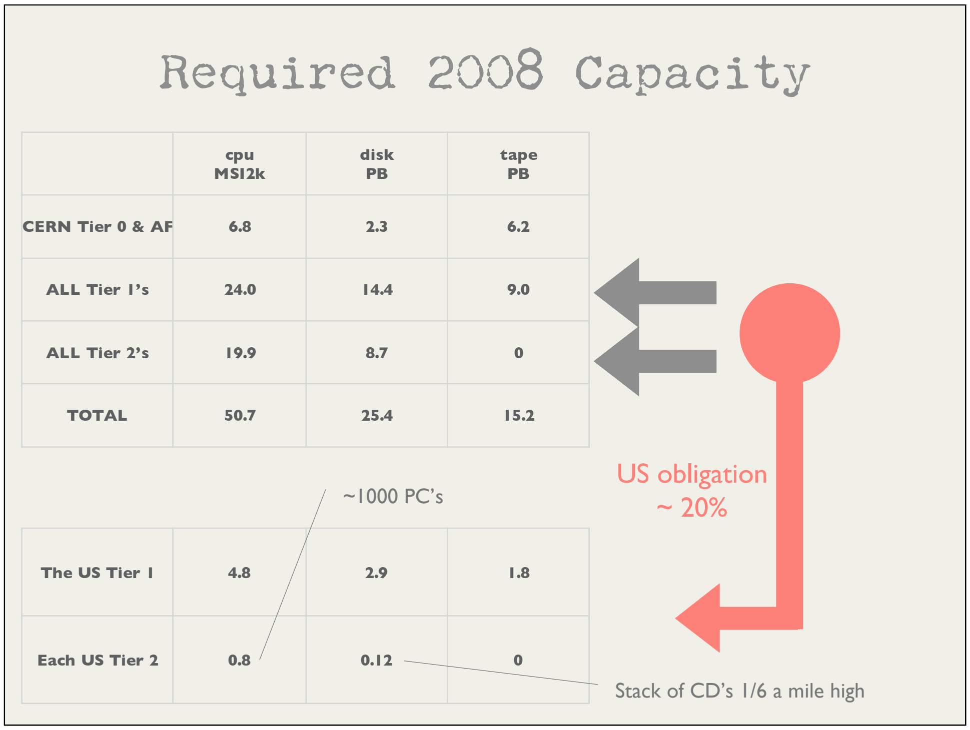

	cpu MSI2k	disk PB	tape PB
CERN Tier 0 & AF	6.8	2.3	6.2
ALL Tier 1's	24.0	14.4	9.0
ALL Tier 2's	19.9	8.7	0
TOTAL	50.7	25.4	15.2

The US Tier 1	4.8	2.9	1.8
Each US Tier 2	0.8	0.12	0

The ATLAS experiment has a computing model that requires minimum resources at each stage of Tiering. The sum total capabilities are in the upper table, while the U.S. obligation to support roughly 20% of the whole results in per-Center requirements in the bottom table.

Planned AGL-T2 2009 Capacity

~2500 PC's

	cpu MSI2k	disk PB	tape PB
AGL-T2	~2	0.7	0

Stack of CD's about 1 mile high

MSU-UM Proposal

Four distinct components:

1. cpu availability
2. RAID storage
3. network capability
4. technical personnel

From 3 defined sources:

A. university-dedicated
B. university-shared
C. project funds

Our proposal will address specifically four components which will be offered to US ATLAS members. There will be provision for university dedicated resources in those four areas, any shared resources that might be occasionally available, and expectations for how the award would be allocated.

we bring now

considerable expertise

high-density commodity Linux cluster development

real-time coding/scripting

physics simulation, electronic trigger design, and detector construction/calibration

large project management

considerable group resources

We are already capable in a number of relevant areas, both in our current responsibilities on Fermilab experiments (both universities have significant groups on both Fermilab experiments) and also on ATLAS. We have for a long time been responsible for large-scale project management.

