



I-Light

A Foundation for Innovation

I - L I G H T

Optical Fiber Infrastructure Supports Collaboration and Fosters Economic Development



Office of the Vice President
For Information Technology & CIO





Agenda

I - L I G H T


- Background on the project
 - Who was involved (i.e., *who is 'we'*)
 - Why we did it
 - How we did it
 - What is I-Light?
- The I-Light Advantage
- Initial Pay-offs
- What's Next?





I - L I G H T

Why I-Light?



Lots of fiber crosses Indiana

I-Light helps more of it stop here!






I - L I G H T

Why I-Light?

- Needed capacity for collaboration
- Needed capacity for inter-campus connectivity
- Wanted control over infrastructure/services
- Needed to stretch scarce infrastructure funding
 - Cost of multiple gigE connections (*if purchased from carriers*) ~\$2-million/yr per campus connected!





I - L I G H T

What is I-Light?

- Very high speed optical fiber network
- Connects IU Bloomington, IUPUI, and Purdue University West Lafayette
- Connects all three campuses to the national Internet infrastructure, including Internet2
- Co-Location facilities in prime Indy carrier hotel





I-Light Project Partners

I - L I G H T

- Indiana University
- Purdue University
- Intelnet Commission (IHETS)
- Verizon (optical fiber infrastructure)
- Juniper and Cisco (hardware)



IUPUI
Why not both?

PURDUE
UNIVERSITY


Research



Intelnet Commission

ihets.

CISCO SYSTEMS



verizon

Juniper
NETWORKS



Who owns I-Light?

- **Indiana and Purdue Universities do!**
 - Most all other networks like this are leased
 - IRU to State Network (Intelenet) for 4 fibers on each route, for State Use (TBD)
 - MOUs cover the various roles and responsibilities





How is I-Light Managed?

- IU and Purdue manage respective connections to IUPUI; IU responsible for oversight of operation at IUPUI and for connections through the Indiana GigaPoP.
- Roughly two-thirds of resource managed by the Universities; the rest available for broader State use
- Steering committee (IU, Purdue, Intelenet, IHETS) sets policy and pricing for State-Use fibers





I - L I G H T

Project Overview

- Idea first discussed in 1998
- \$5.3-million State appropriation in 1999
- Period of planning, design, and contract negotiations through 2000
- Construction began in Spring 2001
- Construction completed November 2001
- Fiber lit and put into use in December 2001






I-Light Project Specifics

I - L I G H T

- State appropriation of \$5.3-million in one-time funds
 - *Support of Indiana Higher Ed. Commission*
 - *Strong backing from Governor's Office*
 - *Leveraging IU's Internet2/Abilene role*
 - *Intelenet (State Commission) acted as Fiscal Agent*
- Sought Advice/Counsel of others
 - CA*NET in Canada – Bill St. Arnauld
 - I-Wire project in Illinois – Charlie Catlett
 - Internet2 – Doug vanHouweling and Steve Corbato






I-Light Project Specifics

I - L I G H T

- RFP Process to investigate options
 - Lease Existing?
 - Commission New?
 - Buy vs. Lease?
 - Vendor Selection?
 - Buried versus Arial?
- Developed MOUs
 - With State, as fiscal agent for appropriation
 - Between Partners – IU and Purdue





I-Light Project Decisions

I - L I G H T

- **Build to own** – not lease from others
 - *Believed advantage in asset control*
 - *Saw as a resource for now – and later*
- **Buried fiber plant** – *virtually none above ground*
 - *One look at the tornado damage in Martinsville tells you why this was important!*
- **Selected a Vendor (Verizon)**
- **Hired a consultant to assist on Plant Design**
 - *Money well spent!*
- **Selected Technology to Deploy**
 - *GigE selected for initial deployment*
 - *DWDM still a bit pricey – will probably go there at some point*
 - *Mix of Juniper and Cisco hardware*






I - L I G H T

I-Light Project Implementation


- Installed (2) conduits on each route – one with fiber and *one empty* for future use
 - Purdue Telecommunications Center (West Lafayette)
 - IUPUI University Library (IUPUI)
 - IUB Wrubel Computing Center (Bloomington)
 - Connection to two (2) Carrier Co-Location facilities in downtown Indianapolis:
 - Qwest (for Internet2)
 - Switch and Data (all purpose carrier facility)
 - 100% in public right of way





I-Light Project Implementation

I - L I G H T



- IUPUI-Purdue – 12 strands of Lucent True Wave
- IUPUI-IUB; IUPUI-Carrier Co-Los – 36 strand composite cable of single-mode (24) and Lucent True Wave (12)
- ~700,000 feet of fiber – 99.9% underground
- ~3.4-million feet of 1” conduit
- ~600,000 feet of warning tape buried above conduits
- 3000 feet of river bore (*yes, under the river-beds!*)
- Over 1200 fiber ends spliced
- 300 hand holds
- ~200 Slack Tubs
- 2 in-route repeater huts





I-Light Project Issues

I - L I G H T

- Needed to understand our responsibilities as owners
- Needed to understand that we are now, in some ways, a utility (but we aren't – have to be careful of that!)
- Knowing the true cost
 - Co-location is not cheap (~\$60K/year)
 - Liability Issues
 - Cost of ownership (cable monitoring, etc.)






I - L I G H T

I-Light IN-progress

Construction along the way





The I-Light Advantage

What is the capacity of I-Light?

I - L I G H T


- Previously, IU and Purdue suffered through limited network capacity
 - ~30-million bits per second between IUPUI and Purdue
 - ~200-million bits per second between IUB and IUPUI
- I-Light increased capacity immediately to multiple 1-**billions**-bits-per-second connections on both routes and provides similar capacity in the connection path to Internet2 (*which today runs at OC-12, but we're looking to expand soon*)
 - *Purdue is actually testing a 10-GigE connection right now*
- I-Light has the potential to expand to 100s and 1000s of **billions** of bits per second in the future





I - L I G H T

A representation of the
growth of **actual** network
capacity between IUB and
IUPUI as of January 2002







I - L I G H T

A representation of the
growth in **potential** capacity
of the connection between
IUB and IUPUI as of
January 2002

*Assumes all
fibers lit at
10Gigbps*





I - L I G H T

A representation of the
growth in **theoretical**
capacity of the
connection between IUB
and IUPUI as of
January 2002

*Assumes all fibers
lit using advanced
DWDM
technology
running multiple
10Gigbps lambda
on each fiber*





The I-Light Advantage

How does I-Light Relate to Internet2?

- I-Light acts as a digital on-ramp to Internet2
 - *Extends the high-speed research networks into the heart of the State*
- Resources at IU that manage the Internet2 network Abilene also manage I-Light
- The Indiana GigaPoP, located at IUPUI, is the focal point of I-Light, and the main access to Internet2 for IU, IUPUI, Purdue, and several other regional institutions





The I-Light Advantage

Importance to Internet2

I - L I G H T

“Indiana University, Purdue University, IUPUI, and the City of Indianapolis have been integral to the success of Internet2. With I-Light, the State of Indiana is taking a lead role by providing the foundation upon which the future of Internet technology can be built.”

Douglas van Houweling

President and CEO

University Corporation for Advanced Internet Development.





I - L I G H T

The I-Light Advantage

University Ownership

- Represents long-term investment by the State in research infrastructure
- Investment made during good economic times will help retain and strengthen State's IT advantage
- Provides enough networking capacity for the next 10-20 years between the three main research campuses (IU, Purdue, IUPUI)





I - L I G H T

The I-Light Advantage

Collaboration

- I-Light significantly reduces the barriers to digital collaboration -- Ushers in a new era in collaboration between Indiana's two great State Universities
- Three research campuses will have greater leverage and potential for federal grants
- More grants to campuses will make Indiana a more substantial player in the Information Economy





The I-Light Advantage

Uses

- Supports High-end Research Applications – Virtual Reality/Visualization, Tera-scale processing, massive data storage
- Supports Operations – Voice, E-mail, and Video Conferencing
- A unparalleled platform for collaborative research and distance education





The I-Light Advantage

Indiana Virtual Machine Room

I - L I G H T

- I-Light also allows the universities to pool their high-end computational resources to build larger, more effective facilities
- Supercomputers, massive data storage facilities, and visualization environments are easily assembled into 'grids' of virtually unlimited capabilities

That these resources - *some of the largest in the world* - are now available in the State of Indiana *cannot be ignored*





The I-Light Advantage

Being First!

- Indiana is the *first state* to deploy such a high-performance data network
- Few other states have Indiana's geographical advantage when it comes to tapping into existing fiber pathways/crossroads





I - L I G H T

The I-Light Advantage

Important Nationally!


“By investing in their own high performance regional network infrastructure, Indiana has taken a very strong forward looking position in the development of high performance networking for the academic community.”

“This unique new capability is highly complementary to the Federal role in national and international connectivity for the research and education community. Indiana, with this new capability, builds on its already well established leadership role, and ensures a strong foundation for future developments both regionally and nationally.”

Aubrey Bush

Division Director, National Science Foundation





I - L I G H T


I-Light

Beyond the Universities

Building not just for the Universities, but for the City of Bloomington too!

- State/City ‘piggy-backed’ on the University project
- Got a great deal on empty conduits connecting Bloomington Fiber Ring and Carrier Hotel to Indianapolis Fiber Concentration Points
- City Looking for ‘Anchor Tenant’ to start the process; or State will bid





Initial Payoffs


I - L I G H T

- Commodity Internet Bandwidth
 - Was able to consolidate IU, Purdue, and part of the State's demand for Commodity Internet drainage
 - Due to location, volume, and competition, we were able to **DRASTICALLY** cut our CI drainage cost
 - *For the same investment, was able to nearly quadruple our bandwidth!*
 - Was able to meet pent-up demand on our campuses – this Fall, bandwidth usage and internet performance are NOT issues!
- Made it possible to more fully use existing I2 'on-ramp' and to pave the way for expanding our connection






I - L I G H T



I - L I G H T


A representation of the **actual**
growth of internet drainage for
the IU network (IUB/IUPUI-
Regionals) capacity in the Fall
of 2002






Commodity Internet Use


I - L I G H T





Internet2 is also impacted

I - L I G H T





I - L I G H T

What does the future hold for I-Light?

- Expansion to build a Regional (Great Lakes) fiber optic backbone
 - Link to Chicago network ‘hubs’
 - Link to I-Wire in Illinois
- Becoming part of a national fiber backbone infrastructure
- Expanding the ‘distribution side’ of I-Light
 - Connecting other ‘partners’ in the State of Indiana

I-Light2 – Steve Mayo of Purdue will be detailing this next!





I-Light

A Foundation for Innovation

I - L I G H T

You can find more information about I-Light on the web at:

<http://it.iu.edu>

<http://www.indiana.edu/~ovpit>


<http://www.i-light.iupui.edu>

Presentation Prepared by:

- David E. Jent, Group Manager, Network Engineering & I-Light Project Manager
- Brian D. Voss, Associate Vice President for Telecommunications

I - L I G H T





I - L I G H T



December 4, 2002

Stephen Mayo

**Associate Vice President,
Telecommunications at Purdue University**



- \$5M dollar investment by State of Indiana to purchase optical fiber infrastructure.
- This week added 10Gig pipe to the existing 1Gig connection to IUPUI.
- Allows for unprecedented collaboration.

- Planning process with IU, ITN, IHETS, Ball State, Notre Dame, Purdue to identify options for deploying broadband connectivity to all Indiana colleges and universities.

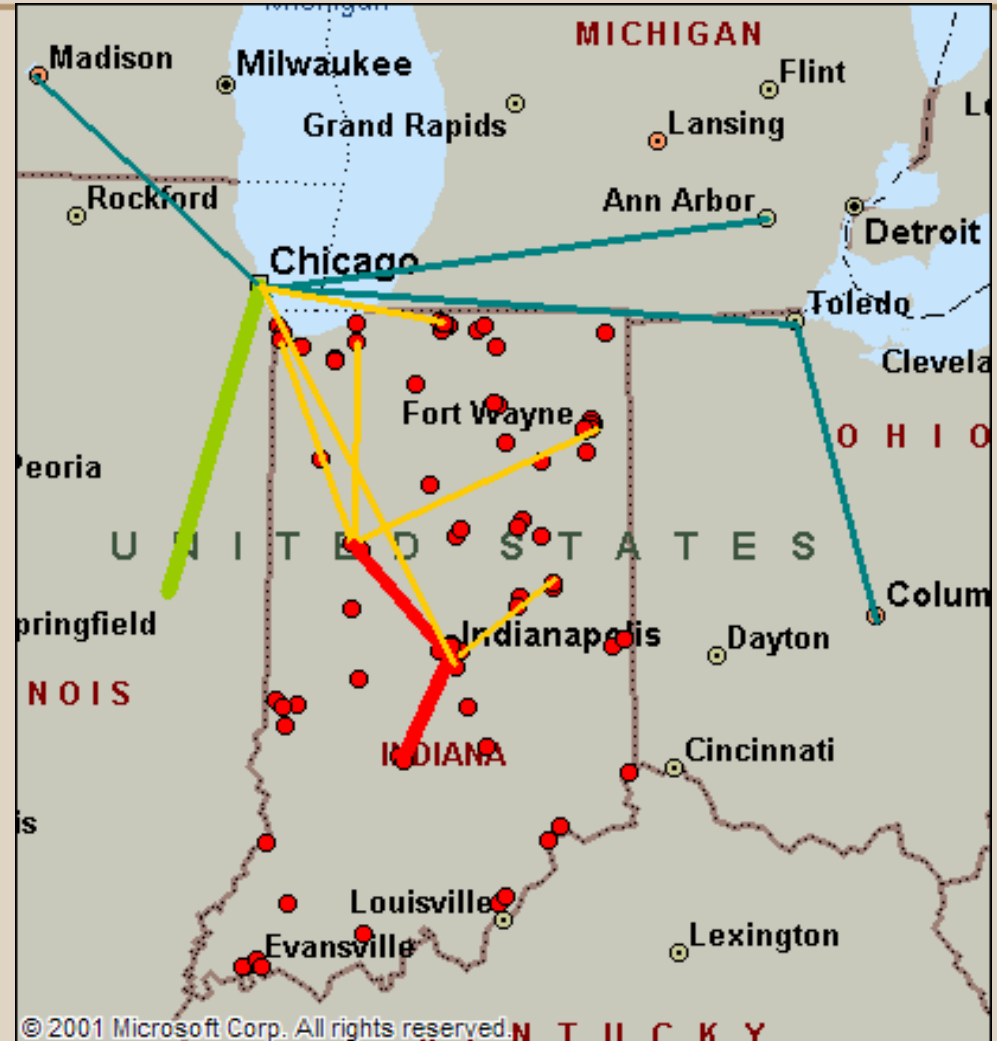
Vision and Objectives

- Create excitement at the campus, state, and national level.
- Create a dialogue among and between researchers/educators.
- Create strategic partnerships between local and state government, education and telecommunications providers to design and implement broadband connectivity.
- Create momentum toward an Indiana broadband strategy that engages colleges/universities, residents/parents, businesses/customers, students/teachers, and state/local government.

- Purdue, IU, and IHETs are members of the Great Lakes Light Rail Initiative evaluating connectivity options for Higher Ed networks in the Great Lakes States (Merit, WisNet, OarNet, IWire, and ILight).
- An RFI was distributed to potential network providers in October, 2002 to solicit recommendations for broadband connectivity to Indiana colleges and universities.
- Responses arrived November 22, 2002.
- Review committee will meet on December 10, 2002 to review the responses and decide upon next steps.

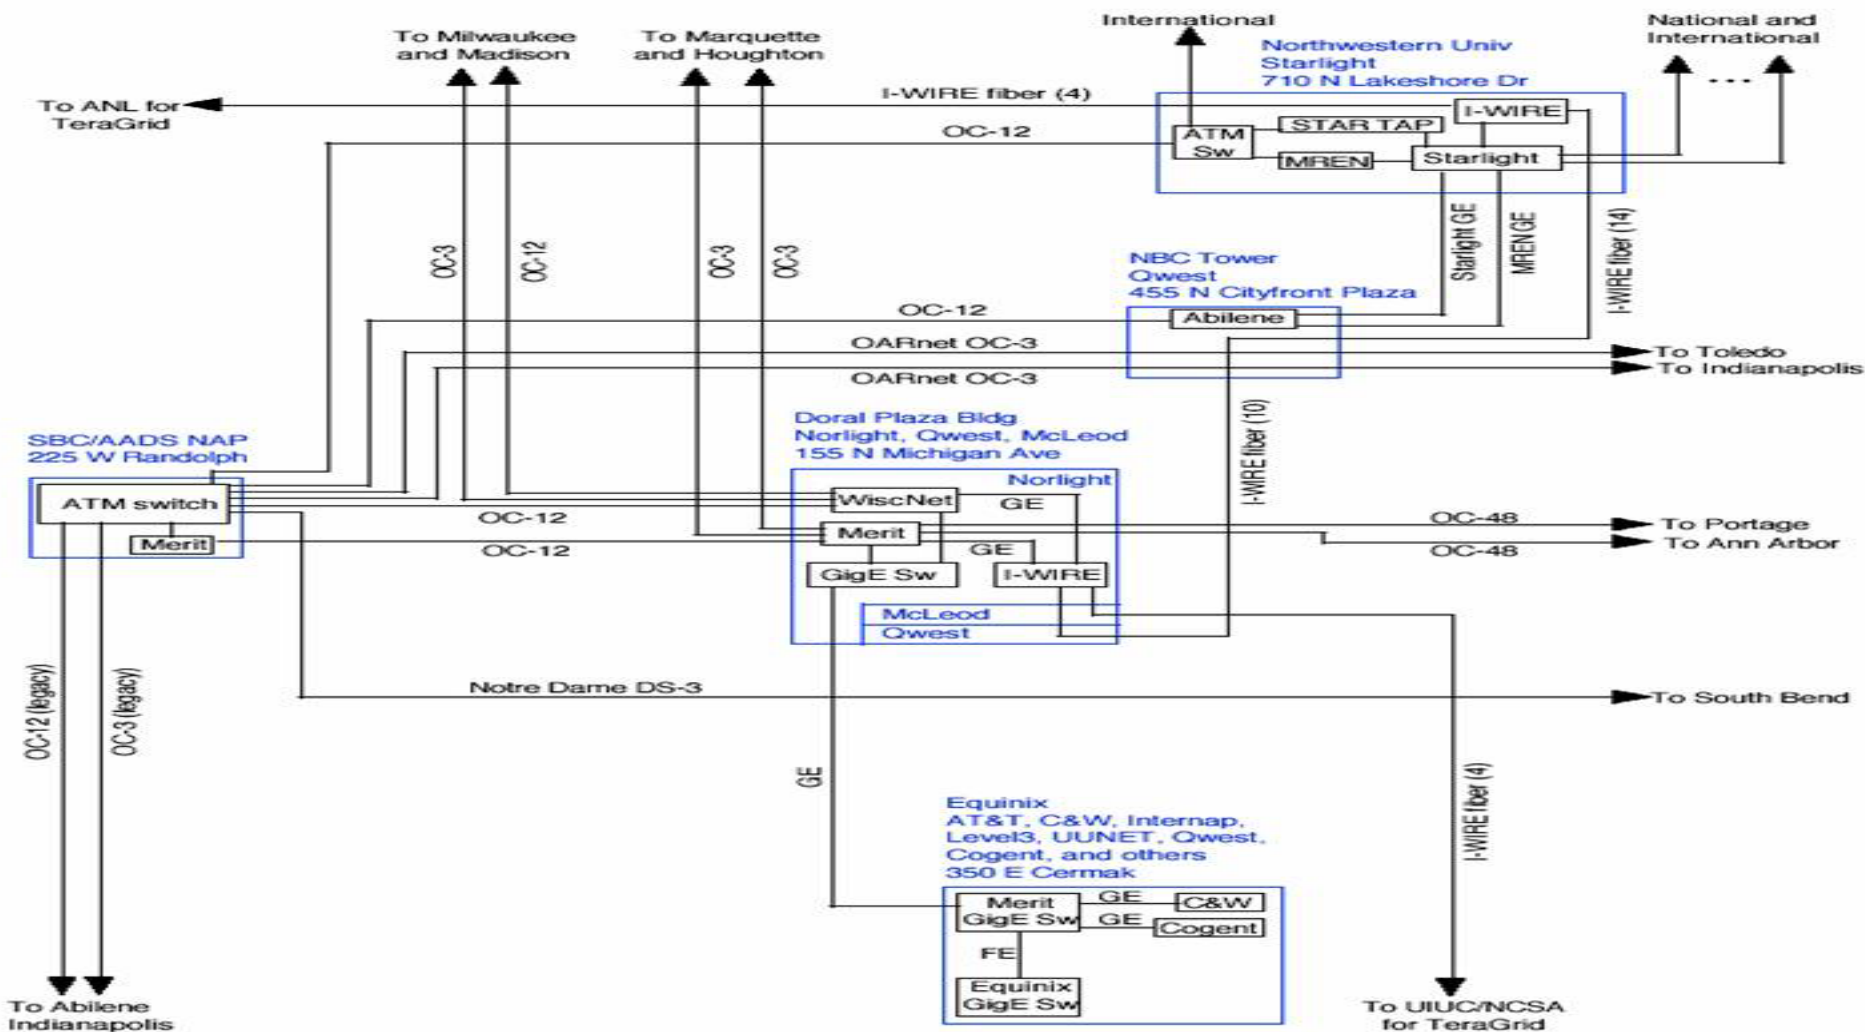
PURDUE
UNIVERSITY

- Indy GigaPop to Starlight and IWire via Chicago.



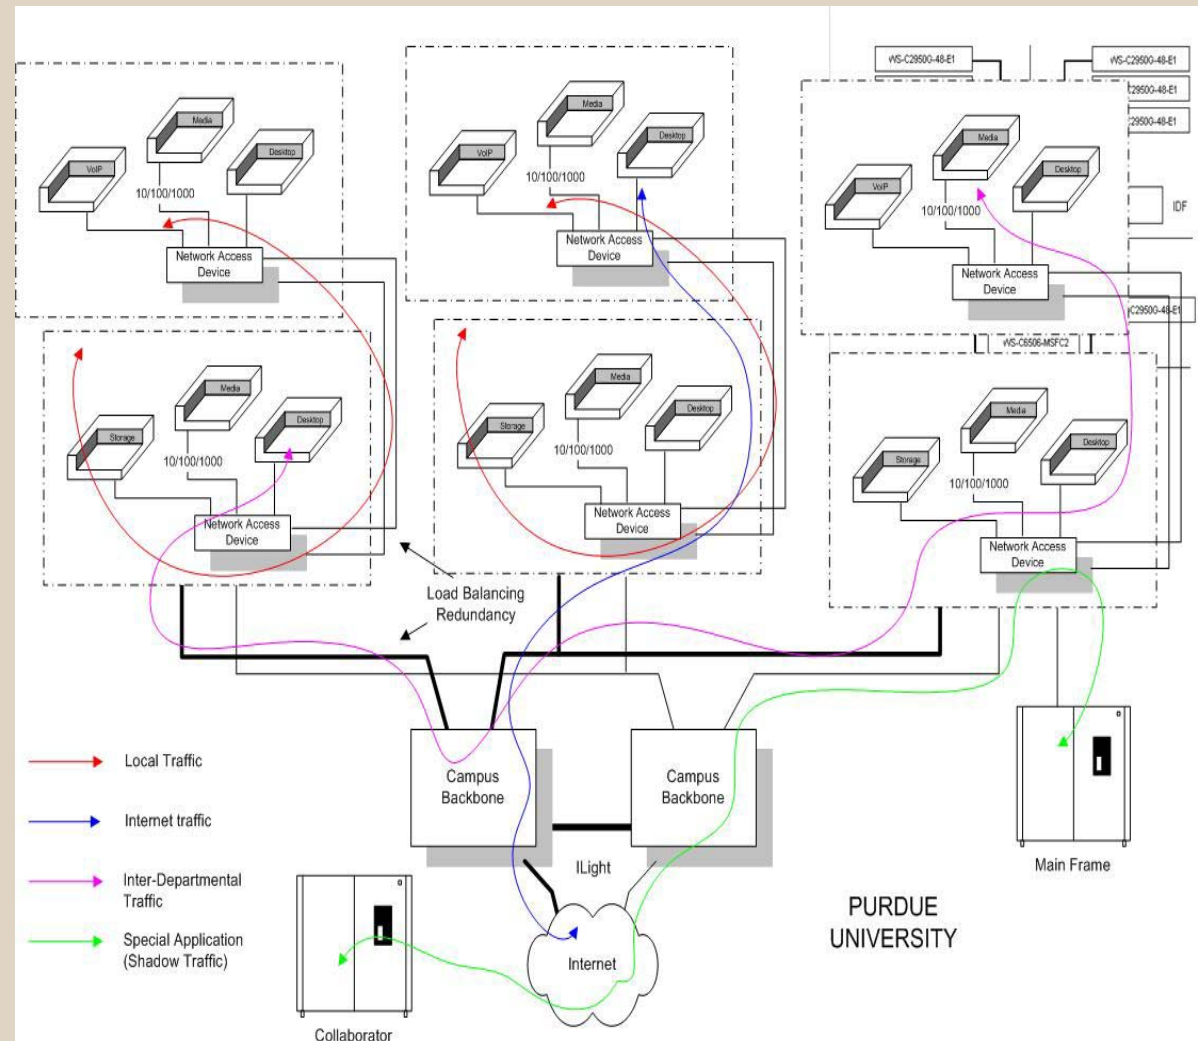
Chicago Connectivity

Chicago Connectivity



Purdue Shadow Network

The goal of the architecture is to minimize the risk of technology failure by establishing redundant, critical paths and to incorporate flexibility in extending or expanding connectivity to meet immediate needs (i.e.. point to point fiber).



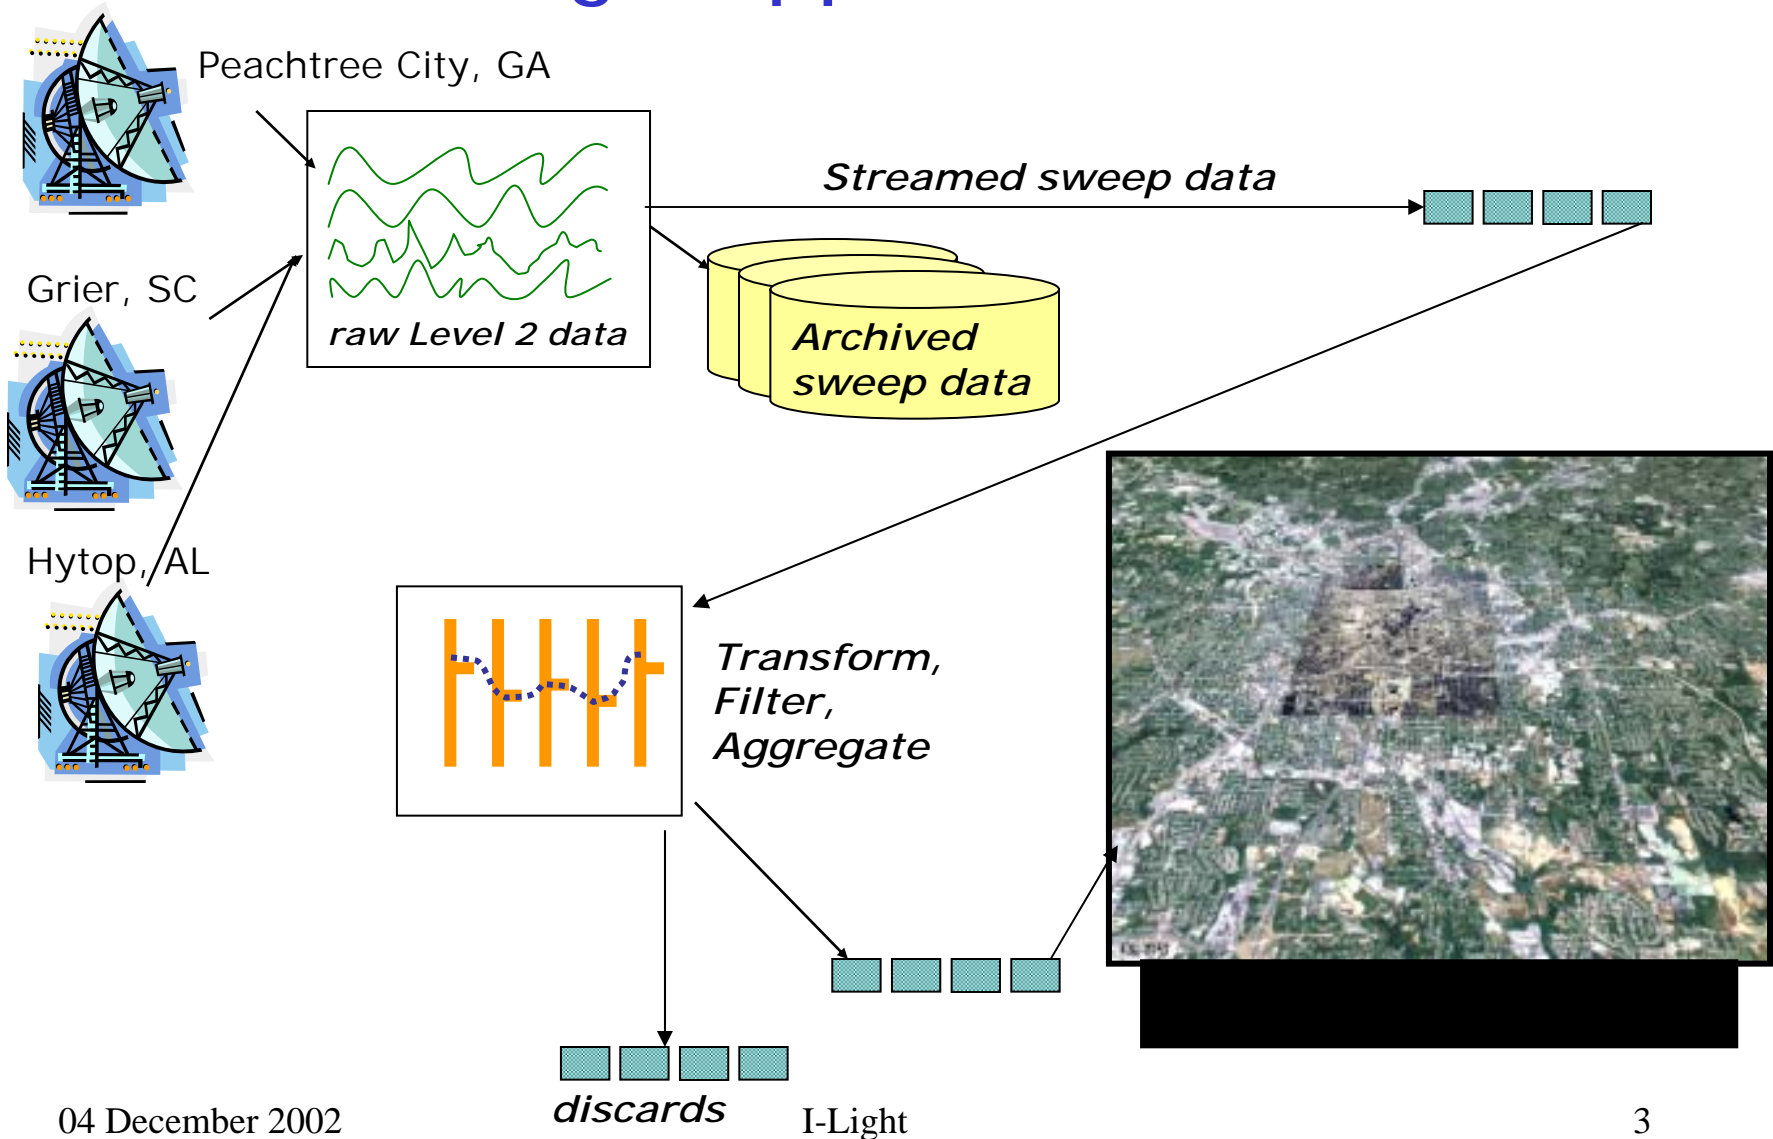


dQUOB: decision-making over streaming data with SQL queries

Beth Plale
Computer Science Dept.
Indiana University
Bloomington, IN

- Data streams are flows of timestamped events that are manipulated by means of continuously executing queries
- Continuous Query Work
 - Fjords (Franklin), NiagraCQ (DeWitt), STREAM (Widom), dQUOB (Plale)

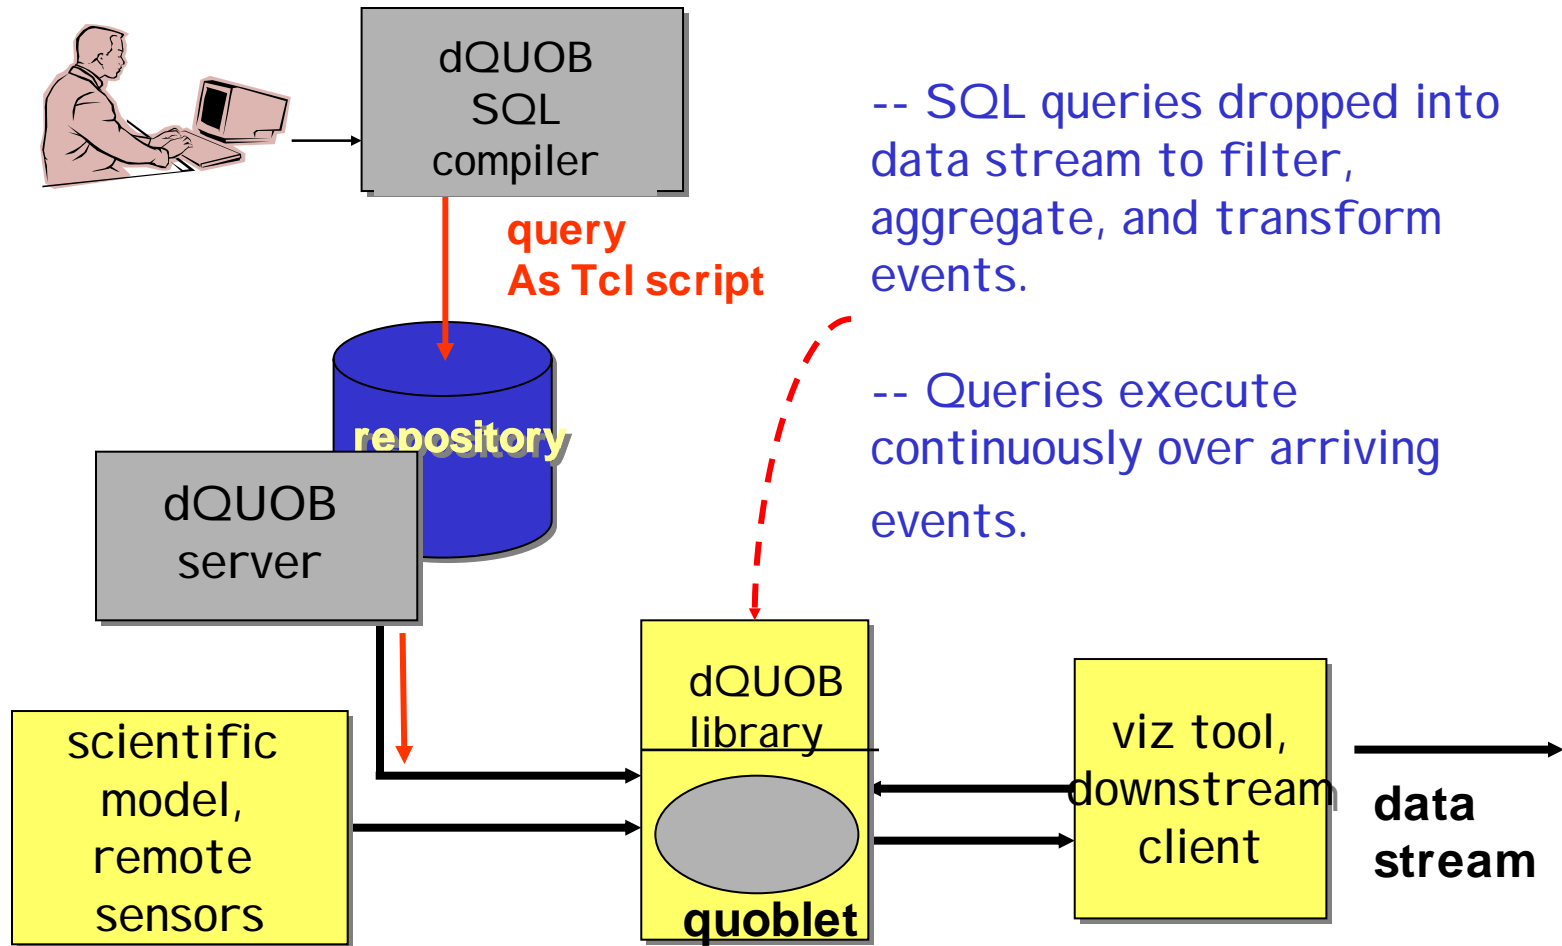
Visualizing Doppler Radar Flows



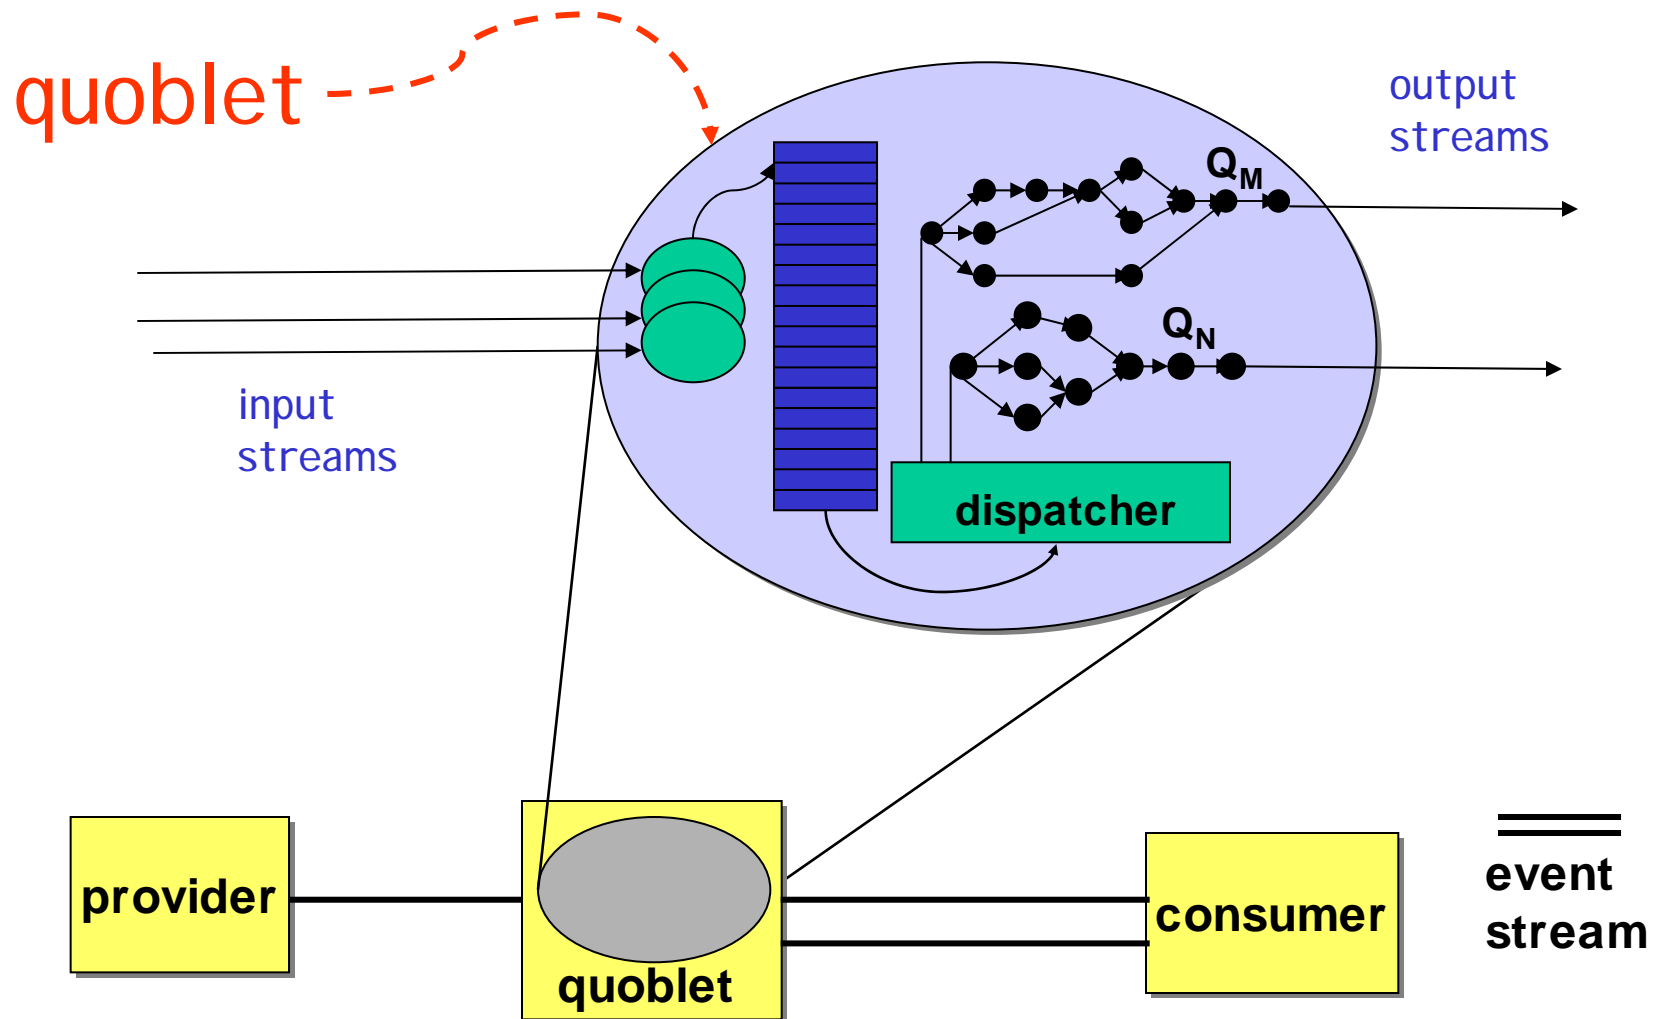
Quick Summary of dQUOB toolkit features

- SQL queries coupled with user-defined functions (e.g., FFT, data reduction).
- Assumes data stream is timestamped sequence of events.
- Event == tuple, data stream == relation
- Supports time-based stream join
 - Two events satisfy a join if they 'happen at the same time'
- Applied to: large scale scientific instruments, scientific visualization, hazard detection

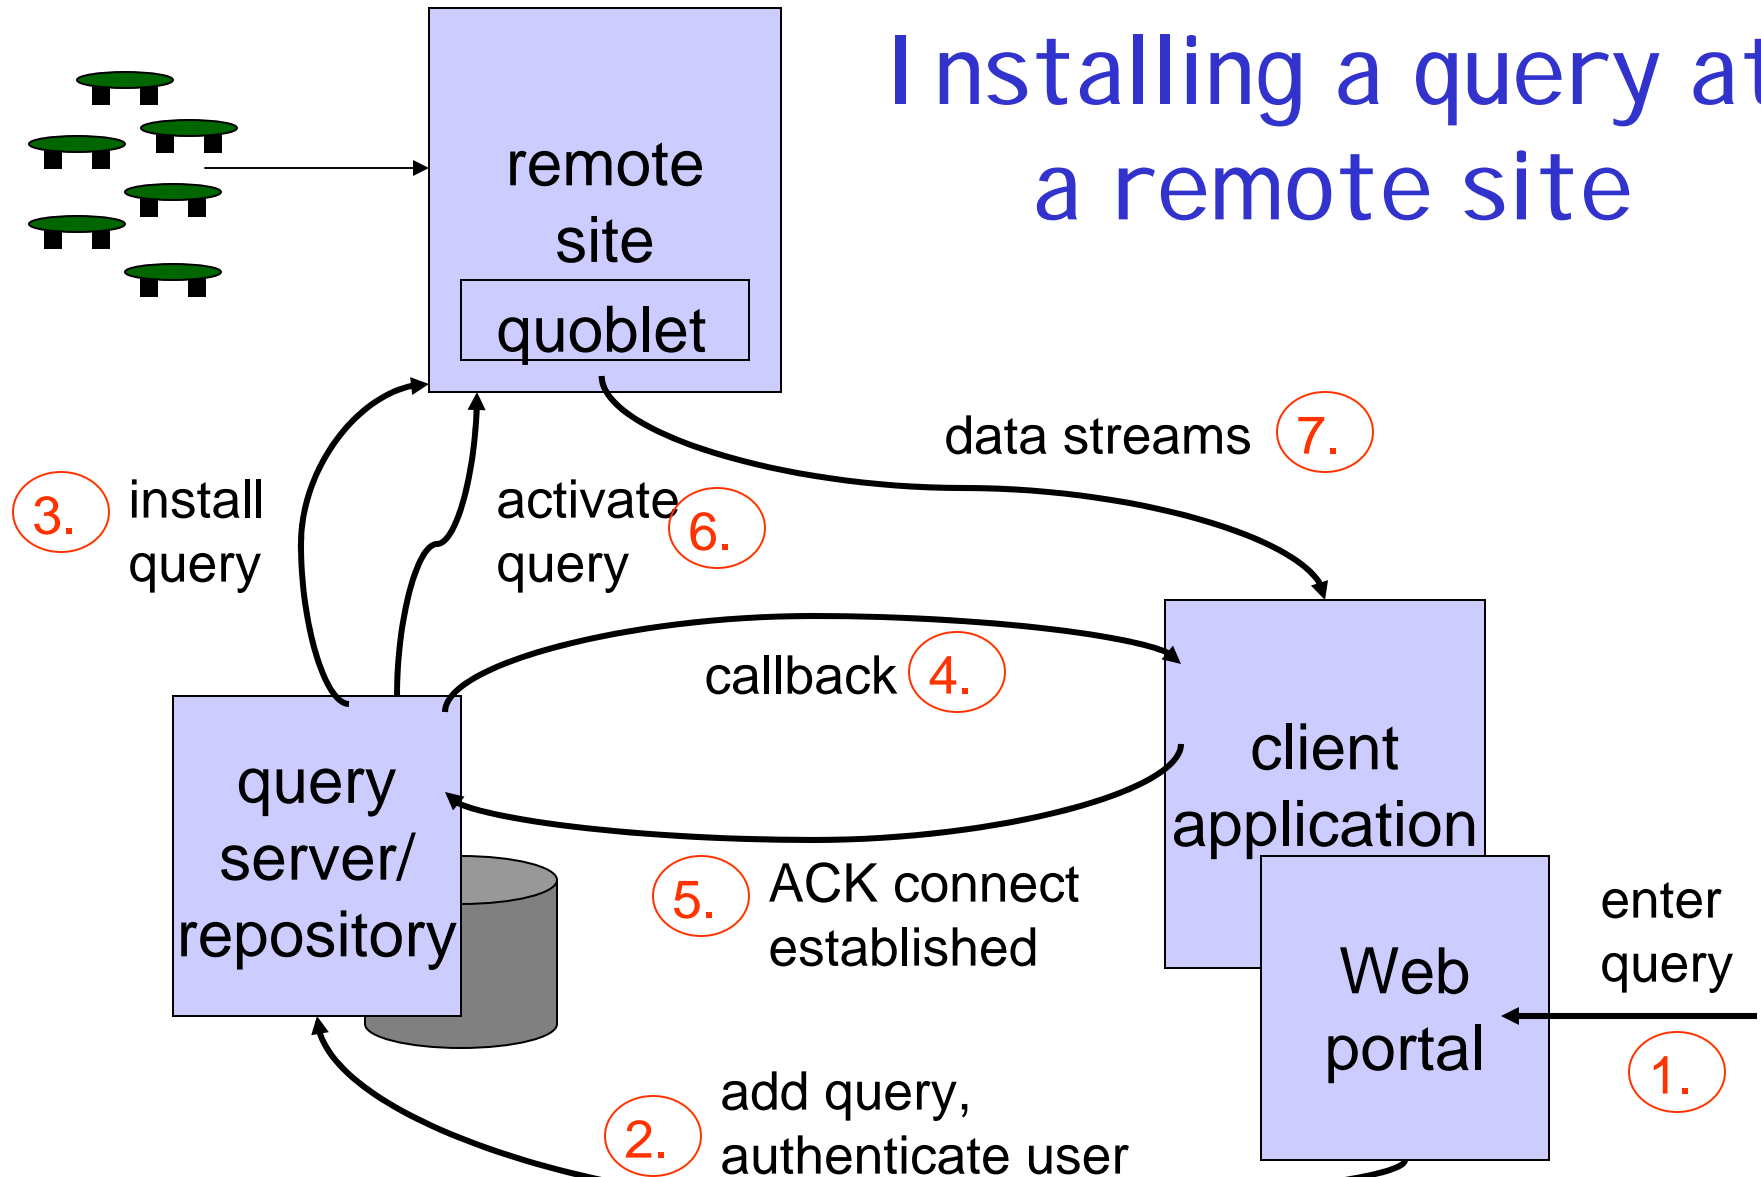
dQUOB toolkit components



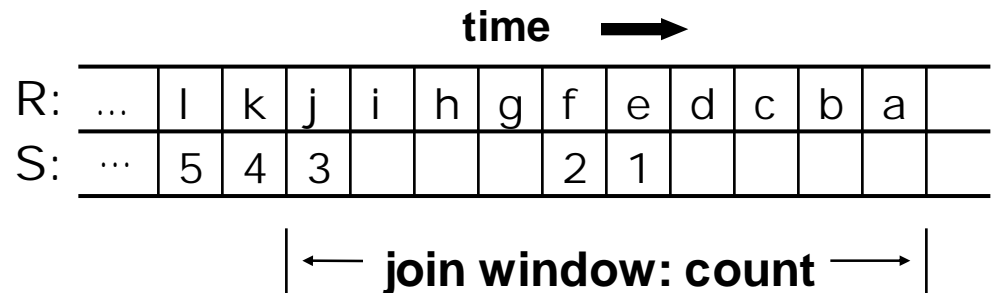
Architecture of quoblet for executing queries at runtime