UM has been involved longer in the computing/grid side of ATLAS than we have.

our plan:

1. protected wavelength through MiLR to Starlight
2. 10Gb/s connectivity on each campus
3. Two identical sites → single virtual presence to the Grid
4. dedicated, professional staffs at each campus

1. MiLR access

UM has been using MiLR in "data challenges"

stressing the high-throughput connectivity to BNL/CERN

anticipate a DC level at ~2-3Gb/s rate, bursting to 10Gb/s

dedicated MSU-UM wavelength for all research?

doable with negotiated use of either regeneration point

Some in the State and at the universities were brilliant a number of years ago when the move to expand the already unusually capable state network to 10Gb/s was approved. None of our colleague-institutions have this sort of locally owned network and its existence is a unique capability of the UM/MSU proposal. With the right switching and negotiations, we can guarantee a completely protected 10Gb/s capability, insuring uninterruptable service. In order to accomplish this we require minimally a DC level of usage on one wavelength of about 2-3Gp/s. We will peak occasionally to the full bandwidth, but we need to get into actual production and data-taking modes to understand this better.

With little effort and expense, **we can establish a dedicated research link for one wavelength at 10Gb/s linking exclusively UM and MSU** and U.S. ATLAS can be the first of many research efforts that enjoy this capability.

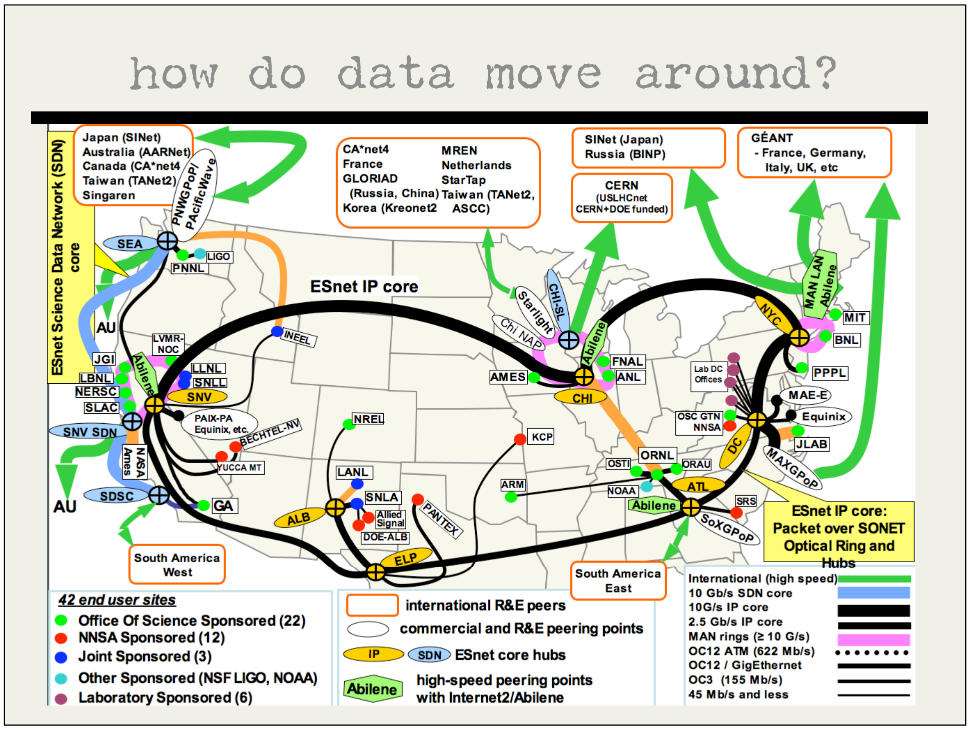

The DOE ESnet infrastructure is the backbone of the data pathways through the U.S.

MSU's and UM's gateway to both Brookhaven and CERN is through StarLight in Chicago, which peers with many international networks, but in particular "USLHCnet," a joint US-CERN project. 10Gb/s fiber networks is the necessary resource.