Installing a query at a remote site





Research Issue: when R and S are asynchronous streams and stream S is **slow** and **erratic**, unneeded memory consumption



Issue 2: difficult for user to pick right join window size. Cost of error is great: too large, **consumes memory**; too small, increases **false negatives**

Approach to asynchronous stream problem: express join window size as interval of **time**



Startup Latency Evaluation

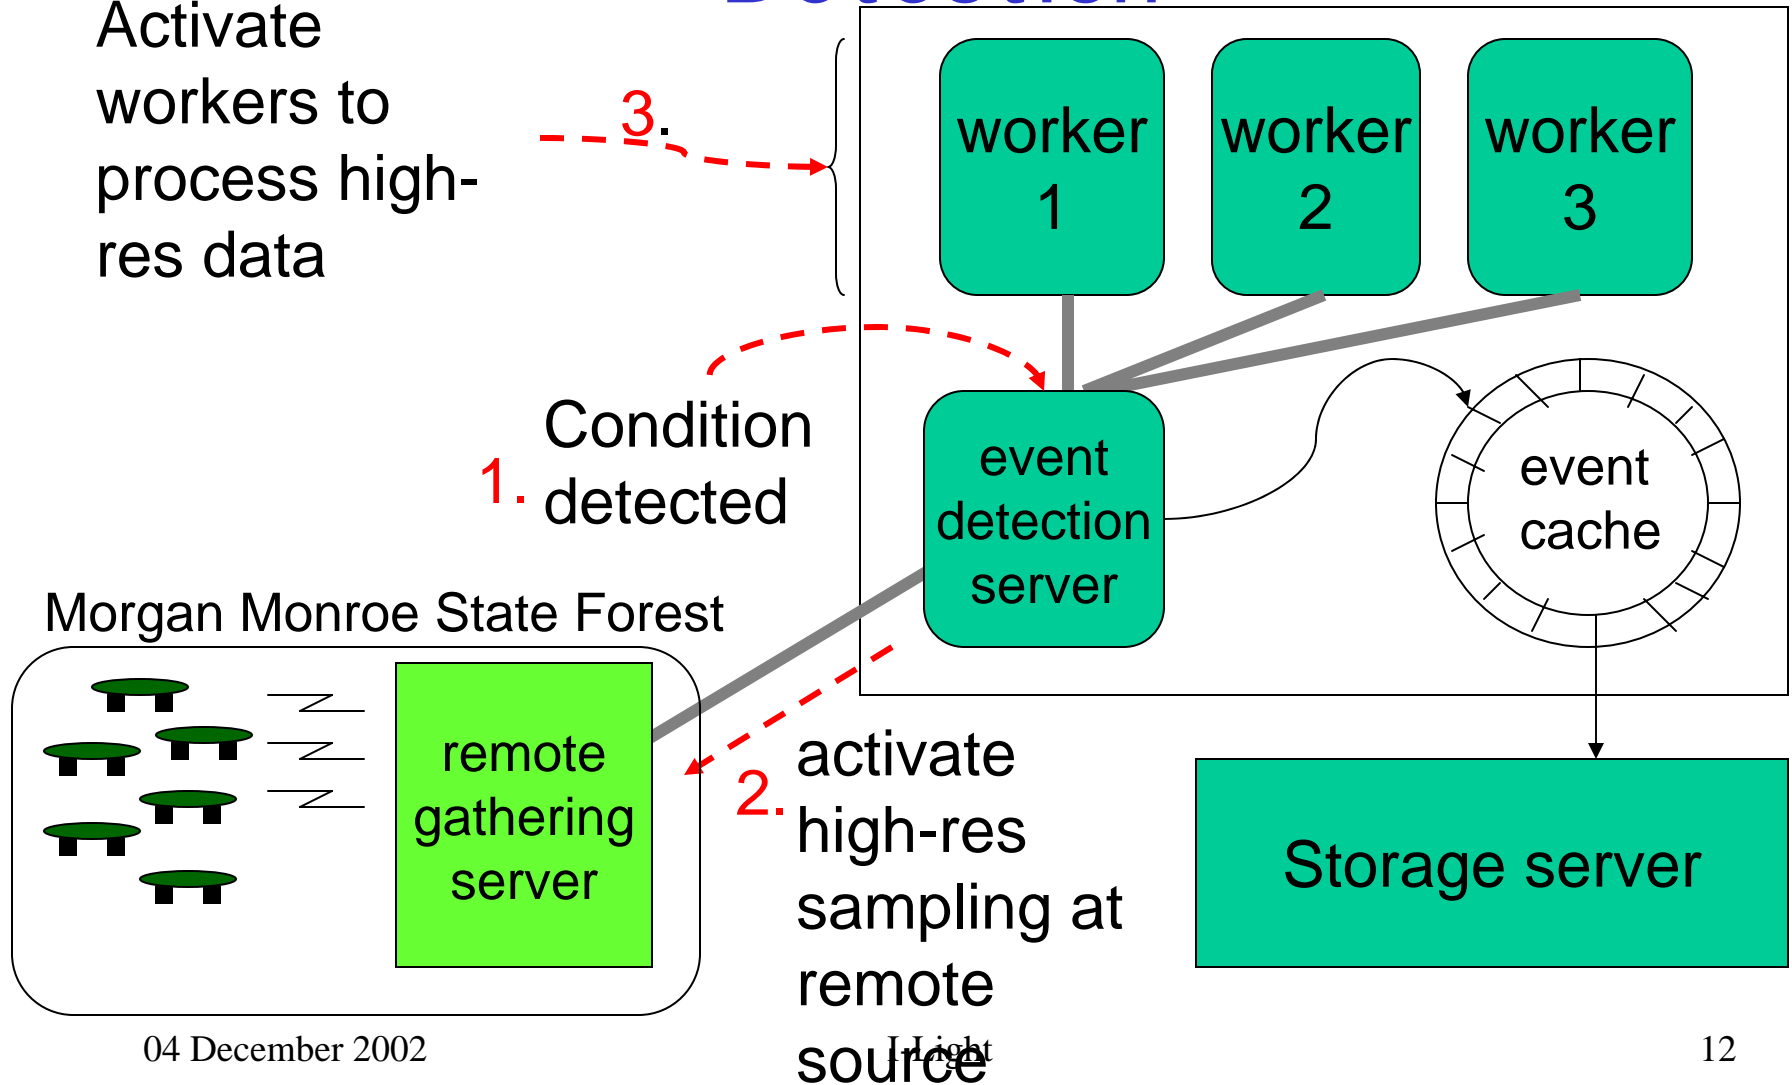
- Purpose: measure startup latency under various runtime environments.
 - Response latency – interval between when a complex event occurs and when a response is initiated at the remote site.
- Runtime environment: worker processes started up using either PBS or Linux OS real time scheduling features.

Experimental Environment

- Proto-AVI DD cluster
 - 5 IA32 nodes
 - Housed in IU University Information Technology Services (UITS) building
- Remote sensor server at Morgan Monroe State Forest (20 miles north of IU)
 - Radio link to IU
 - 4.3 Mbits/sec (iperf 6.2)
- Remote sensor server at IU Computer Science
 - Dual Poweredge 6400 server
 - 94.1 Mbits/sec (iperf 6.2)

Response to Weather Condition Detection

Activate workers to process high-res data



Methods Evaluated

1. PBS/Qsub – submitting via PBS using Qsub command line interface, no running jobs
2. PBS/Qsub/preempt – Qsub interface, existing jobs running under PBS on worker processors that must be preempted
3. Linux real time/RSH – Linux real time scheduling and RSH
4. Linux real time/SSH – Linux real time scheduling and SSH

PBS

- PBSpro 5.2 preemption
- Queue priority
- SIGSTOP/SIGCONT
 - Quickest
 - No swapping to disk
 - Safest
- Checkpointing

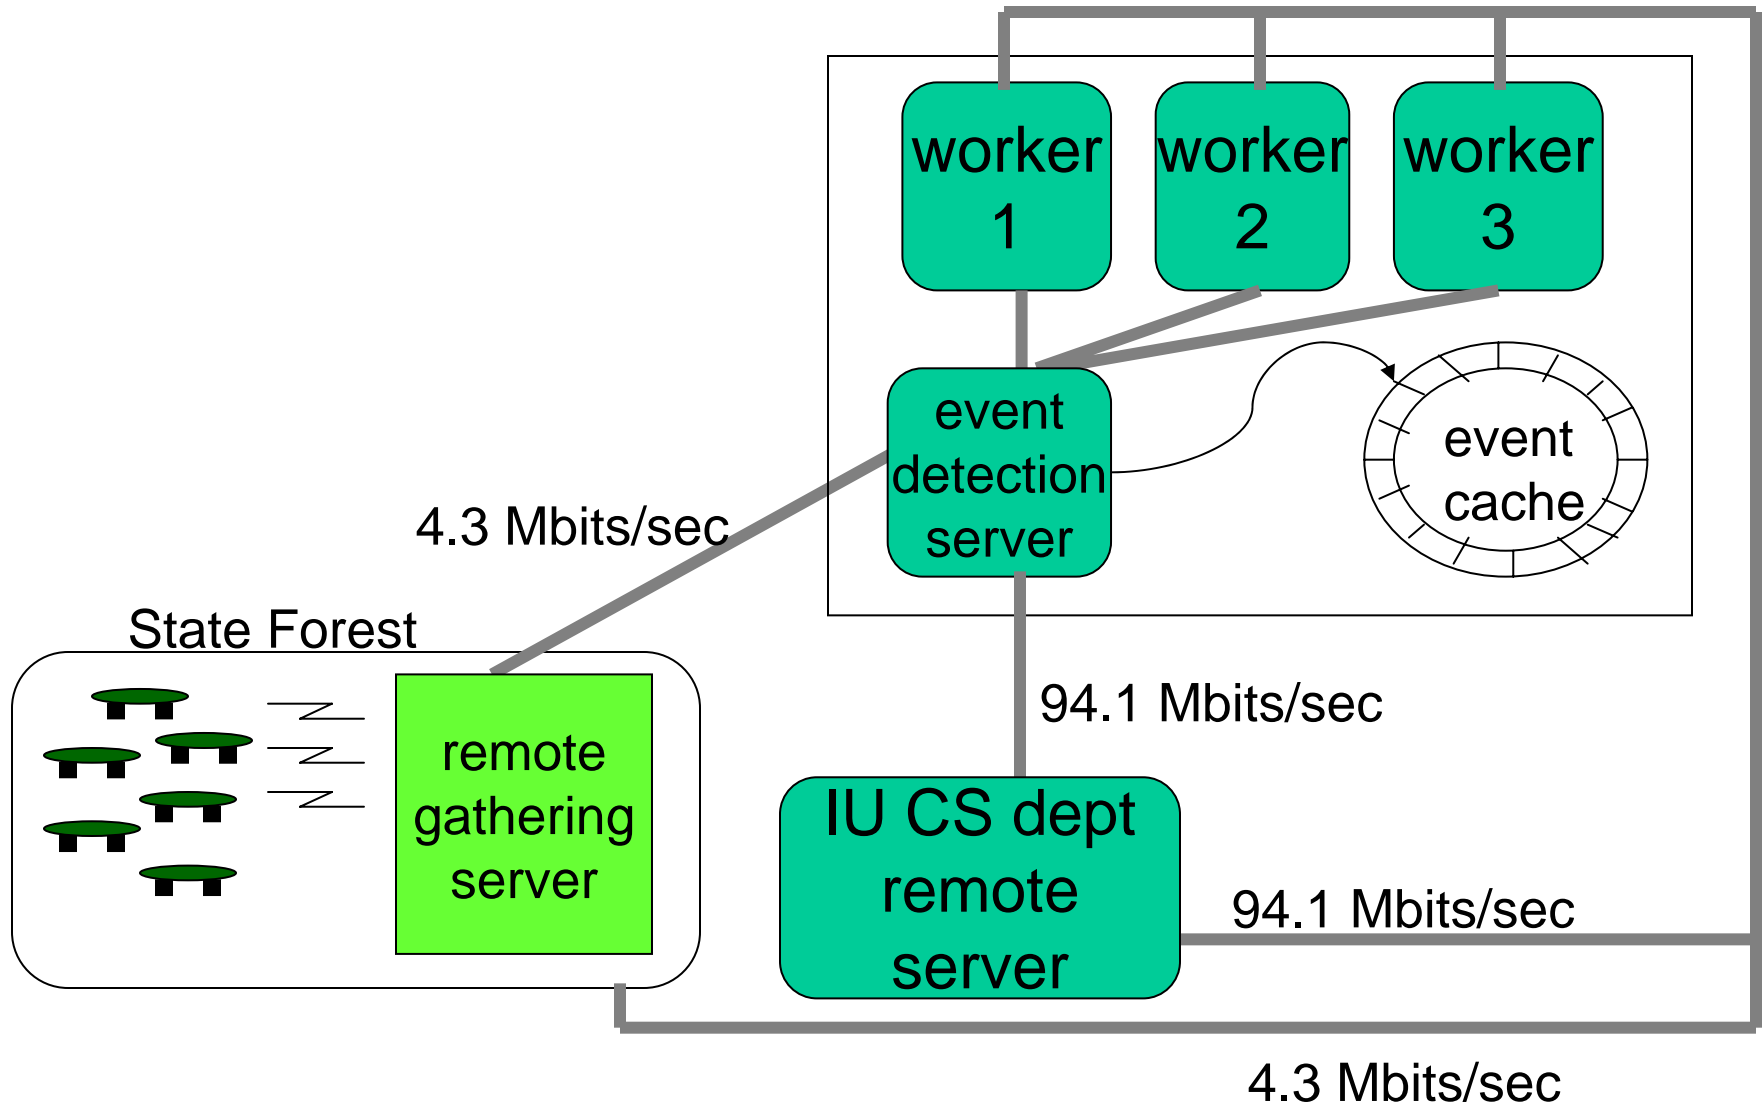
Real time scheduling in Linux

- Using `sched_setscheduler()`
- Not “Hard” real time.
- `rt_priority`: like “nice”
 - integer 0-99,
 - higher number=higher priority
- Preempts non-real time processes
- round robin scheduling within priority (`SCHED_RR`)

Key Events in Detection/Response Processing

Timestep	Event Definition
t_0	Weather condition first detected at server determined from low resolution event data
t_1	First high resolution weather event received at AVI DD
t_2	Worker 'up and ready'; instance prior to opening connection back to remote client
t_3	Worker receives first high resolution weather event

Experiment Scenario



Startup latencies for PBS Scheduling case

Interval	PBS/Qsub LAN (sec)	PBS/Qsub remote (sec)
$\{t_0 - t_1\}$	0.008414 ± 0.001685	0.022217 ± 0.004000
$\{t_0 - t_2\}$	1.264183 ± 0.463738	1.265159 ± 0.490130
$\{t_2 - t_3\}$	0.003163 ± 0.006258	0.088769 ± 0.182738

$\{t_0 - t_1\}$ interval between detection and activation of high-rate stream

$\{t_0 - t_2\}$ interval between detection and worker startup,

$\{t_2 - t_3\}$ worker delay before receiving first event

Startup Latencies under Linux real time

Interval	SSH (sec)	RSH (sec)
$\{t_0 - t_1\}$	0.053161 ± 0.001464	0.056945 ± 0.000361
$\{t_0 - t_2\}$	0.757201 ± 0.017996	0.057351 ± 0.000181
$\{t_2 - t_3\}$	0.012534 ± 0.006820	0.013329 ± 0.006804

$\{t_0 - t_1\}$ interval between detection and activation of high-rate stream

$\{t_0 - t_2\}$ interval between detection and worker startup,

$\{t_2 - t_3\}$ worker delay before receiving first event

Summary of Startup Latency

	Remote source response latency (sec) $\{t_1 - t_0\}$	Worker response latency (sec) $\{t_3 - t_0\}$
PBS LAN	0.008414	1.267346
PBS remote	0.022217	1.353928
RSH MAN	0.056945	0.070681
SSH MAN	0.053161	0.769735



Beth Plale

<http://www.cs.indiana.edu/~plale/projects/dQUOB>

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under Grant No. 0116050. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).



LABS



Measured Response

**The Application of Terascale Computing in
a Homeland Security Simulation**

Prof. Alok Chaturvedi
alok@purdue.edu

Dr. Shailendra Mehta
mehta@purdue.edu

Krannert School of Management

Funded By:
Indiana State 21st Century Research And Technology Fund
National Science Foundation

Measured Response is a multi-organization simulation that provides the infrastructure to develop emergency response solutions to a bio-terrorist attack during a major spectator event

- A tool for decision makers to test potential resource allocations and planning options in virtual environment
- For example, it creates a “virtual world” to determine how actions taken by HLS agencies would influence the outcome of the bio-terrorist attack

Measured Response

The Disaster:

The release of a biological pathogen (disease causing germs) into the general audience at the concert...

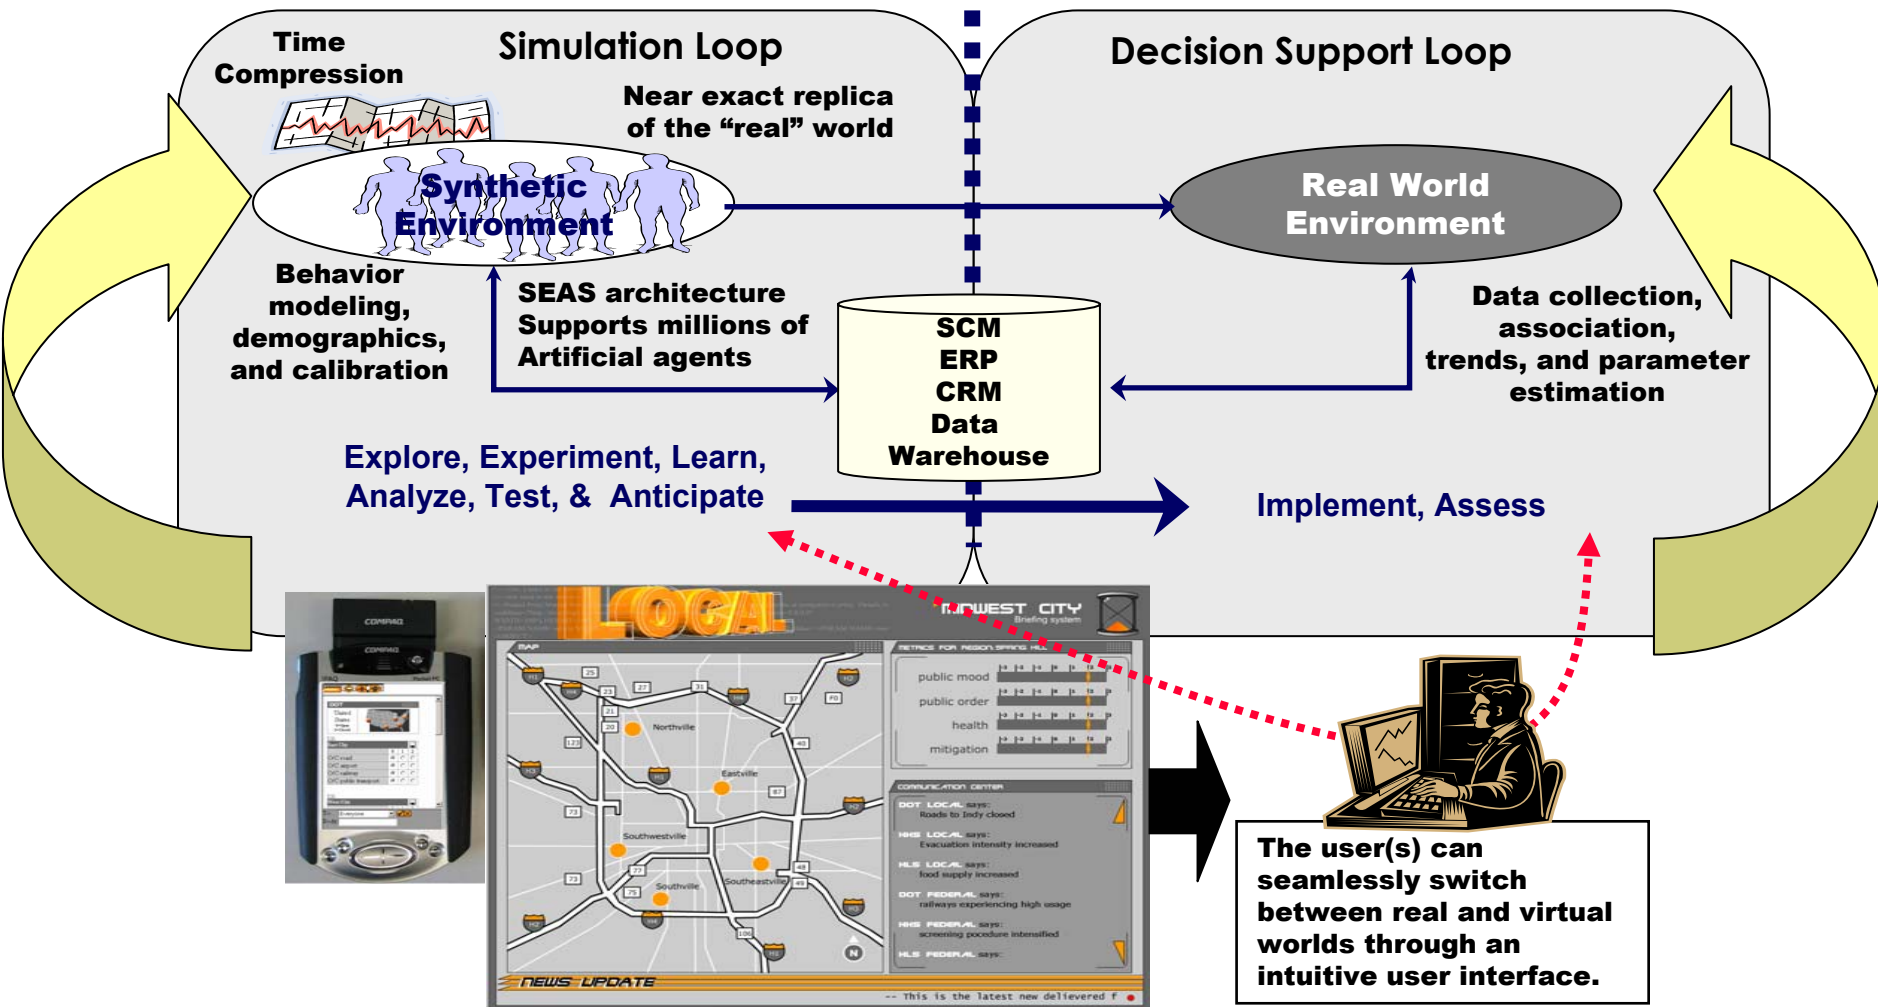
The Simulation:

Participants were presented with the situation, took actions to reduce the spread of the disease and quell public panic, then reviewed the results of their actions

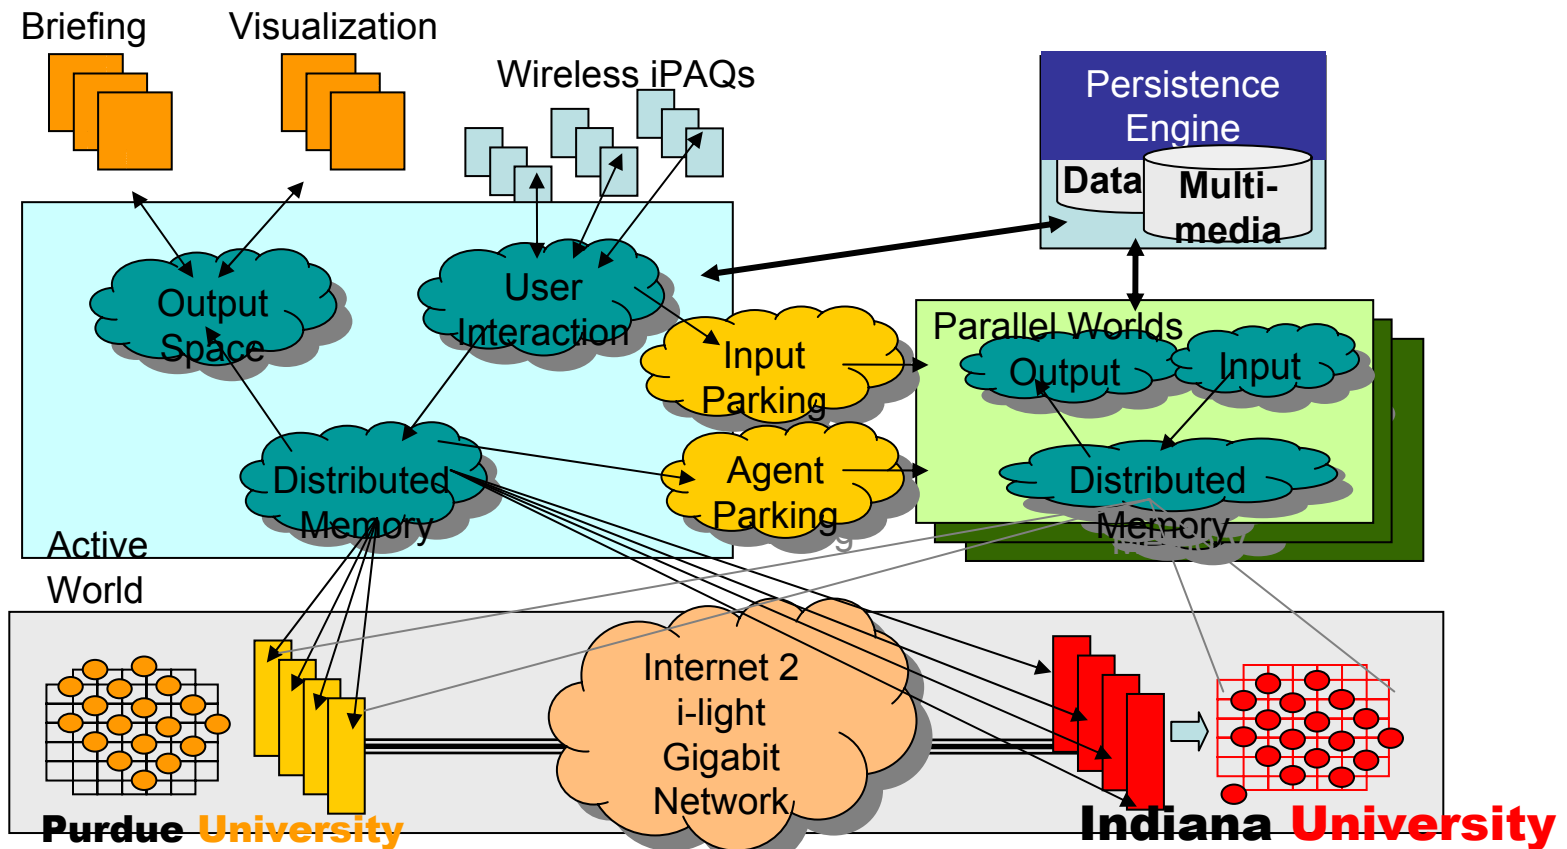
The value:

The value of any simulation is not the technology used for the simulation, but the lessons learned from its outcome...

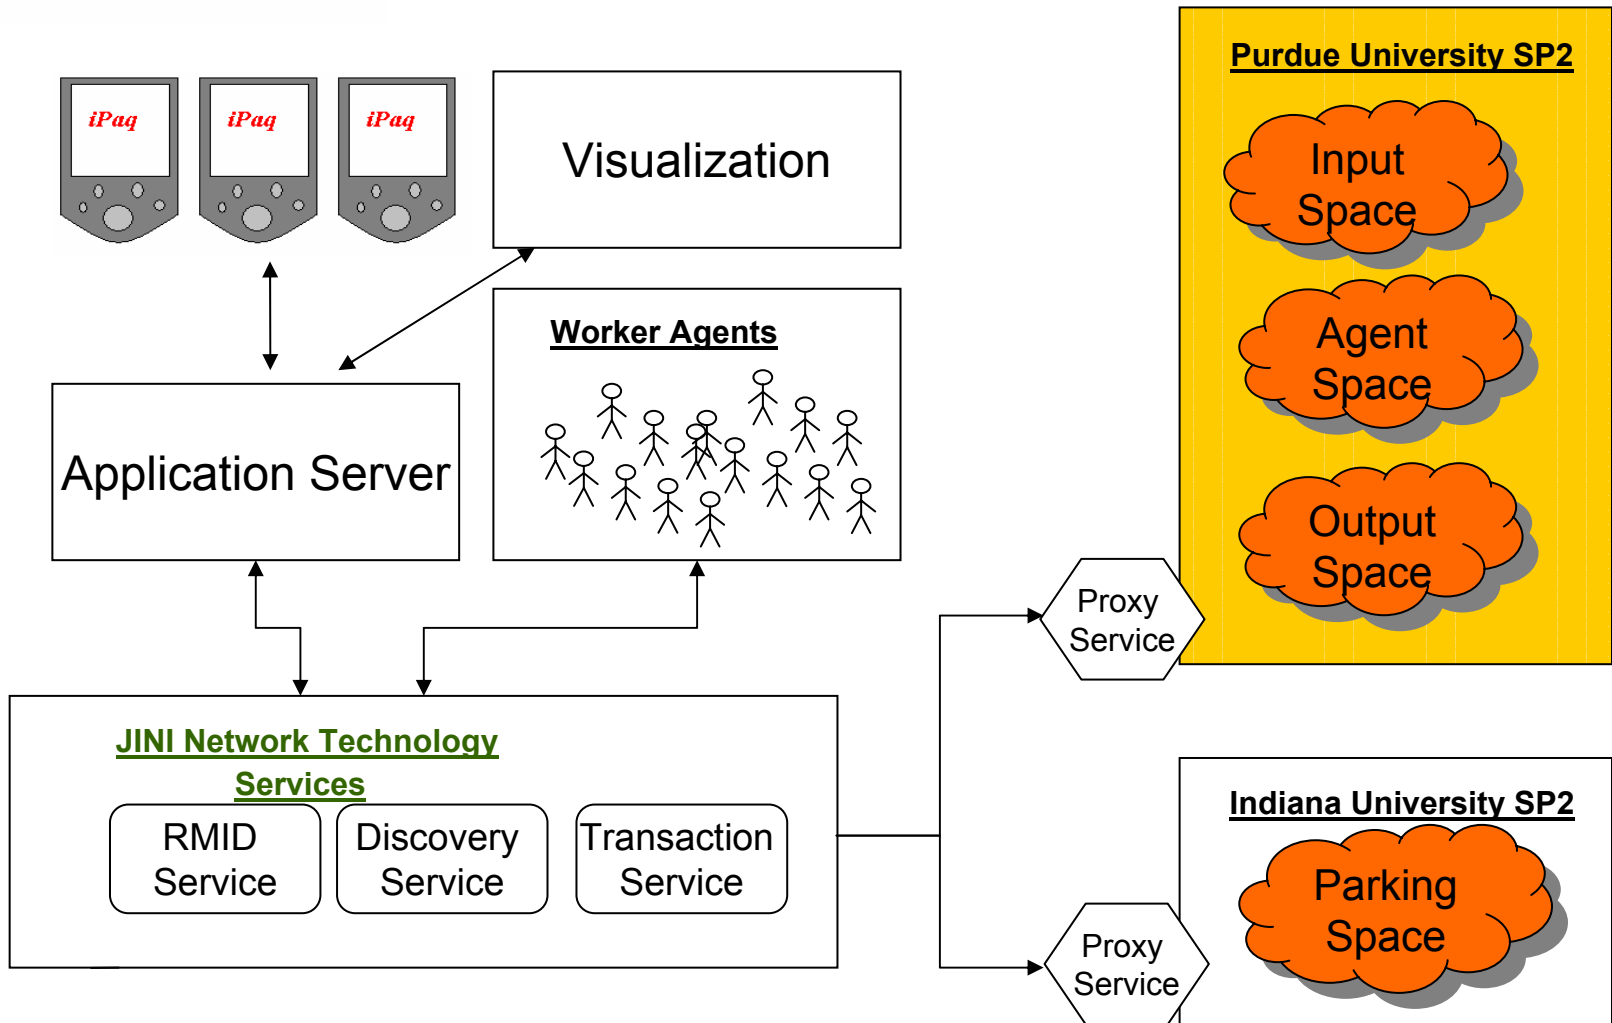
Parallel Worlds



Measured Response used configurable, distributed, grid computing on linked IBM SP2 supercomputers to connect palmtops to teraflops in a seamless simulation



Technical Architecture

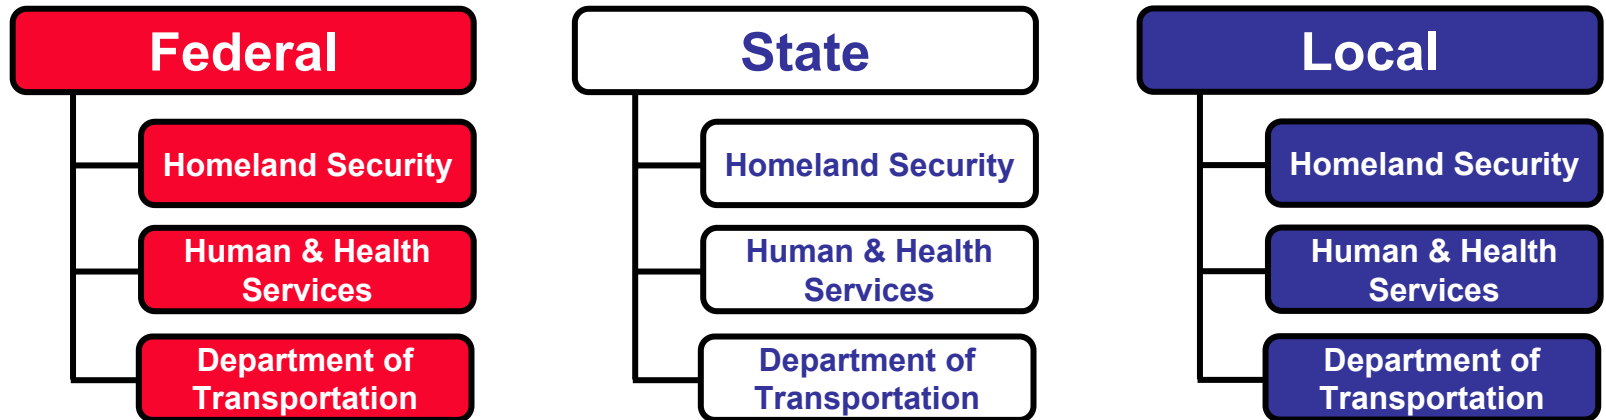





Grid computing divided each complex simulation into small processes that could be simultaneously computed...

- Model population of 250,000
- Model the spread of the disease
- Model the impact of the disease on infected people
- Model the impact of participant decisions on public perception

Human players



Artificial agents



250,000 artificial agents

“Real time” information provided to decision makers to help them determine their next steps

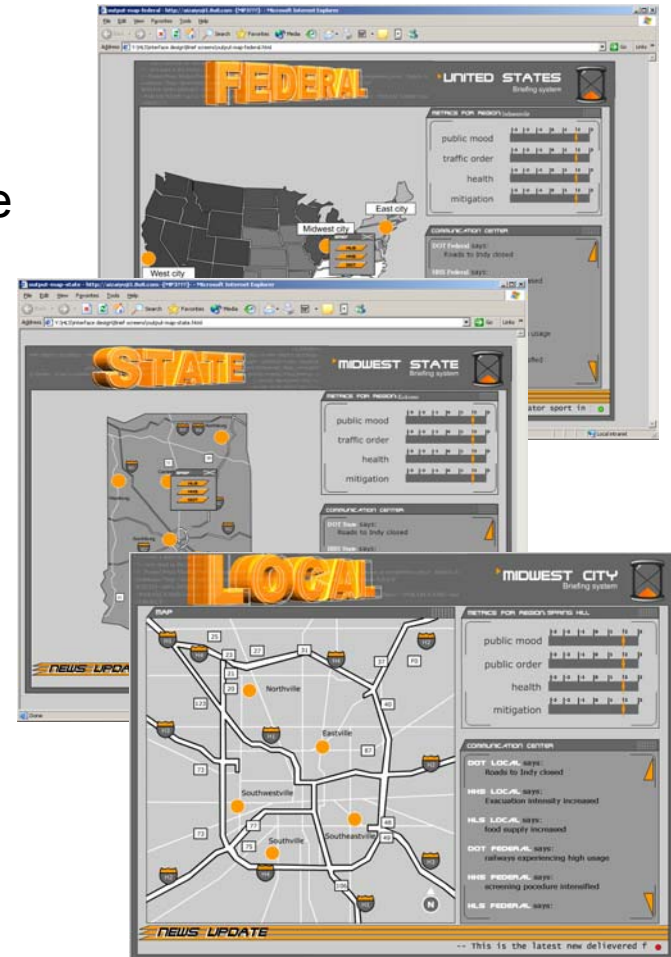
Information – What do we know? What more do we need to know? Where can we acquire it?

Analysis – How reliable is our information? What does it mean?

Judgment – What should we do about the situation? What can we do about it? What resources are available?

Communication – Whom should we inform? How? Within government? Outside government?

Decisions? Who makes them? What range of options are available? What resources are available?




FEDERAL


UNITED STATES

Briefing system

Midwest City



METRICS



COMMUNICATION CENTER

DOT Federal says:
Roads to city closed

HHS Federal says:
Evacuation intensity increased

HLS Federal says:
food supply increased

DOT Federal says:
railways experiencing high usage

HHS Federal says:
screening procedure intensified

HLS Federal says:

NEWS UPDATE

| News reports 100 dead in Midwes

<!-- URL's used in the movie-->

<!--Posted Price Market Period: Mean Price for this Period: Standard Deviation: Surplus at competitive price: Details for
codebase="http://download.macromedia.com/pub/shockwave/cabs/flash/swflash.cab?version=5,0,0,0"

WIDTH=100% HEIGHT=100%-->

<PARAM NAME=movie VALUE="postedPriceMarket.sw?=<PARAM NAME=logo VALUE=false"> <PARAM NAME=mc

</OBJECT>

UNITED STATES


Briefing system

Midwest City



Back

STATUS



SITUATION REPORT

Health threat increasing in region 3

Riots on the rise, Civilians dying

Accidents increase by 200% on the
interstate highway in city

Health threat increasing in region 3

Riots on the rise, Civilians dying

Accidents increase by 200% on the
interstate highway in city

THREAT REPORT

Health threat increasing in
southeast area of state

Riots on the rise, Civilians dying

Accidents increase by 200% on the
interstate highway in city

Health threat increasing in
northwest region of state

Riots on the rise, Civilians dying

PLAN OF ACTION

Health threat increasing in region 3

Riots on the rise, Civilians dying

Accidents increase by 200% on the
interstate highway in city

Health threat increasing in region 3

Riots on the rise, Civilians dying

Accidents increase by 200% on the
interstate highway in city


NEWS UPDATE

14 dead in local hospital | | Airplane crashes due to serious illness of pilot | | 4 arrested in




The Demonstration - Results

“officials should have dealt with life and death issues more quickly”




Time (Day) Vs. Number of People infected

The next step for **Measured Response** is two state coordinated simulation running at *three* locations on *three* supercomputers




Physico-Chemical Cell Modeling and Information Theory Data Integration



Peter Ortoleva, Distinguished Professor
Director


Center for Cell and Virus Theory
Department of Chemistry
Indiana University, Bloomington



An integrated, multi-disciplinary approach is needed to develop a quantitative, predictive understanding of the nature and origin of cellular life. In our *Karyote*[®] project, the integration is achieved through information theory.

Research Support

- Department of Energy
- Department of Defense
- Pharmaceutical Industry
- Petroleum Industry
- The State of Indiana



A eukaryotic cell is divided spatially into regions of biochemical specialization (organelles). *Karyote*[®] is a comprehensive, physico-chemical, compartmentalized cell model (Weitzke and Ortoleva 2003; Navid and Ortoleva 2003; Ortoleva 2003). Special intracellular features (e.g. mitochondria and proteosomes) evolve via molecular mechanics, mesoscopic equations solved on finite element grids, and macroscopic equations for overall mass balance.


Computational Techniques

- Slow manifold projection and multiple timescale analysis
- Transcription/translation windowing
- Finite element simulation of mesoscopic equations
- Multi-gridding and adaptive time stepping
- Matrix manipulation efficiencies
- Parallelization of information theory algorithm
- Space warping molecular mechanics
- Tree algorithms for atomic forces
- Atomic force hardware

Space Warping Method

$$\vec{r}' = \prod_n \Gamma_n^{(m)} \vec{f}_n(\vec{r}), m = 1, \dots, M$$


In the space warping method a transformation is applied to the macromolecule or parts thereof so as to allow efficient large-scale conformation changes (Jaqaman and Ortoleva 2002).




(a) Initial configuration of 1,715 atoms.

(b) Final configuration determined using local coordinates only, corresponding to a local minimum of the potential energy.


(c) Final configuration reached using the space warping method, corresponding to the global minimum of the potential energy or to a close-lying local minimum (Jaqaman and Ortoleva 2002).



Energy-minimizing configuration of $[AlO_{16}]^+$ computed using our M3 software. Predicted conformation and atomic-scale features are key to identifying drug targets (Jaqaman and Ortoleva 2002). M3 is based on our novel space warping method for macromolecular mechanics and field theoretic, mesoscopic host medium model.



A virus (HIV seen schematically) is a hierarchical, free energy-minimized structure wherein macromolecules with their three levels of internal structure are organized into a composite with the genetic and other molecules. HIV consists of some 300 million atoms.



Rendering of structure of poliovirus Capsid Type I Mahoney using NIH single protomer structure data and our *VirusX* simulator (Jarymowycz, Hubbard and Ortoleva 2003). All atoms (713,580) are shown in CPK coloring.

Mesoscopics

$$F[\xi, c] = U(\xi) + \int d^3r f$$


$$f = f^{cl} + \frac{1}{2} \Lambda |\vec{\nabla} c|^2 + \varphi(\xi, \vec{r})$$

$$\frac{\partial F}{\partial c(\vec{r})} = \bar{\mu}$$

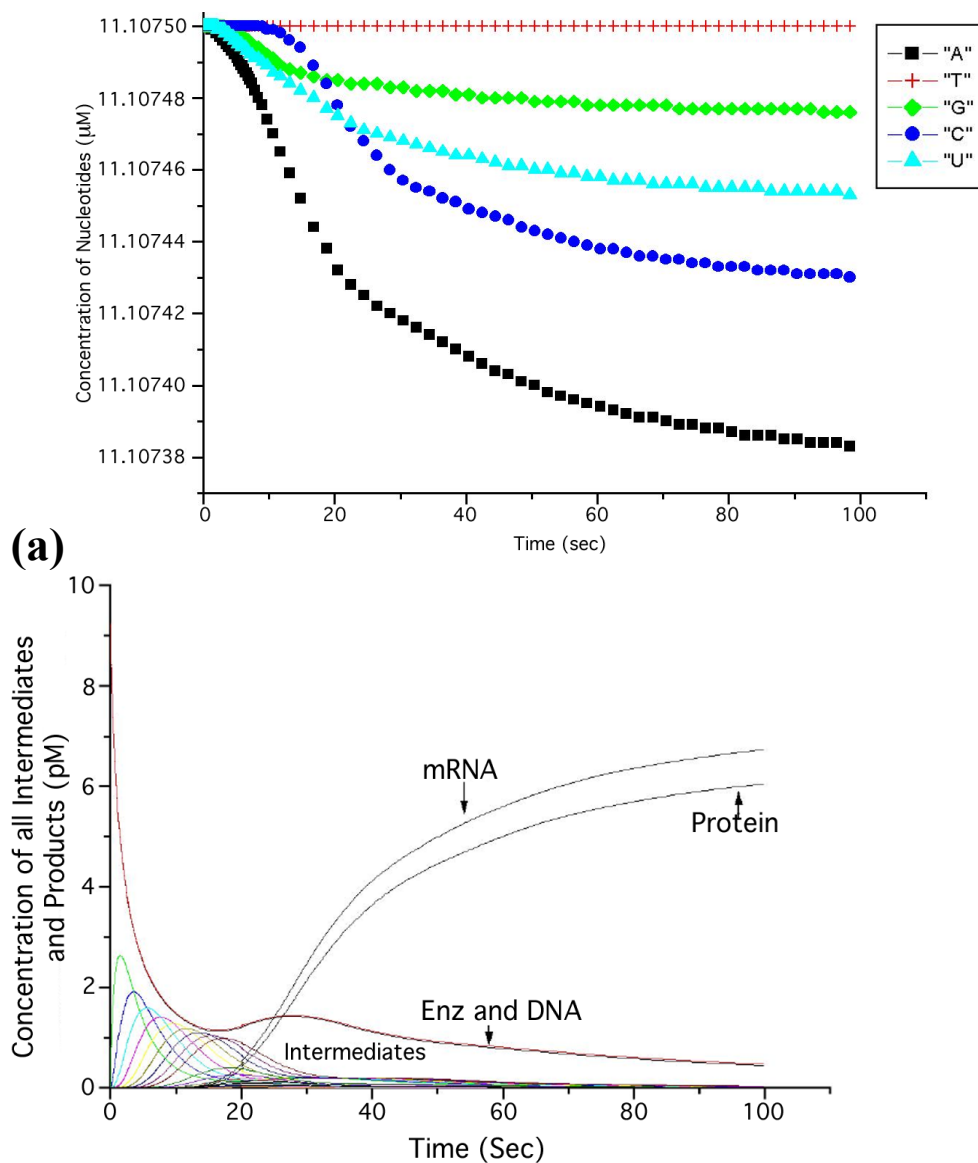
$$\frac{\partial F}{\partial \xi} = 0$$

Equilibrium structure is obtained by minimizing F with respect to the distribution of host medium density in the vicinity of the macromolecule of structure ξ (Ortoleva 2003).

Prokaryote Genomic/Proteomic Reaction Network



A schematic depiction of prokaryotic, coupled transcription/translation. The above complex transcription/translation network is simulated using chemical kinetic rate laws. The result is a set of ordinary differential equations for the evolution of the populations of *RWE*, proteins, *mRNA* and other species and their complexes (Weitzke and Ortoleva 2003; Ortoleva 2003).



(a) *Karyote*[®] predicted time dependence of nucleotide concentrations during transcription for the gene *TACTTTTAGGGG*. As nucleotides are depleted, *mRNA* synthesis slows down, illustrating an important feature of coupled genome, proteome, metabolome dynamics captured in *Karyote*[®].

(b) Evolution of concentration over time of *DNA*, *RNA* Polymerase I, all of their complexes involved in transcription/translation and *mRNA* created under the same condition as seen in (a) (Weitzke and Ortoleva 2003).