Michigan LambdaRail (MiLR)

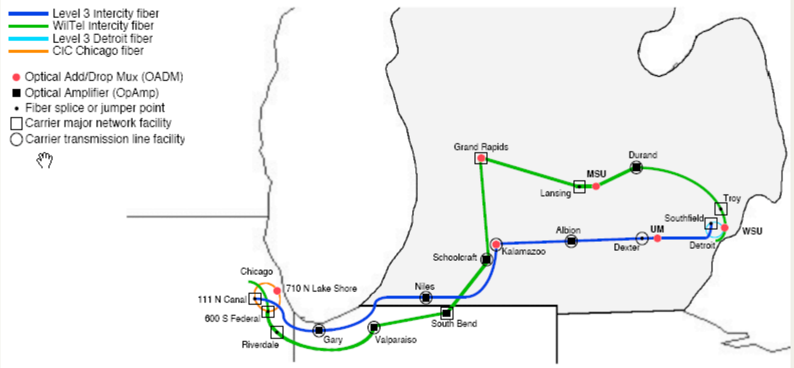

We have a single "/23" network for the AGL-Tier2

Internally each site (UM/MSU) will has a /24 (254 hosts)

Our network will have 3 10GE wavelengths on MiLR in a "triangle"

Loss of any of the 3 waves doesn't impact connectivity for both sites

Network Details (Almost Operational)

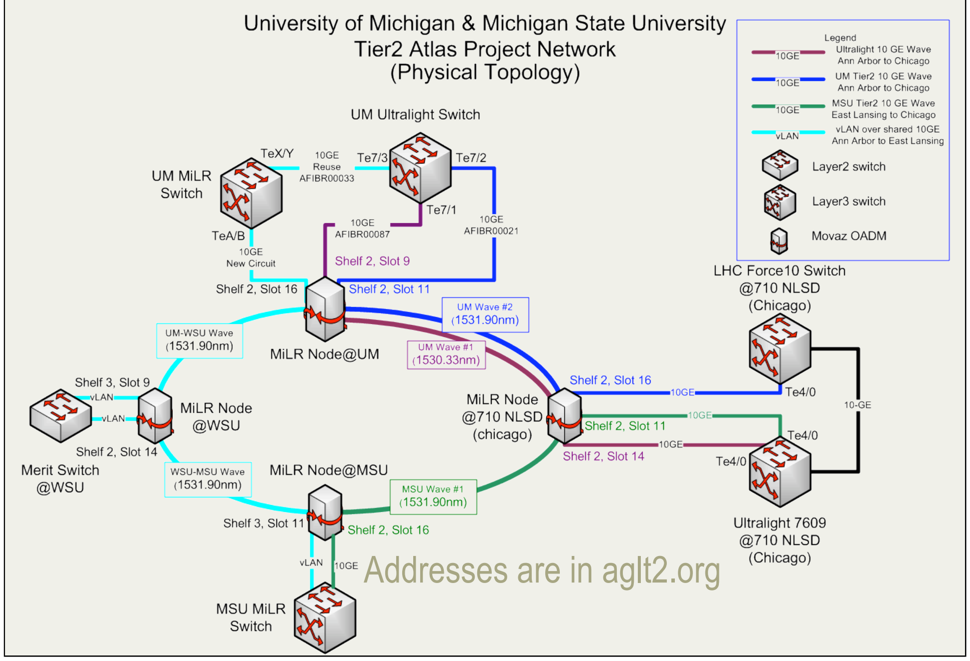

3. sites

- **computing equipment**
  identical in MSU and UM
  we would distribute cpu and storage purchases
  coordinating bidding with UM for best price
- **networking equipment**
  campus 10Gb/s equipment

We are prepared to front-load with CPU resources. That fits our current plans, and our startup obligations. However, that also puts a strain on heating.

We need to find a blend of up-front computing and the necessary RAID disk storage facilities necessary to the early and the later tasks. There will be lots of simulation early, less data...while more data will always stimulate more simulation needs. So, disk growth will expand non-linearly.

4. Personnel

We anticipate the need for 6-7 FTE technical people:

A. 1-2 at each site to **maintain hardware**, patching, repairing, backups, upgrading both processors and storage equipment

B. 1 for the T2 Facility for **grid software integration support**

C. 1 for the T2 Facility for **middleware**, T2 script creation/maintenance, facilitating transfers to and from T3 facilities, dealing with the (awful) job of maintaining world-wide software releases within the ATLAS Virtual Community

D. 1 for the T2 Facility for **Cluster Operations**

Conclusions.

We are very excited about planting this flag for our shared HEP Future at MSU and UM

Success in becoming a T2 site will further leverage individual grants at both institutions

We believe it makes the best use of current and future Michigan resources

tailor-made for MiLR