Compartmentalized Cell Dynamics

$$V^{\alpha} \frac{dc_i^{\alpha}}{dt} = \sum_{\alpha' \neq \alpha}^{N_c} A^{\alpha\alpha'} J_i^{\alpha\alpha'} + V^{\alpha} \sum_{k=1}^{N_r} \nu_{ik}^{\alpha} W_k^{\alpha}$$

$A^{\alpha\alpha'}$ = shared boundary surface area
separating compartments α and α'

$J_i^{\alpha\alpha'}$ = net flux of species i from α' to α ($= -J_i^{\alpha'\alpha}$)

N, N_c, N_r = numbers of chemical species,
compartments and reactions, respectively

V^{α} = volume of compartment α

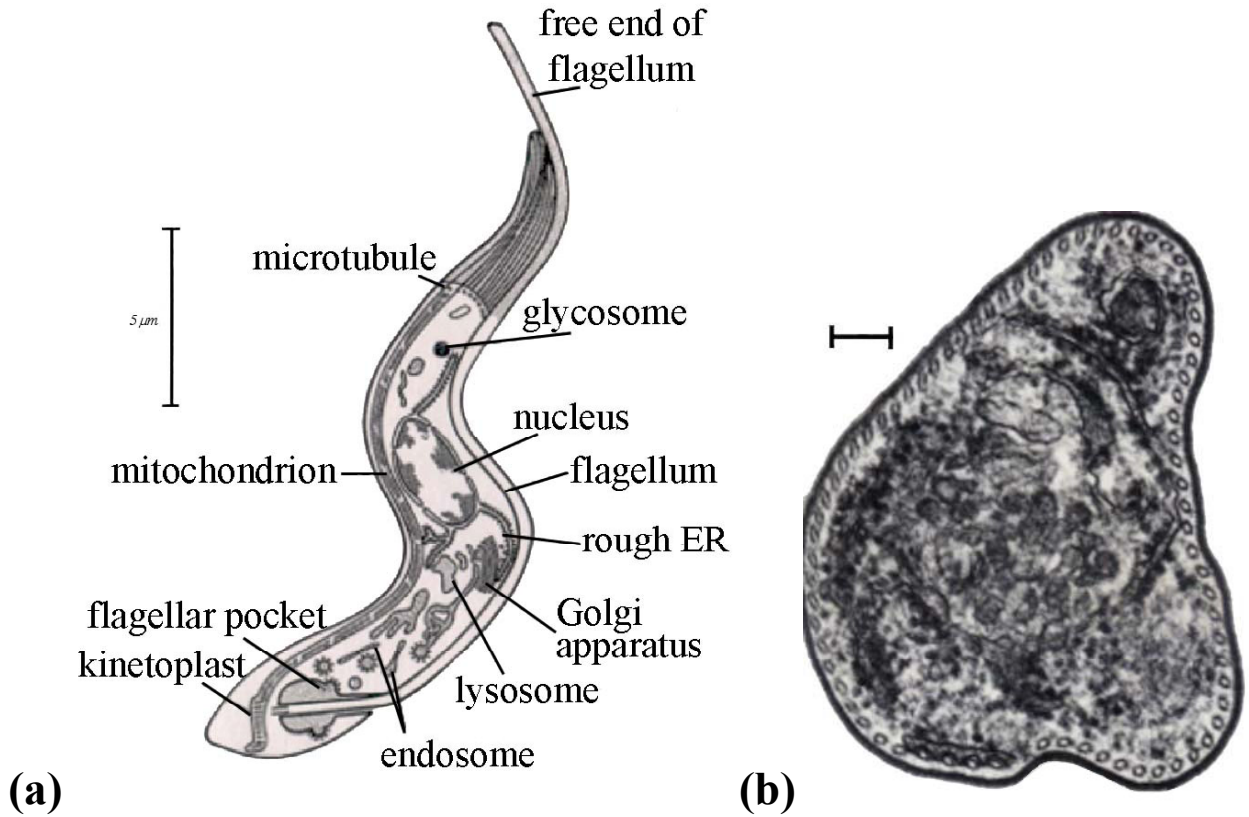
W_k^{α} = rate of reaction k in compartment α

ν_{ik}^{α} = stoichiometric coefficient for species i
in reaction k in compartment

α ($\nu_{ik}^{\alpha} > 0$ for products).

(Weitzke and Ortoleva 2003; Ortoleva 2003)

Forms of the Pathogen *Trypanosoma brucei*



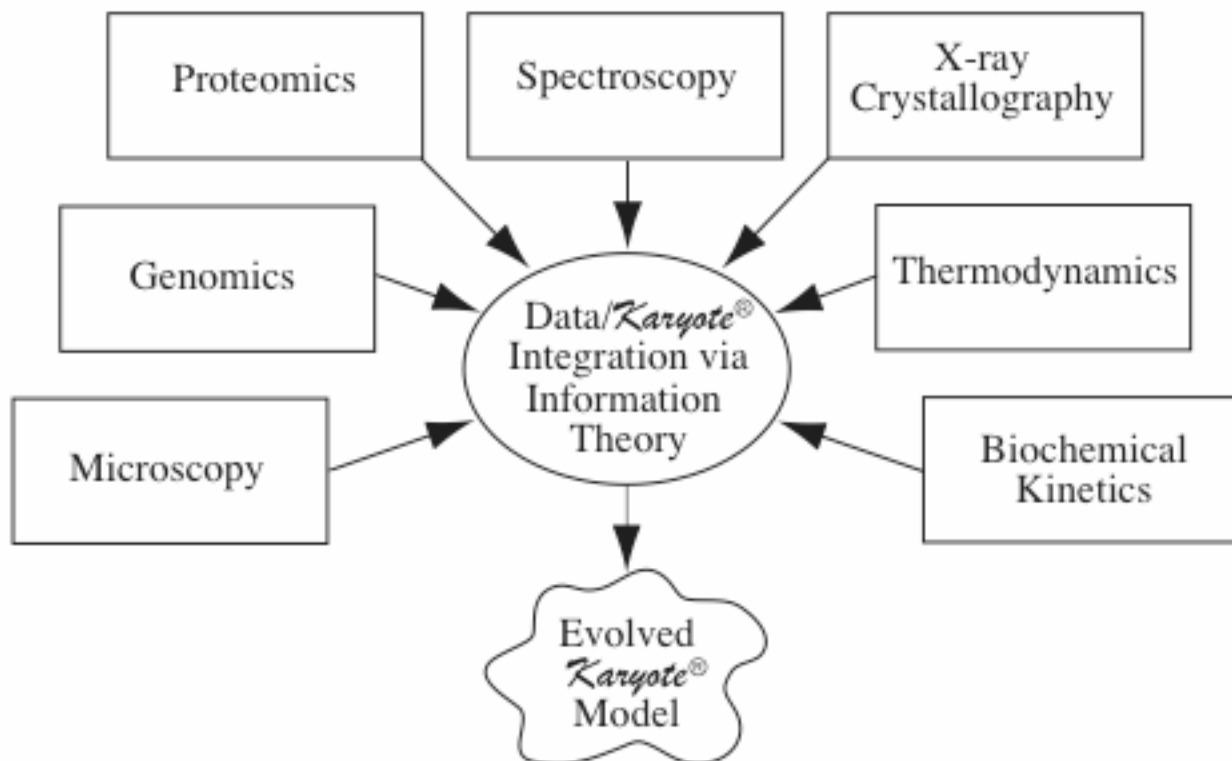
(a) In the slender form with a long flagellum (Kang 2001), the single mitochondrion is forced into a peripheral canal and has few cristae, no cytochromes, and the citric acid cycle does not function.

(b) The stumpy form (Aman et al 1985) has an expanded mitochondrial canal and thereby participates in metabolism (see Oppenheimer 1987).

Using *Karyote*[®], we have modeled the aerobic and anaerobic glycolysis of the above parasite (Navid and Ortoleva 2003).

Species	Experimental concentration (aerobic) (mM)	<i>Karyote</i> ⁰ concentration (aerobic) (mM)/ % error (Navid and Ortoleva 2002)		Calibrated <i>Karyote</i> ⁰ concentration (aerobic) (mM) / % error		Bakker et al. concentration (aerobic) (mM)/ % error	
G6P	4.4	1.0	77	4.40	0	0.44	90
FBP	2.4	0.55	77	2.41	1	0.13	95
F6P	1.9	1.4	26	1.93	2	26	1268
GAP	0.47	0.25	46	0.28	40	0.074	84
DHAP(g/c)	2.6	3.8	46	4.26	64	1.6	38
1,3-BPG	0.77	0.2	74	0.74	4	0.028	96
3PG(g/c)	4.8	1.7	65	4.98	4	0.68	86
2PG	0.59	0.3	49	0.60	2	0.13	78
PEP	0.85	2.0	135	0.91	7	0.85	0
Pyruvate	21	21.6	3	20.7	1	21	0
Gly-3-P(g/c)	2	0.4	80	1.68	16	1.1	45

Comparison of calculated and measured steady state metabolite concentrations for aerobic glycolysis in *T. brucei*. In column 1 are measured concentrations (Visser and Oppendoes, 1980). In column 2 are results for *Karyote* simulation of aerobic glycolysis in *T. brucei*. In column 4 are results of a Bakker et al (1997) simulation for the same system. *Karyote*[®] results have smaller average margin of error in comparison to those of Bakker et al. Column 3 shows improvements in results after calibration of *Karyote*[®] using 13 parameters (3 membrane transport, 8 slow reactions rate coefficients, and 2 fast reactions equilibrium constants). (g) and (c) denote glycosomal and cytosolic concentrations respectively.



Multiple types of data are integrated with *Karyote*[®] via our information theory methodology to automatically yield improvements in *Karyote*[®] as new data becomes available (Sayyed-Ahmad, Tuncay and Ortoleva 2003).

Information Theory


Through the maximum entropy principle, we construct the probability ρ

$$S = - \sum_{\psi} \rho \ln \rho$$

ψ : Set of least-well constraint parameter in the model

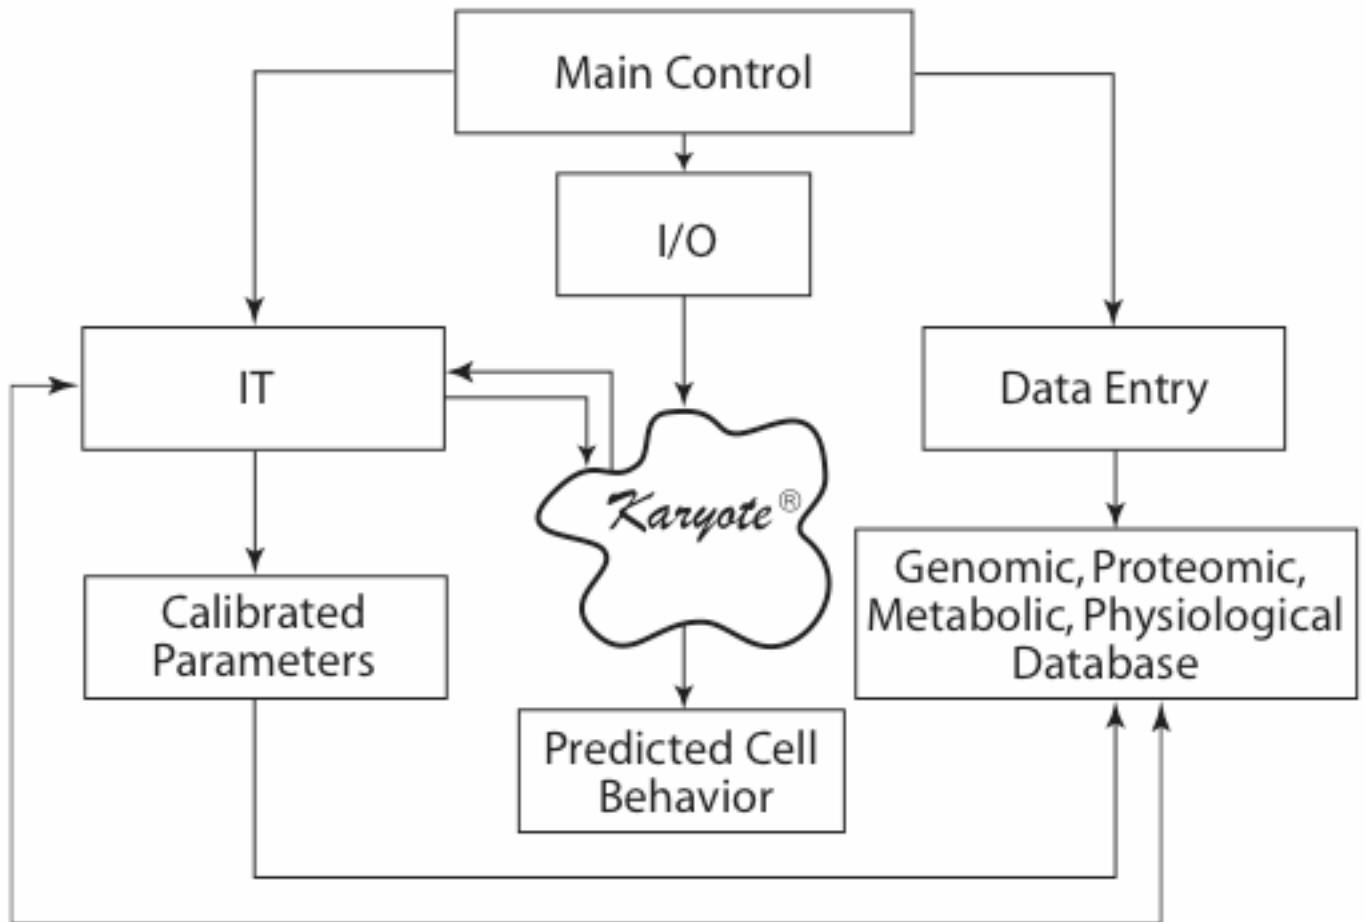
(Sayyed-Ahmad, Tuncay and Ortoleva 2003)

Predicted and Observed Concentration Showing Uncertainty Envelope



Time dependence of the concentration of a species in an incomplete model using an information theory-based data model integration approach (Sayyed-Ahmad, Tuncay and Ortoleva 2003).

Overall *Karyote*[®]/Data Integration Flow Chart



This structure underlies our web-based cell modeling system that allows for a cell database, automated running of the cell model, and the use of information theory to calibrate and guide the model via integration with experimental data.

Multi-mode Architecture for Distributed Thin-client Collaborations

Karthik Ramani

Center for Information Systems in
Engineering

Purdue University, West Lafayette, IN

Acknowledgements

STUDENTS

- **Abhishek Agarwal**
- **Mahendra Babu**
- **Nikhil Joglekar**
- **Aliasgar Ganiji**
- Ajoy Raj
- Vineet Agarwal
- Srinivasan Turuvekere
- Kuiyang Lou
- Natraj Iyer
- Liu Min
- Zhanjun Li
- Jayanti Subramaniam
- Chan Woo Chung

- Sitaram Karthik
- Joon Hong
- Alexander Lee

FACULTY

- **Chris Hoffmann (CS)**
- Sunil Prabhakar (CS)
- Suresh Jagannathan (CS)
- Xianfan Xu (ME)
- David Anderson (ME)

- Jim Bottum (ITap)
- David Moffet (ITap)
- Abhyankar (Maths)

COLLABORATORS

- Ram Sriram (NIST)
- Simon Szykman (DC)
- Suresh Sitaram (GTech)

Acknowledgements

Sponsors

- Defense Logistics Agency
- NSF – Partnership for Innovation Award (NSF-PFI)
- NSF-IGERT (Innovation Realization Lab)
- NSF-MRI (Envision Center for Perceptualization)
- American Foundry Society
- State of Indiana 21st Century Award
- University Faculty Scholar Award
- E-Enterpirses@Discovery Park

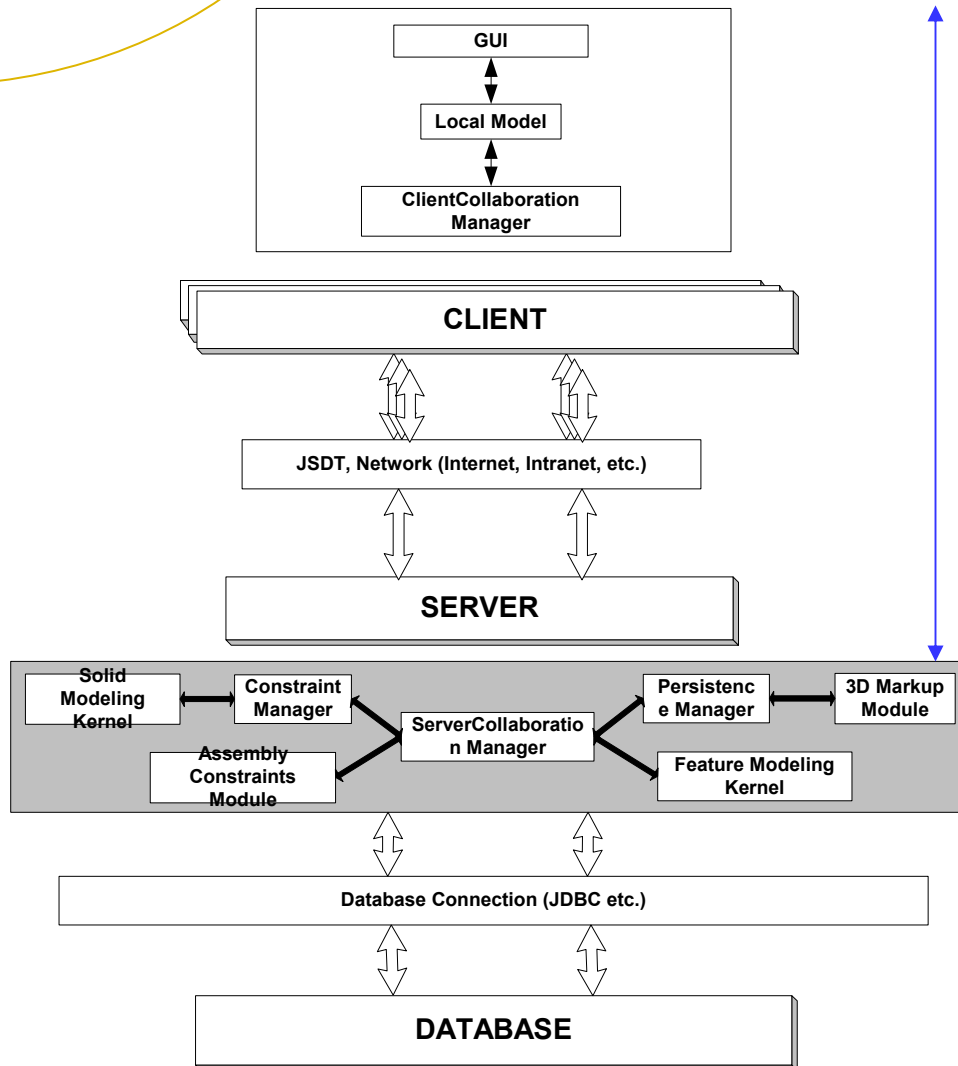
Structure of the Presentation

- CADDAC design goals
- Architecture of CADDAC
- Components of CADDAC
- Modules developed using CADDAC
- Conclusions

CADDAC design goals

- Collaborative shape generation and management
 - Support remotely located team of designers
 - Immediate feedback from other people
- Web-based system
 - Easily accessible, cost effective, platform-independent
- Thin-client
 - Easy download, installation, minimal resources at the client
- Restricted privilege
 - Corporate security, intellectual property

CADDAC architecture



Three-tier architecture:
 Client, Server, Database


Server-side:
 Computational power,
 Collaboration

Client-side:
 Interface with user

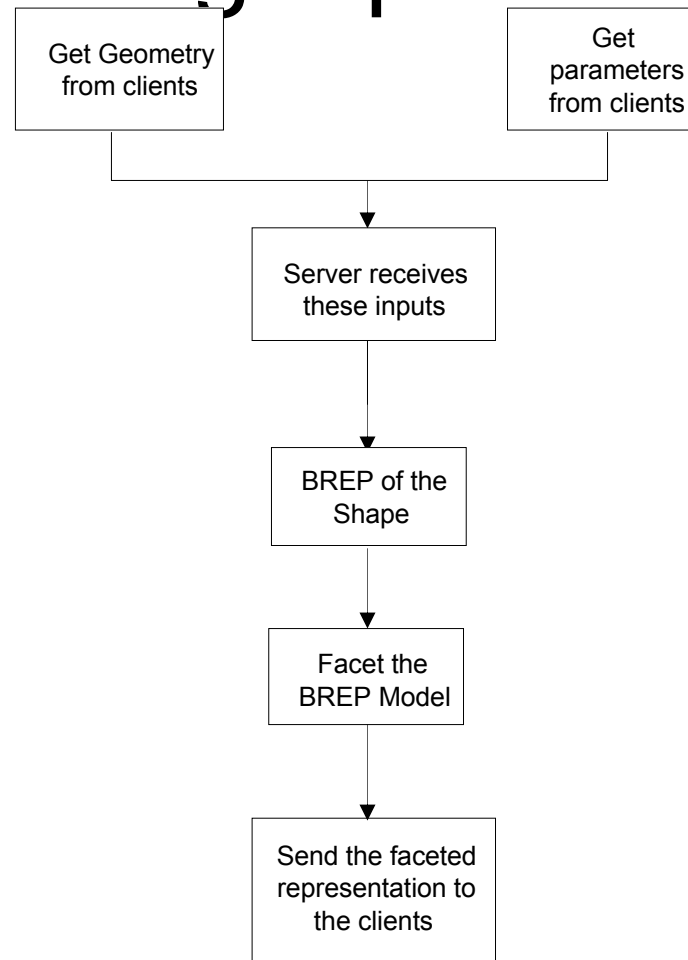
Database:
 Store, search, and retrieve

Java <-> C++

Expandable Plug-in

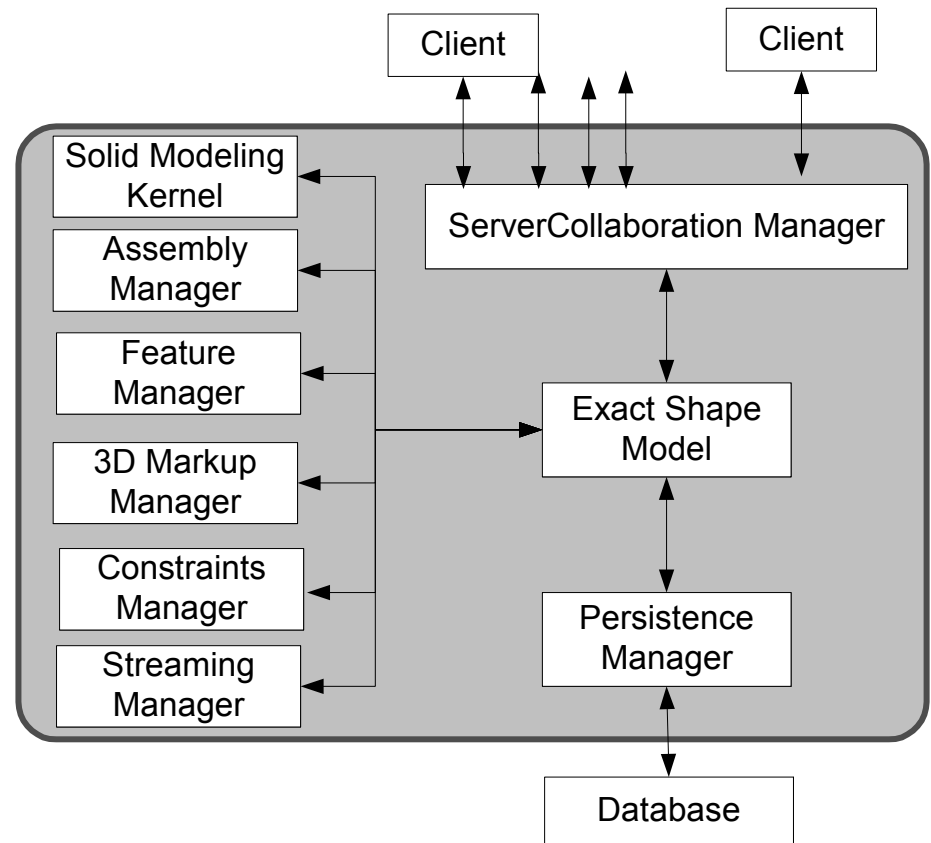


Information flow with Solid Modeling Operations

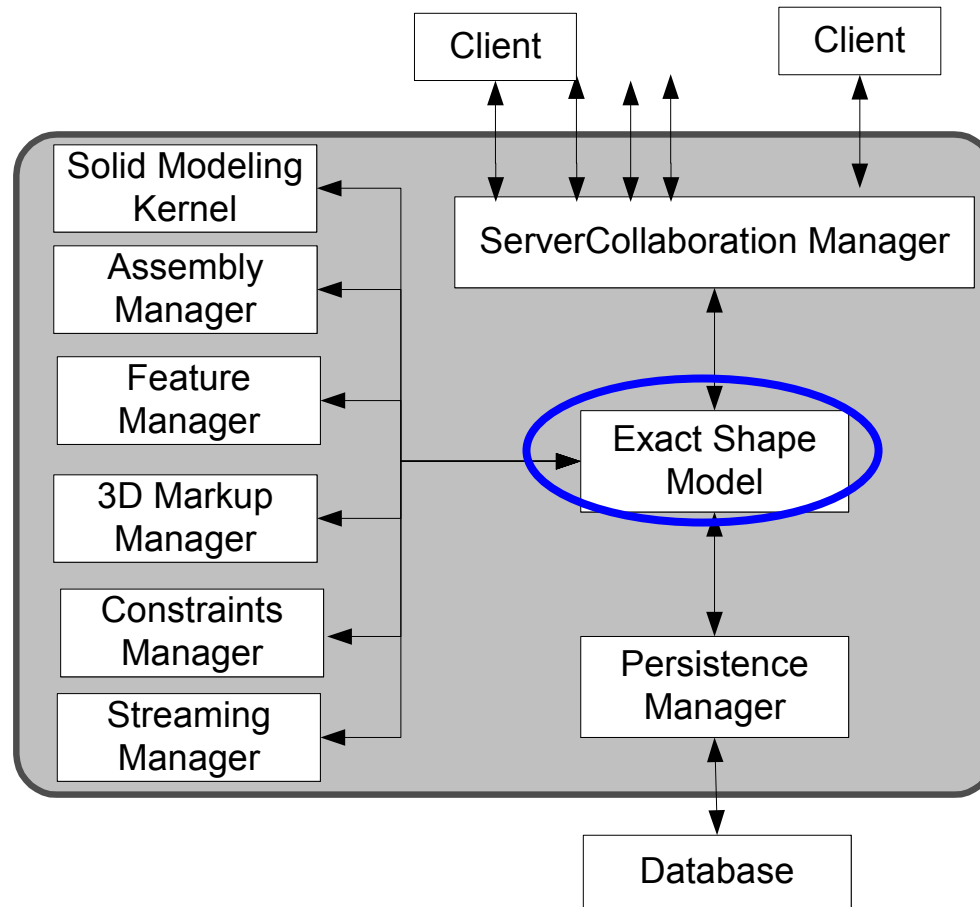


CADDAC Server

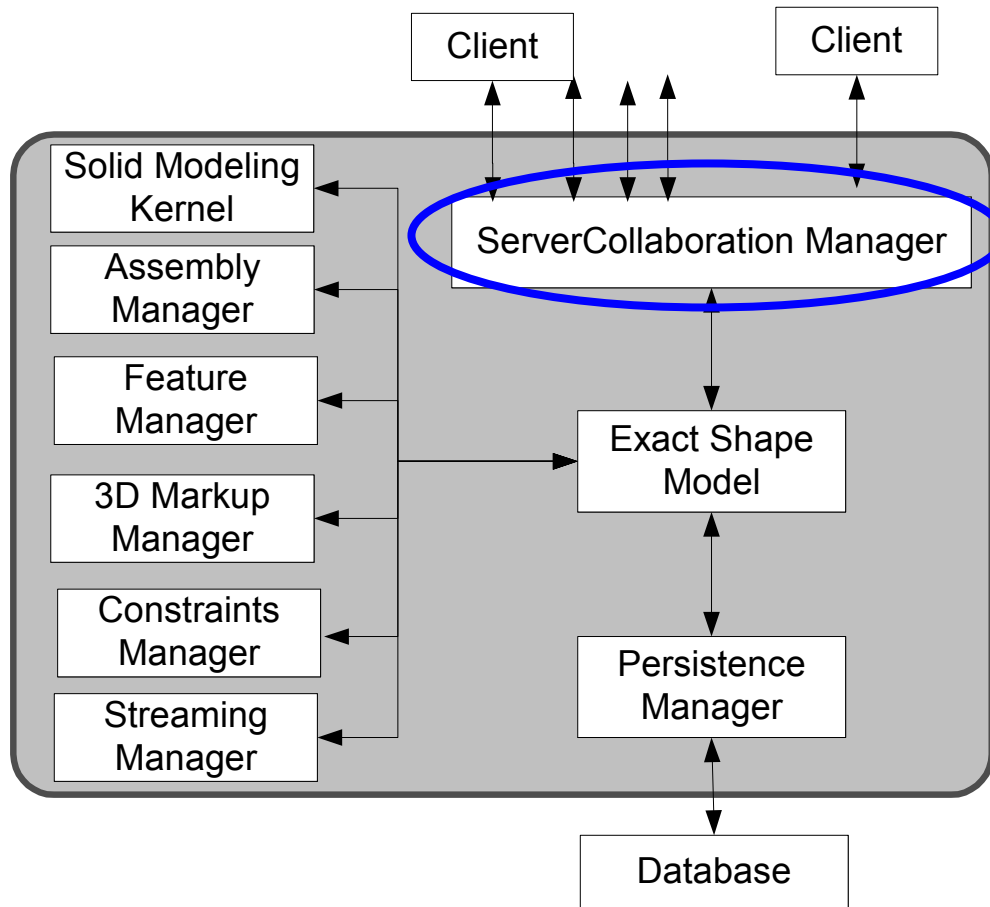
- Master Model
- Geometry Kernel
- Assembly Manager
- Feature Manager
- 3D Markup Manager
- Constraints Manager
- Streaming Manager
- Persistence manager



Exact shape model



Servercollaboration Manager



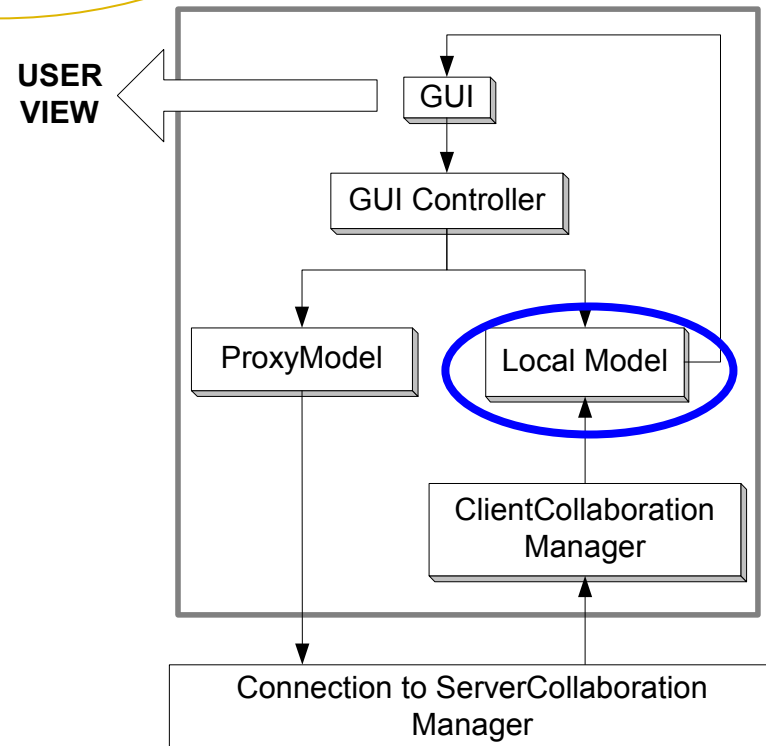
Servercollaboration Manager

- Multiple Session Management
- Receives Command objects from clients
- Command objects represent a method with arguments
- Method from the model is called if the corresponding Command object is executed

CADDAC client

- Functionalities
 - Creating, editing, deleting geometry and geometric constraints
 - Viewing and navigating through 3D models
 - Rotation, zoom in-out, panning
- Thin-client
 - Only Magician (graphics routines) installation
 - Standard Internet browser based

Client architecture



- Model/View/Controller based
- Model: Local Model
- View: GUI
- Controller:
 - GUI Controller & ClientCollaboration Manager

The diagram illustrates the Client Collaboration Manager architecture. It features a central box containing several components: **GUI**, **GUI Controller**, **ProxyModel**, **Local Model**, and **ClientCollaboration Manager**. The **ProxyModel** and **ClientCollaboration Manager** components are highlighted with red and blue ovals, respectively. The **GUI** component is connected to the **GUI Controller**, which in turn connects to both the **ProxyModel** and the **Local Model**. The **ClientCollaboration Manager** is connected to the **Local Model**. A large arrow labeled **USER VIEW** points from the **GUI** component to the left. A **Connection to ServerCollaboration Manager** is shown at the bottom, with arrows pointing to the **ProxyModel** and the **ClientCollaboration Manager**.

CADDAC client

- ClientCollaboration Manager
 - Receives incoming Command objects from server
 - Modify the local model based on the received Command object
- ProxyModel
 - Sends Command objects to the server
- Two modes of synchronization
 - Sending data from one client to another
 - Operations involving computations at the server

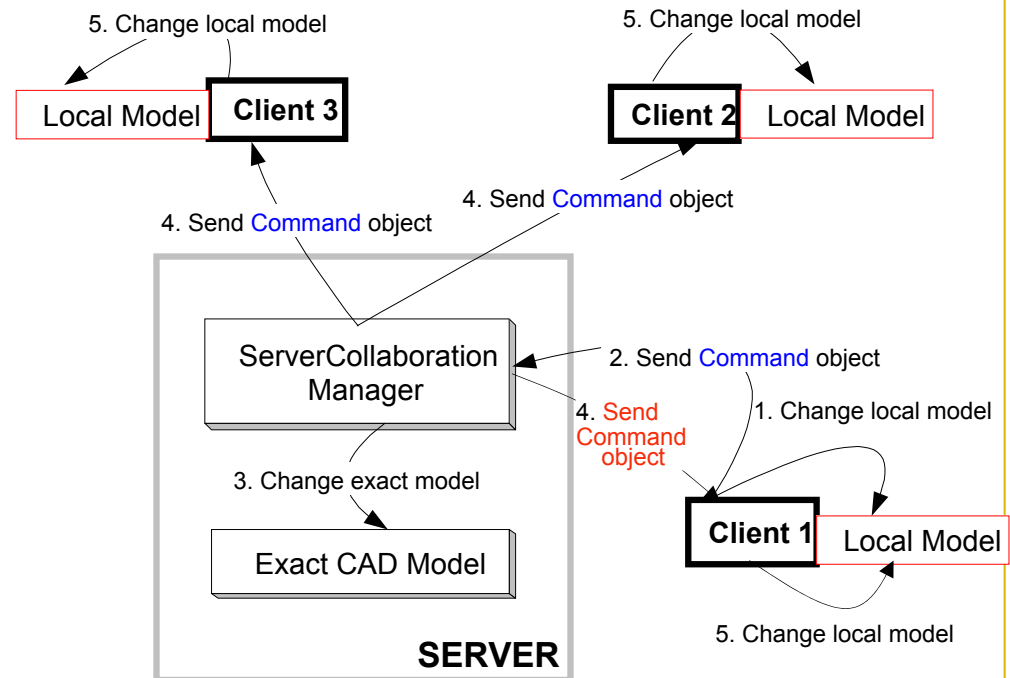
No computation at server

-
- The diagram illustrates the interaction between a **SERVER** and three clients (**Client 1**, **Client 2**, and **Client 3**) during a model update process. The **SERVER** contains a **ServerCollaboration Manager** and an **Exact CAD Model**.
- Sequence of Events:**
- Client 1** sends a **Command** object to the **ServerCollaboration Manager** (labeled "1. Change local model").
 - The **ServerCollaboration Manager** sends a **Command** object to **Client 2** (labeled "2. Send Command object").
 - The **ServerCollaboration Manager** sends a **Command** object to **Client 3** (labeled "3. Change exact model").
 - The **ServerCollaboration Manager** sends a **Command** object to **Client 2** (labeled "4. Send Command object").
 - Client 1** changes its local model (labeled "1. Change local model").
 - Client 2** changes its local model (labeled "5. Change local model").
 - Client 3** changes its local model (labeled "5. Change local model").
- The diagram shows that the **ServerCollaboration Manager** acts as a central hub, distributing **Command** objects to the clients and managing the **Exact CAD Model** within the **SERVER**.


Synchronization process

- Computation at server



1. Client1 → Local Model
2. Client1 → Server
3. Server → Exact Model
4. Server → All Clients
5. Other Clients → All Models



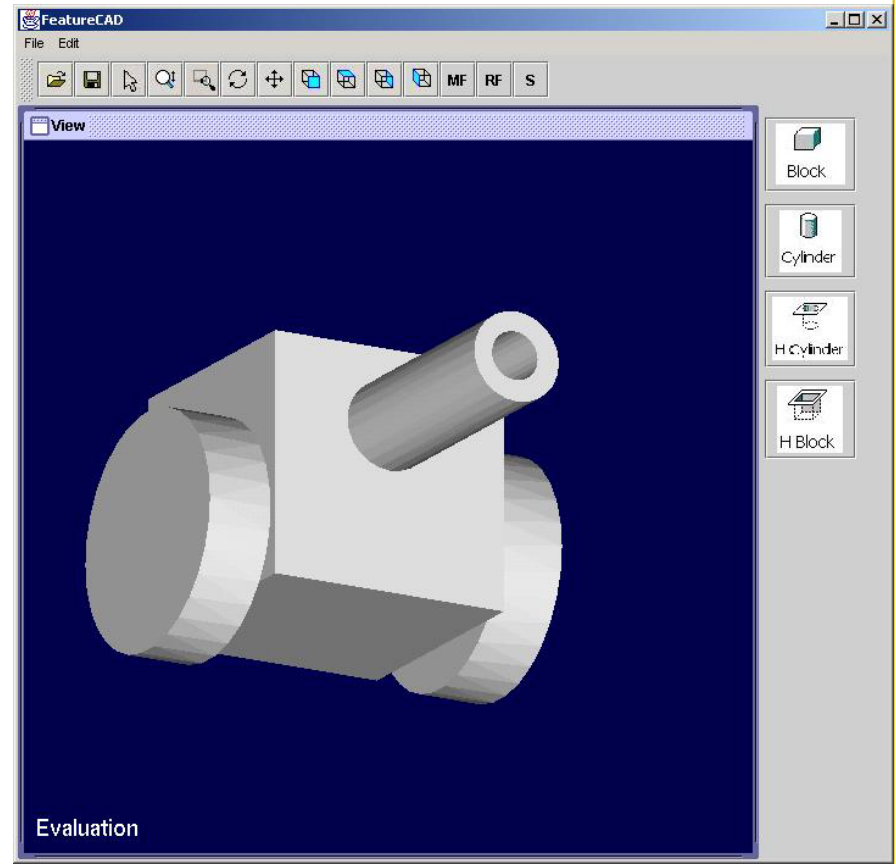
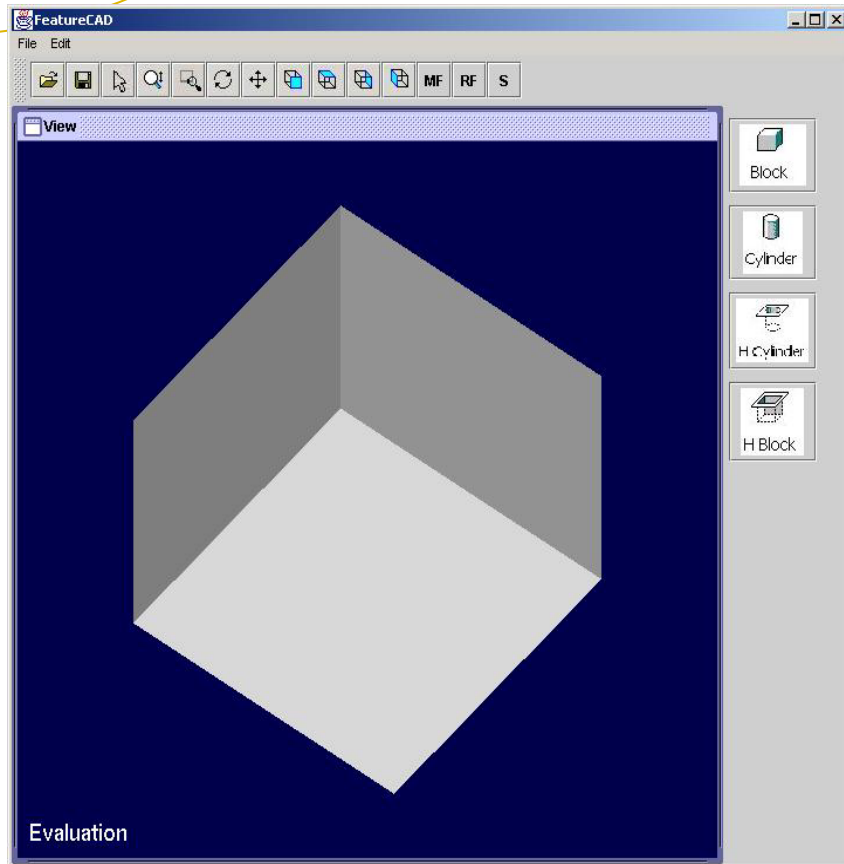
Sample Part created




Assembly



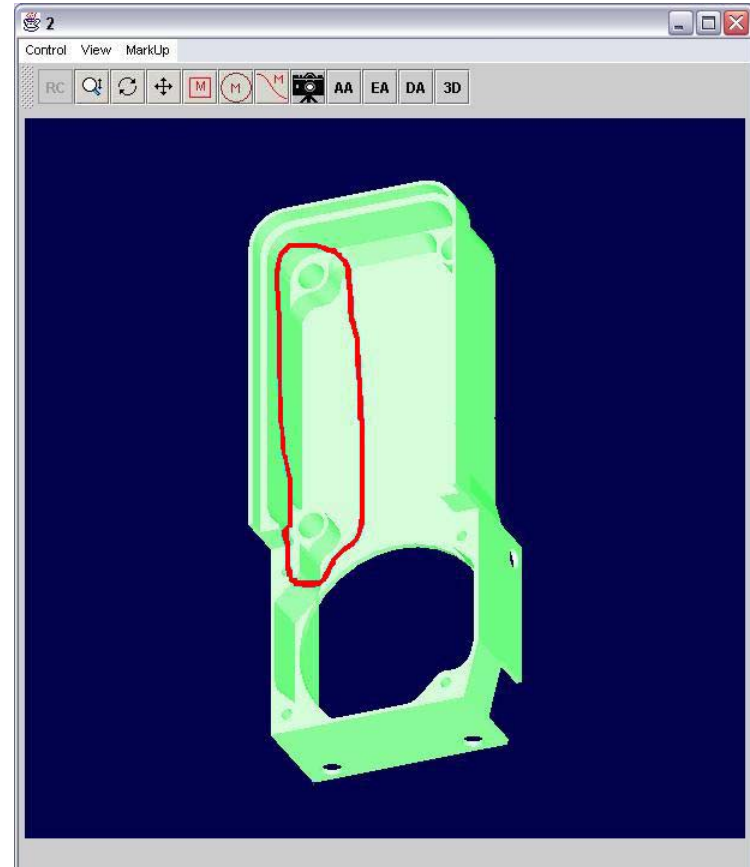
Feature Modeling



3D Markup



Unmarked object




Object with 3D markup

Conclusions

- Thin client-side achieved
- Use of MVC, Command, and Proxy patterns allows expandable architecture
- Web-based collaborative 3D modeling system with centralized geometry kernel
- Ability to quickly develop new plug-in features with the same collaborative framework

Questions ?


Thank you !



Pixels Everywhere!

Dr. M. Pauline Baker

Associate Professor,
Computer Science and School of Informatics
Distinguished Scientist and Director,
Visualization and Interactive Spaces Lab




pervasive**technology**labs

AT INDIANA UNIVERSITY


Visualization and Interactive Spaces

- Design and build deployable physical environments for interactive data display and exploration, and for learning
- Explore use of computer-generated graphics to answer questions about data, to support data exploration and learning
- Experiment with user interaction techniques



VisBox

- Developed by VisBox Inc.
- 8' x 6' screen, passive stereo





*Images from Paul Rajlich,
Albert Khakshour (VisBox Inc.)*



Visualization and Interactive Spaces
Pervasive Technology Labs at Indiana University



TableTop Display

- Top-down projection to table
- Plasma panel for additional display
- Optical tracking of user symbols




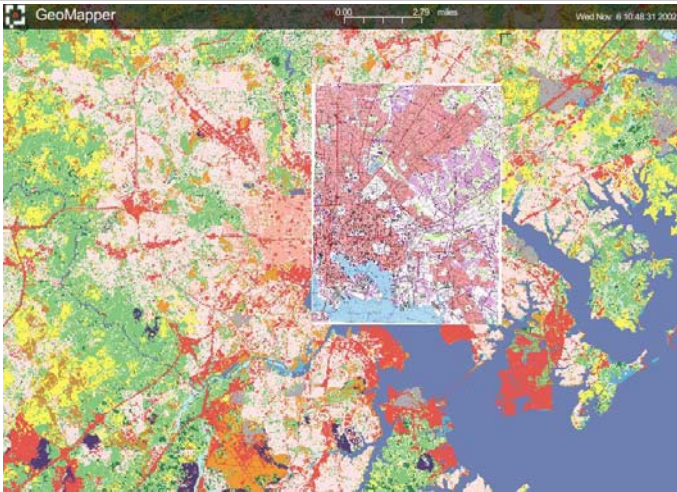
Multi-User Tool Tracking

- Cardboard symbols with retro-reflective tape shapes
- Infrared illuminators
- Infrared-only camera
- Image-processing to identify shape




Applications?

- Satellite imagery
- Microscope images
- Maps of all kinds, especially multi-layer maps
- Data collected from sensor networks
- Etc.




Restaurant Finder Application

- For Supercomputing 2002 demonstration
- Users position symbols to select restaurant categories and select individual restaurants
- Triggers browser to get Mapquest review of restaurant



- Thanks!
- For more information, see:

vis.iu.edu



Who

Genomics Data and Computational Sharing with I-Light



Rick Westerman
Genomics Center

westerman@purdue.edu

765-494-0505

Service not research

Make datasets and computational resources ...

- Easily accessible to the lab-bench Biologist
- An invisible fit into existing data generation & analysis pipeline [8000+ sequences/month]

20% done?



- De-novo 3D folding (protein & RNA)
- Phylogenetic analysis (evolution)
- Array analysis
- **Sequence comparison (string matching)**
 - Align → Predict same 3D shape → Predict same function
- **Multitude of smaller programs**
 - Server resources on the desktop

Datasets

- Many different biological public datasets
 - ✓ Sequence (nucleotide and protein)
 - ✓ Structural
 - ✓ Enzymatic
 - ✓ Ontology (relationships between data)
 - ✓ Motifs and families
- 100+ Gigabytes.
- Need to format for different programs
- Keeping up to date – space is not a problem
- Private & semi-public

Most of the data needed to run the GCG and Emboss sequence analysis suites of programs – hosted by IU.

- 68 GB [compressed] NFS-shared from IUbio to Purdue Genomics' Sun computers
- Share in existence from I-Light's inception.
- Speed is close to local hosting.
- Not complete data but close.
- Updates approximately quarterly.
- Saves ~\$3000/year.

For the BLAST sequence comparison program

- Purdue picks up datasets from IU's bio-mirror.
- Combines these with other datasets [ex: Rice genome]
- Purdue reformats the datasets
- Purdue provides NFS and SMB mounts
- 42 GB [compressed]
- Permanent mounts on Purdue's supercomputers & Linux cluster
- Permanent mount on Genomics' Suns
- Have used with PCs and Macs [OS 10.2]

Possible future data sharing

- Grid technology – pick up small chunks as needed
- Wider use of the biodata disk
- Private & semi-public datasets
- HMMer program data
- More GCG & Emboss data

2 main sequence analysis programs

BLAST

HMMer

BLAST example

64 sequences [40 KB] with largest dataset [HTG – 10GB uncompressed] & longest program [TblastX]

2 hours, 37 minutes (wall clock) on 16-CPU
Winterhawk SP supercomputer – 34 CPU hours

Pipeline: 8000 sequences/month [400 working day]

Many datasets

Re-analysis as updated datasets or
combinations of sequences become available

Need to explore more of Purdue's resources


- “Krypton” Linux cluster [limited memory]
- Regatta (soon)
- Modern Linux cluster (eventually)
- Teaching lab clusters (soon)

Tying into IU's excellent GCG and Emboss interfaces?

Grid technology for resource sharing?

Final

Turn sequence, array, and proteomic data into information




Rick Westerman

westerman@purdue.edu

765-494-0505


Discovery in the Classroom:

New Stars in Galaxy Andromeda!



IUB freshman search
for novae in the
Andromeda Galaxy

Catherine Pilachowski/IUB Astronomy
I-Light Applications Workshop, 4 Dec 2002



With thanks to:

- UITS Telecommunications Division
- Instructional Support Services through an Active Learning Grant.

Their assistance is gratefully acknowledged!

Goals

Integrate education and research in the undergraduate curriculum



Build a course curriculum around the discovery of new knowledge in astronomy

Help students experience the process of science in an authentic and meaningful way





Galaxy Andromeda Students




Finding Novae in Andromeda





Students compare images of Andromeda taken at different times to find "new stars" - novae which explode as new hydrogen gas falls on their surfaces. These novae gradually fade away as the explosion subsides.



A Nova Discovered by Students



Observing from the IUB-Astronomy Remote Observing Center




IT Resources Needed for Remote Observing

- IP Videoconferencing
- Telemetry
- Instrument/Telescope Control
- Data Transfer



IP Video



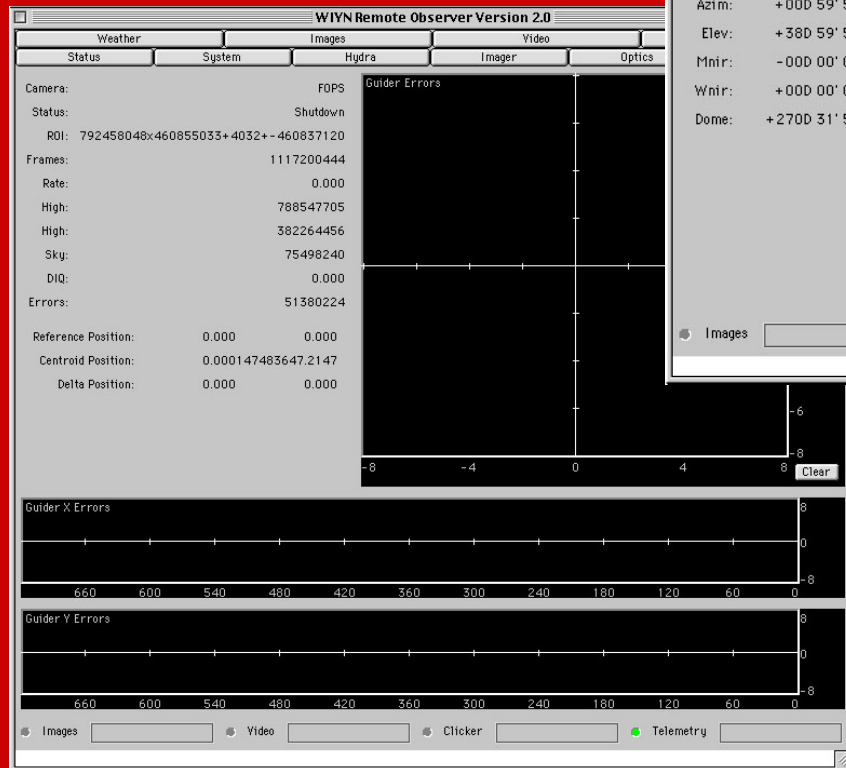
- Real-time interaction with on-site observers
- Students can participate in planning observations
- Just like being there...

Telemetry



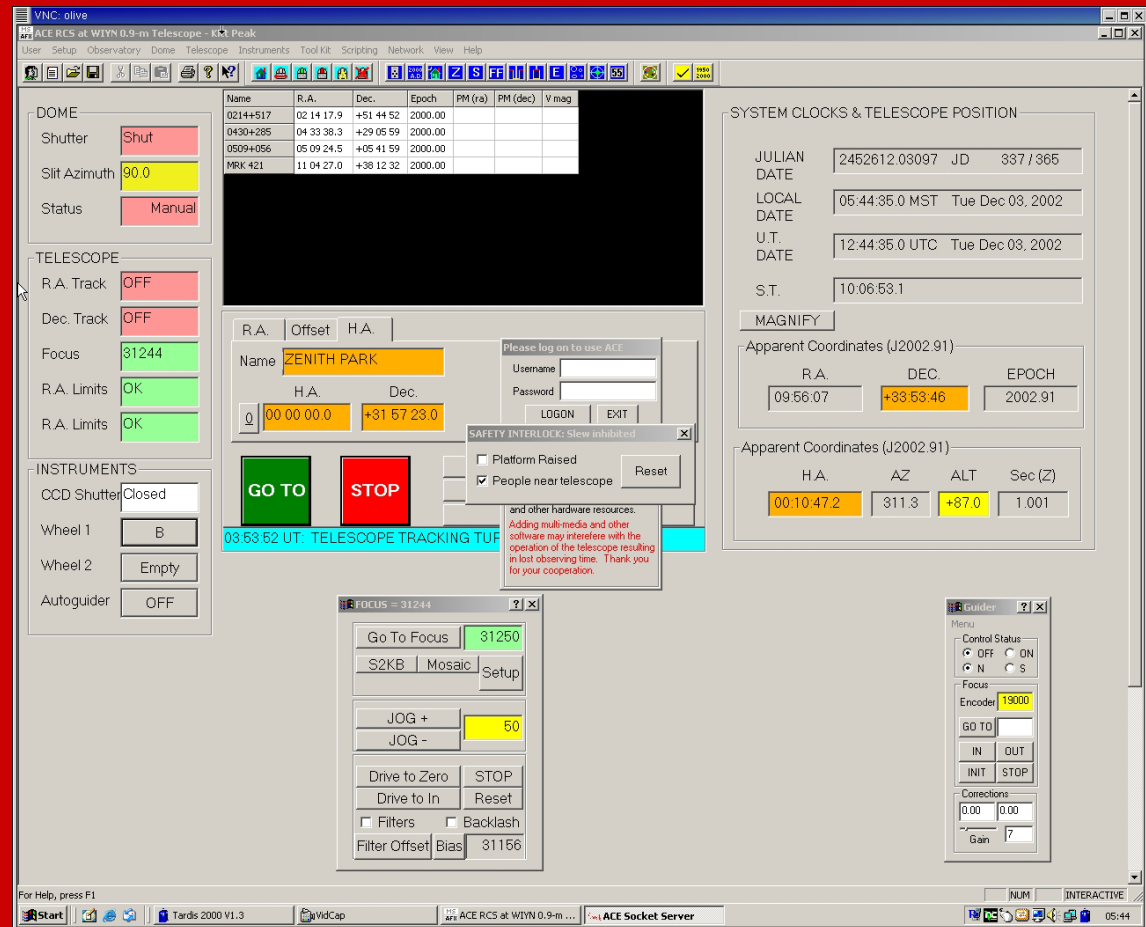
- **WiynRo** - Custom Java tool developed by a team led by Jeff Percival of the University of Wisconsin
- **WiynRo** reaches out across the Web to bring information to the astronomer. It gathers weather maps from NOAA and other sources of weather information, and collects WIYN observatory status information from a number of computers while providing a simple integrated GUI for the astronomer.

WrynRo Remote Observing Tool




Telescope/Instrument Control

- The remote observer can control the instruments on the telescope to take data
- VNCViewer provides a virtual console for instrument control



Data Transfer



- Automated, secure transfer of digital image data
- Instrument CCDs record images and spectra
- Image sizes of 4, 16, and 64 megapixels, 16-bit

Discovery in the Classroom: Goals for the Future

Develop methodology for participation by larger classes

Improve bandwidth for remote observing applications in the classroom

Multicast protocols to enable students to "observe the observers" from their dorms or Technology Centers, and from other campuses



Deployment of Fine-grained Grid Applications using I-Light

Seung-Jai Min, Ayon Basumallik,
Rudolf Eigenmann, Y. Charlie Hu

School of ECE, Purdue University

I-Light Applications Workshop, December 4, 2002

Contents

- Motivation
- SPEC OMPM2001 BenchMark
- OpenMP Compiler
- TreadMarks Software DSM
- Performance Result
- Conclusion and Future Work

Motivation

- High communication latency between Grid points over WAN → Current Grid applications limited to coarse-grained
- I-Light opens new opportunities
 - Round-trip latency between Purdue and IUPUI 2.5 msec
- Early experience with deploying SPEC OpenMP benchmarks over I-Light


SPEC OMPM2001 BenchMark

- Fashioned after SPEC CPU2000 benchmarks
- Designed to be used on shared-memory multiprocessors.
- A collection of 10 floating-point and 1 integer applications with medium size data set.
- The application areas include chemistry, mechanical engineering, climate modeling, physics, image processing, and decision optimization.

Translating OpenMP Applications into Software DSM Programs


- Microtasking Model
 - common translation method for loop-parallel languages
- Subroutine-local Variables
 - shared data may include subroutine-local, stack-allocated variables. Stacks on one process are not visible to other S-DSM processes.
 - compiler places these variables into a shared common block
 - requires inter-procedural analysis.

TreadMarks: Software Distributed Shared Memory




Runtime system provides shared memory image

Performance of PI Kernel Benchmark




4 processors at Purdue plus 4 processors at IUPUI, connected through I-Light

Performance of Jacobi Algorithm



Super-linear speedup due to large aggregate memory

Performance of SWIM Application Benchmark



Modest Speedup over 2-processor, local execution

Conclusion and Future Work

- First demonstration of OpenMP Grid applications over WAN (using I-Light)
- Near-ideal speedup on kernels, super-linear speedup on high-locality algorithms, modest speedup on real applications.
- Opportunities for optimizations on I-Light cluster
 - Latency hiding is most important:
 - Prefetching/Forwarding
 - Aggregation
 - Update protocol (as opposed to invalidate protocol)



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasivetechologylabs.iu.edu

I-Light

Steven Wallace

ssw@iu.edu

Advanced Network Management Lab

Indiana University



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasivetechologylabs.iu.edu

Network Application Challenges

- Speed of light
- Packet loss & MTU



Speed of light

- Cost of latency
 - Average time for CPU to fetch one byte of memory: $\sim 12\text{ns}$
 - Average time for CPU to fetch one of byte memory that's ten feet away: $\sim 1,000+ \text{ ns}$
 - Average time for CPU to fetch one byte of memory that's 60 miles away: $\sim 1,000,000\text{ns}$
 - 3,000 miles away: $\sim 50,000,000\text{ns}$
- Worse case, a distributed I-light application might run 100,000 times slower than on a single CPU



Packet Loss / MTU

- Small rates of packet loss cause dramatic performance changes for TCP connections
- Effect of one percent packet loss on TCP performance:

	1,500 MTU	9,000 MTU
10 feet	1Gb/s	1Gb/s
60 miles	160Mb/s	908Mb/s
3,000 miles	3Mb/s	18.1Mb/s



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasive technology labs.iu.edu

Development of a high-performance network application (cheating)

- Why we built Tsunami
- How Tsunami works
- Tsunami's continued development



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasive technology labs.iu.edu

Why Tsunami

“Necessity is the mother of invention”

- Working toward the launch of the Global Terabit Research Network (www.gtrn.net)
- Demonstration and launch at EC network summit meeting in Brussels
- Collaboration between TransPAC, Pacific Wave, Abilene, GEANT & Belnet
- Wanted to do something both symbolic and flashy



pervasivetechlabs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Why Tsunami (cont.)


- Converged on something we called the “*Blue Riband*”
 - Send the text of the first transatlantic telegraph message from Seattle to Brussels over and over
 - Sustain this sending for the same amount of time the first message took (17 hours 40 minutes)
 - Use only common PC equipment
 - Quantify the amount of information transferred



pervasivetechlabs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.org

The Cable





pervasive^{technology}labs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Why Tsunami (cont.)

- The LAB had just finished an extensive evaluation of Linux TCP performance including driver tuning and Gigabit Ethernet NIC evaluation
- We had demonstrated wire-rate gigabit Ethernet transfers in the lab using normal Ethernet MTUs and understood the issues related to bandwidth delay product
- We were confident we could easily achieve $> 500\text{Mb/s}$



Why Tsunami (cont.)

- The Lab shipped a PC to Belgium, one to Seattle, arrange for the machines to be connected to the network and proceeded with testing...
- The test results quickly eroded our confidence
 - Transfer rates varied from a few tens of Mb/s to an occasional but brief few hundreds of Mb/s
- We pretty quickly determined that there was very slight ($< 0.3\%$) random packet loss



Why Tsunami (cont.)

- We worked for days with four separate NOCs to try and determine the source of the packet loss.
- The definitive source was never identified
 - In retrospect I believe there was very low and infrequent loss in two of the three networks
- Less than one week before the scheduled demo we decided our approach just wasn't going to work



Why Tsunami (cont.)

- It was frustrating and **instructive** to see VERY low packet loss cripple our ability to perform a high-speed transfer
- Less than one week before the demo the Lab decided that we were giving up on doing this with TCP and we'd have to make our own protocol
- A few idea were tossed around and pretty shortly we had a “white board” diagram of something we thought might work



pervasivetechlabs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Why Tsunami (cont.)

- Less than three days after the “white board” diagram, we had a working prototype and it seemed to perform as expected.
- A few days later we ran the demo and managed to average over 800Mb/s for 17 hours and 40 minutes.



How Tsunami Works

- Hybrid protocol – uses TCP and UDP
 - UDP for payload
 - TCP for signaling, including requests for retransmission
- UDP checksum used to insure block is transferred correctly
- Sequence numbers applied by the applications to each block
- Receiver keeps track of missing blocks and transmits this list to the sender over the TCP connection



Tsunami's continued development

- Initial version
 - Optimized for memory to memory transfer
 - Went as fast as possible
 - No authentication
 - Not very usable interface
- The “fast as possible part” cause for concern



Tsunami's continued development (cont.)

- Things we've added
 - Enhanced support for files (including large files e.g. > 2GB)
 - MD5-based authentication
 - Simple user interface
 - Ability to define max speed and max loss rate
 - Max speed controls inter-packet delay rather than a window



Tsunami Screen Shots

Tsunami daemon
started

Transfer request
from client

Negotiated port

Calculated inter-
packet delay

Transfer progress

Transfer completion

```
dajr@localhost:~/dist - Shell - Konsole
[dajr@localhost dist]* ./tsunamid
Request for file: 'testfile'
*** IPD Time = 2621 us
Sending to client port 32776
Transmitted 0 of 3052 blocks (0.00% done)
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
Transmitted 1000 of 3052 blocks (32.77% done)
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
Transmitted 2000 of 3052 blocks (65.53% done)
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
ERROR RATE xmit.
error rate = 0 [threshold = 1000]
new IPD = 2621
Transmitted 3000 of 3052 blocks (98.30% done)
Transmission complete.
Transferred 100000000 bytes in 7.72 seconds (103.7 Mbps)
```



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasivetechologylabs.iu.edu

Tsunami Screen Shots

```
dajr@d-103-188:~/dist - Shell - Konsole
[da,jr@d-103-188 dist]* ./tsunami
tsunami> connect 10.0.0.2
Password:
Connected.
tsunami> get testfile
```

Client started

User-initiated connection request and validation

Transfer request



pervasive^{technology}labs
AT INDIANA UNIVERSITY

www.pervasive^{technology}labs.iu.edu

Tsunami Screen Shots

dajr@d-103-188:~/dist - Shell - Konsole

Elapsed time: 00:00:07.075

Last 5 seconds

Blocks count: 400
Data transferred: 0.01 GB
Transfer rate: 103.69 Mbps
Retransmissions: 0 (0.00%)

Continuously
updated transfer
statistics

Cumulative

Blocks count: 2799
Data transferred: 0.09 GB
Transfer rate: 103.70 Mbps
Retransmissions: 0 (0.00%)

Final transfer
statistics

Mbits of data transmitted = 762.94
Duration in seconds = 7.71
THROUGHPUT (Mbps) = 103.71

Return to
command line

tsunami> █




pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Tsunami's continued development

Tsunami development Test Bed





pervasivetechlabs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Test bed Particulars

- Cheap IDE drives (IBM 120GB DeathStar)
- Dell servers (PIII 1Ghz)
- www.3ware.com RAID controllers
 - High-performance (we get over 100MB/s sustained)
 - Cheap (\$500 for controller)



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasive technology labs.iu.edu

Where to get Tsunami

- <http://www.indiana.edu/~anml/anmlresearch.html>



pervasive**technology**labs
AT INDIANA UNIVERSITY

www.pervasivetechlabs.iu.edu

Credits

ANML staff



National Science Foundation
Award No.: ANI-0129592



Eli Lilly Endowment

METACOMPUTING OF FLUID-STRUCTURE INTERACTION PROBLEMS VIA I-LIGHT

Hasan U. Akay, Xiaoyin He and Resat U. Payli

**Computational Fluid Dynamics Laboratory
Department of Mechanical Engineering, IUPUI**

I-Light Applications Workshop, IUPUI, December 4, 2002



Outline


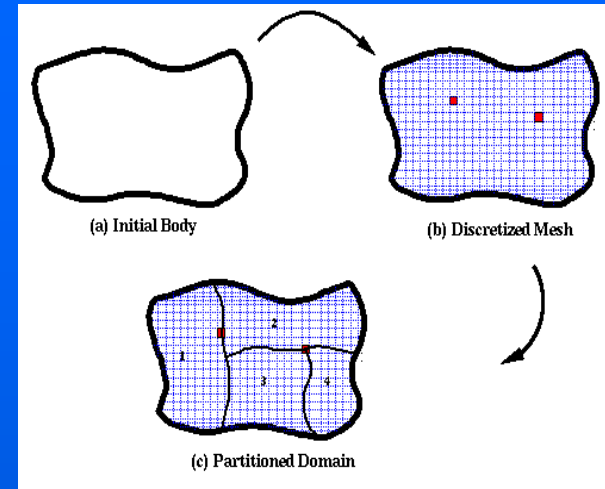
- Computational Fluid Dynamics Laboratory (CFDL)
- Parallel and metacomputing tools developed
- Coupling of multidisciplinary codes
- Example using I-Light

Computational Fluid Dynamics Laboratory (CFDL)

- A multidisciplinary research laboratory in mechanical engineering department, where computational algorithms are developed for solution of large-scale industrial problems, including:
 - Parallel and matacomputing on distributed systems for affordable computing
 - Dynamic load balancing on heterogeneous systems
 - Solution of multi-scale problems encompassing macro, micro and nano-scales
 - Potential applications in advanced propulsion, aerospace vehicles, automotives, micro/nano fluidics, materials processing, and biological phenomena


Parallel Computing

- A computational mesh is generated for the domain
- The mesh is partitioned into *blocks*
- The blocks are distributed to processors on the network and solved concurrently
- Processors communicate data through interfaces between blocks using a message passing library, such as MPI and PVM





Metacomputing with I-Light at CFDL

- Metacomputing is an efficient approach to utilize the resources of geographically distant computers that are connected by a network.
- CFDL uses **I-light**, a high speed optical fiber network connecting IU, IUPUI, and PU and to Internet2, for that purpose.
- I-Light has presently increased the access speed to 30 times than before (1 Gb/s) and is expandable to 100 Gb/s in the future.



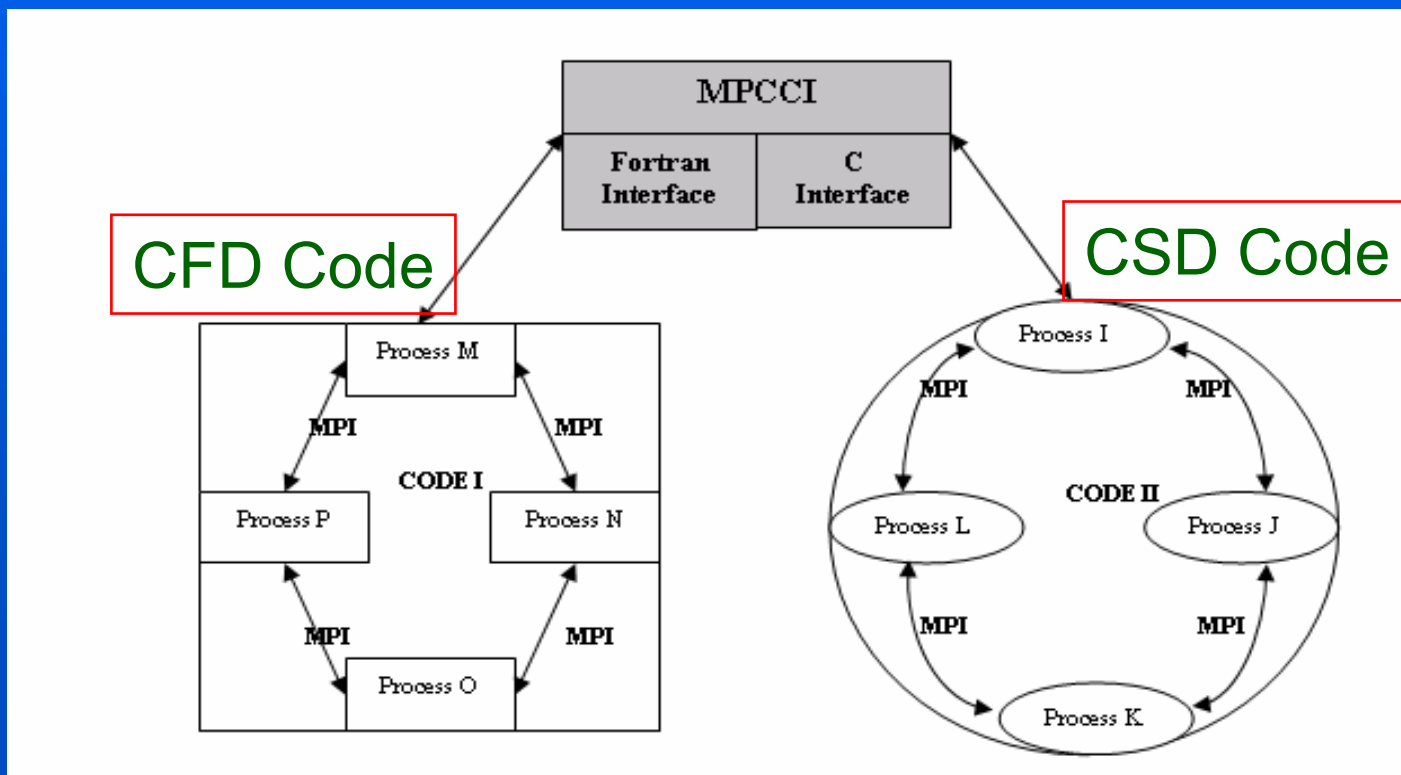
Aeroelastic Coupling Algorithm Developed

- Obtains solution of:
 - Aerodynamic pressures using a CFD (Computational Fluid Dynamics) code
 - Deformation of structures using a CSD (Computational Solid Dynamics) code
 - Movements of the flow mesh using an elastic spring network
- Uses separate computational meshes for flow and structure (CFD and CSD meshes)
- The CFD and CSD meshes are loosely coupled using a code coupling approach across multiple processors
- CFD mesh is subdivided into multiple blocks for parallel and metacomputing on distributed systems





Aeroelastic Coupling Algorithm (cont.)

- A mesh-based parallel code coupling code, MpCCI, is used to exchange information between CFD and CSD codes
- MPI is used for message passing among the parallel blocks of CFD code

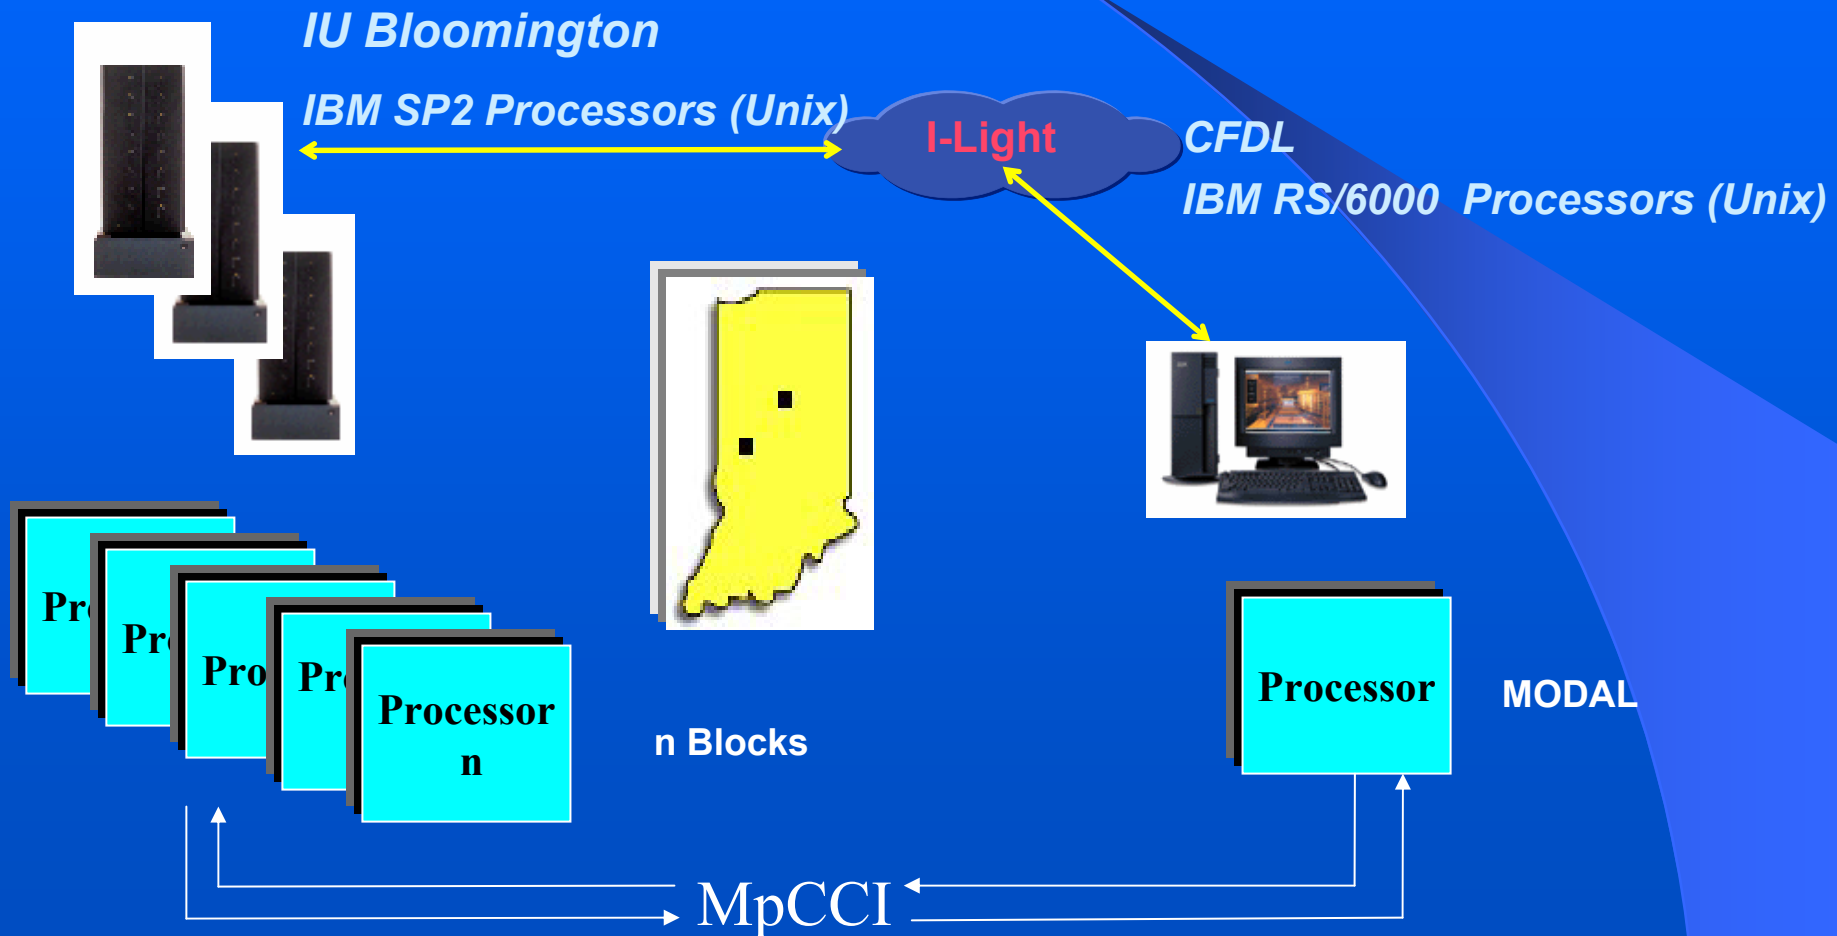


Test Case

- Transient solid-fluid interactions of an aircraft wing is solved using two codes:
 - CFD code: USER3DP
 - CSD code: MODAL
- The CFD domain is solved via partitioned flow meshes for parallel computing
- The codes and their meshes are coupled via MpCCI
- Transient solutions of 500,000 flow and 2,000 structural equations are obtained in a coupled fashion
- Objectives:
 - To demonstrate the feasibility of I-light for metacomputing
 - To test the speed of I-light

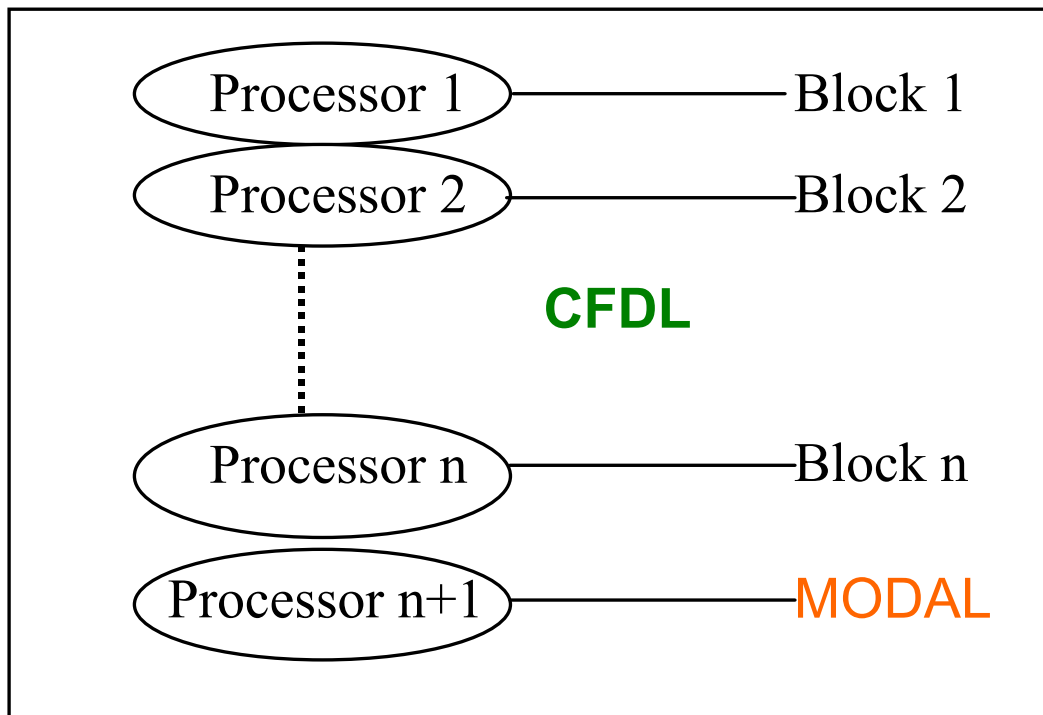


Metacomputing Setup for Test Case



Scenario 1:


Both codes run on CFDL computers



Both USER3DP and MODAL were run on the IBM RS/6000 processors of CFDL. All the processors are connected by the 10/100 Mb/s speed network (Ethernet).

Scenario 2:

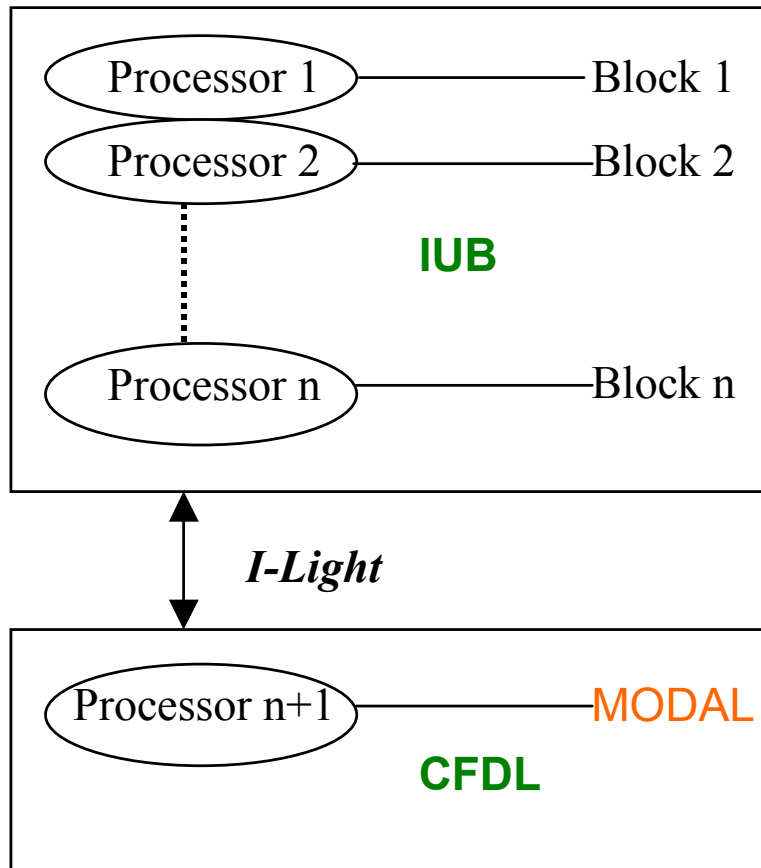
Both codes run on IUB computers



Both USER3DP and MODAL were run on the IBM SP2 system at IUB. All the processors are connected by the high speed network which can provide 150 Mb/s peak bandwidth and 1.2 Gb/second aggregate bandwidth.

Scenario 3:

CFD code runs at IUB and CSD code runs at CFDL




The USER3DP was run on the IBM SP2 parallel system at IUB. The MODAL was run on the IBM RS/6000 computer at CFDL.

I-Light transferred the information between the interface of fluid and solid mesh for two hundred times.


Results of Test Case with I-Light

Total Elapsed Time Comparison




Results of Test Case with I-Light

Communication Time Comparison



Results of Test Case with I-Light

Parallel Speedup: $S_p = T_1/T_p$



Conclusions

- Metacomputing with I-Light employed here has been successful and provided fast solutions.
- I-Light has the potential to make collaborative computing among the three campuses a reality.

Genomes to Grids

Bio Data Distribution for Grid Computing

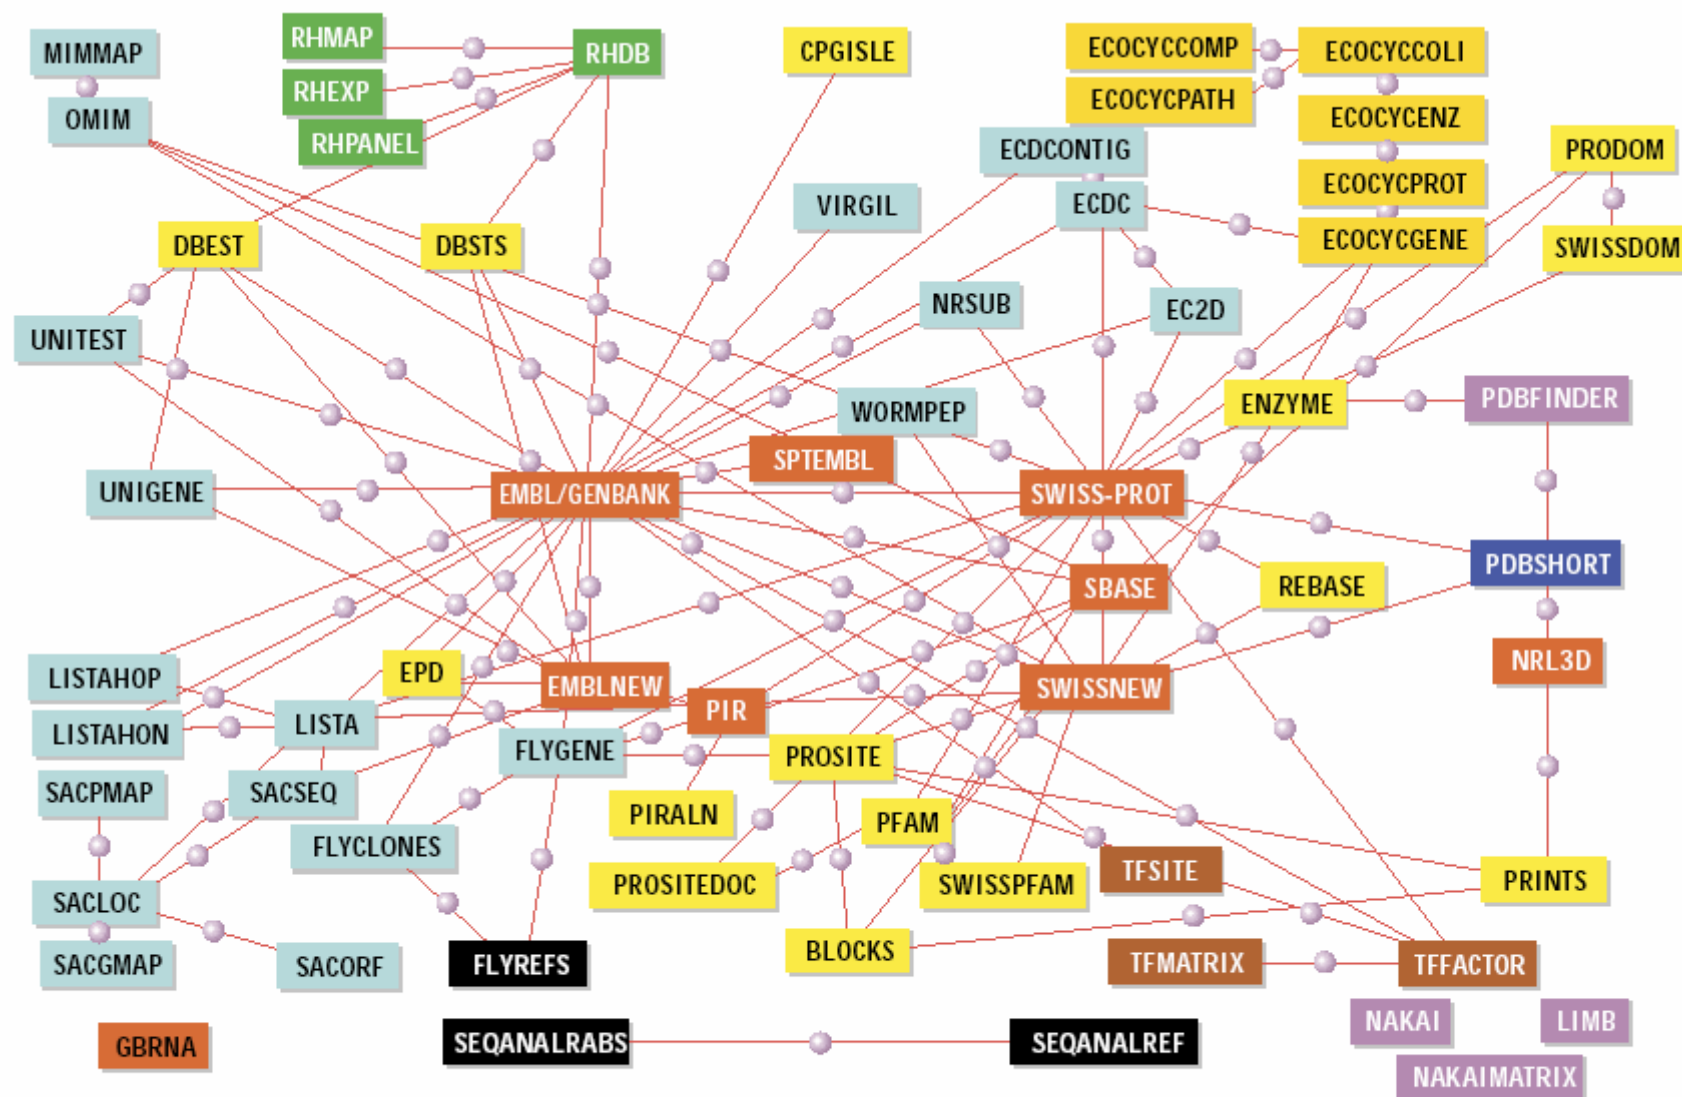
Biologists have discovered many millions of genes and genome features, now part of the bio-data "library" distributed on computers around the world. Grid computing methods for finding and using interesting genome knowledge from this mountain of data are discussed - their promise and practical concerns for building usable bioinformatics grids.

Don Gilbert, gilbertd@bio.indiana.edu, December 2002

Bio Databanks, EBI, Sept. 2002

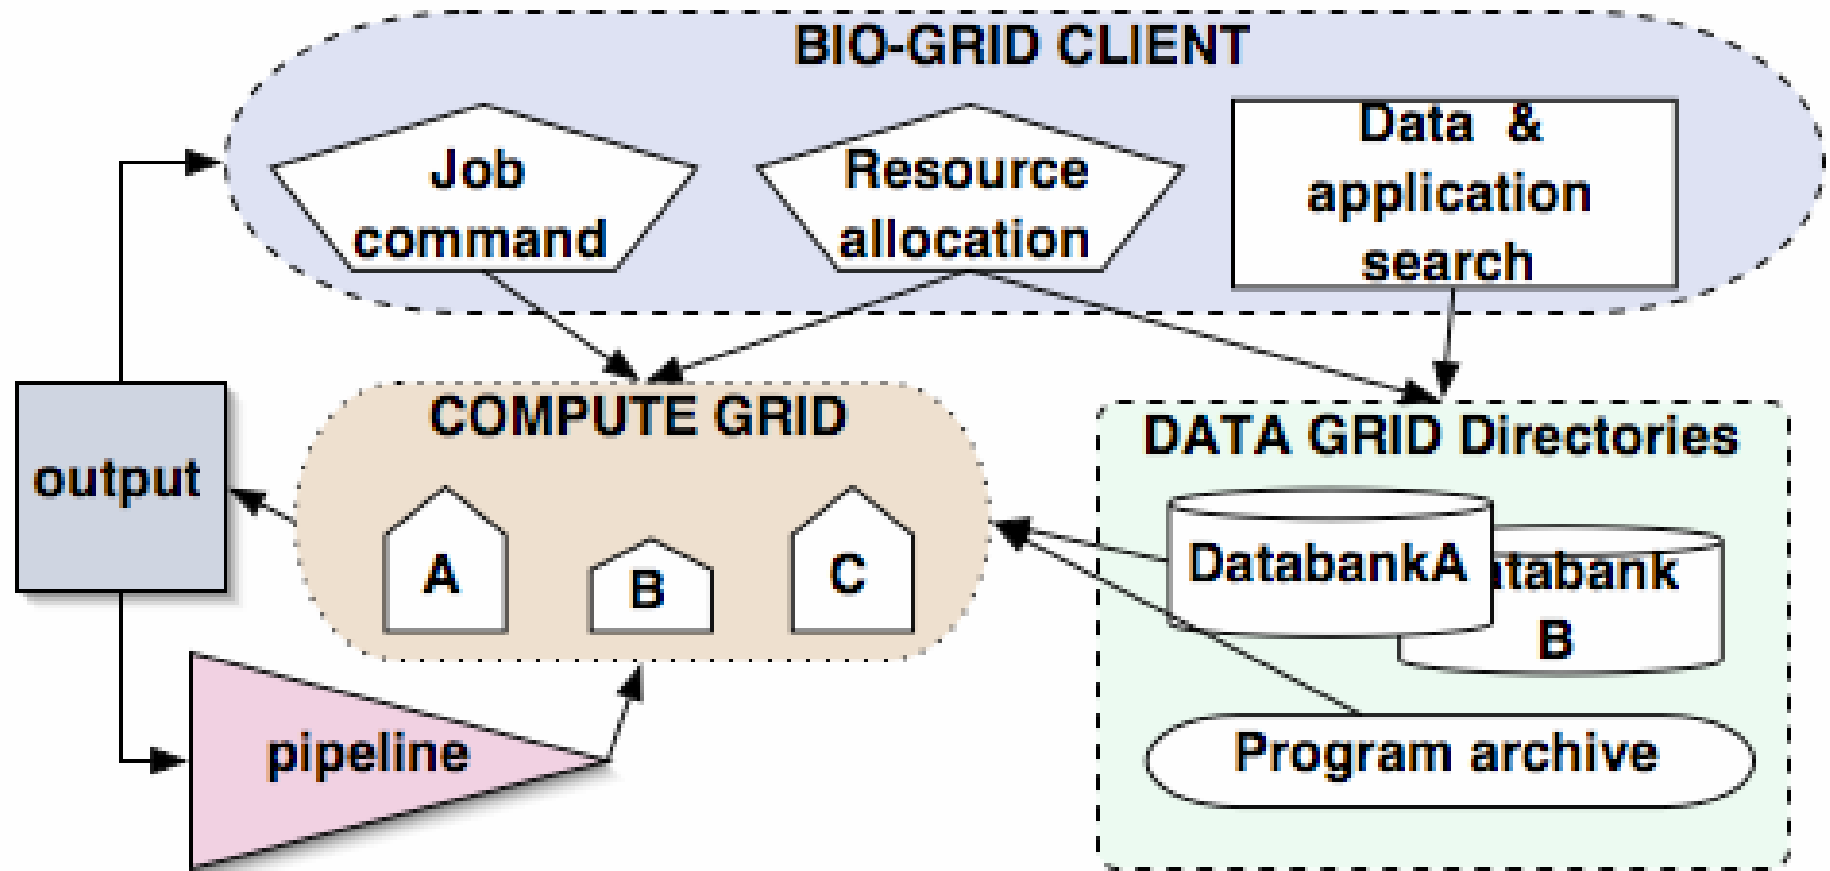
Databank	Contents	Entries
EMBL	DNA Sequences	18,800,000
SWALL	Protein sequences	900,000
InterPro+	Protein motifs	1,000,000
HGBASE	SNP database	1,500,000
	Metabolic Pathways	250,000
MEDLINE	Literature	11,350,000
Total		33,800,000

Constellation of Bio-Data (SRS - Lion Bioscience)



Sequence Protein Structure SeqRelated Genome Metabolic Literature Others Transfac Mapping

BioGrid Schematic



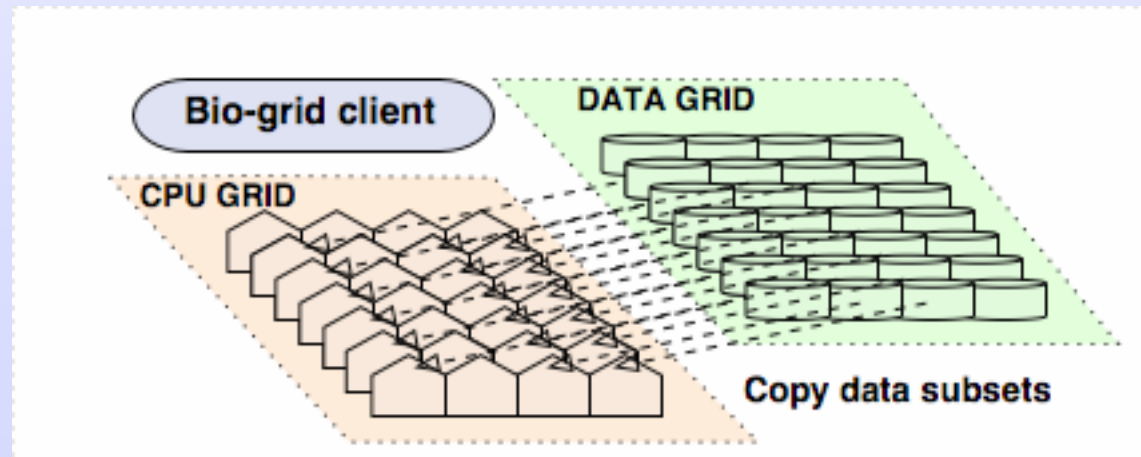
- Grid-aware client software
- Data and software directories
- Grid of processing computers

Directories of Genome Data


- For genome data, *"broad and shallow"* directories federate the *"narrow and deep"* data-bases
- **BioData access tools**
 - *SRS - Sequence Retrieval System; Entrez ; AceDB*
 - *RDBMS ; Ensembl ; IBM DiscoveryLink; BioDAS ; BioMoby*
- **Directory services - Data tools + LDAP , Web Services**
 - **LDAP:** mature, efficient for high volumes, allows federated queries over distributed directories, and works well for SRS databanks and genome annotations
 - **Web Services:** new, simple & complex for XML messages over Web ; has wide industry support , but its many standards are in flux

Moving Data on the Grid

1. @virtualdata= biodirectory "find protein coding sequences for Homo sapiens"
2. @realdata= biodirectory "get locators for @virtualdata split 100 ways"
3. for i (1.. 100) { copydata(realdata[i],gridcpu[i]);
runapp(gridcpu[i]) }



BioDirectory Search/Retreive Efficiency



FTP > 5x LDAP > 5x Web	FTP	LDAP-net	Getz-local	WGetz-net	SOAP-net
Q2: 340K ids, 1.5 Gb	162	506	3138	3653	4680
Q1: 20K ids, 60 Mb	8	84	297	309	300

Q1/Q2 - Query biosequence directories ; **FTP** - no query selection

Q1 = {swissprot trembl refseq}-des:kinase , 20K records; Q2 = genbank-org:drosophila , 340K records

gilbertd@bio.indiana.edu, Oct 2002

Using Bio Directories

Simple client
software

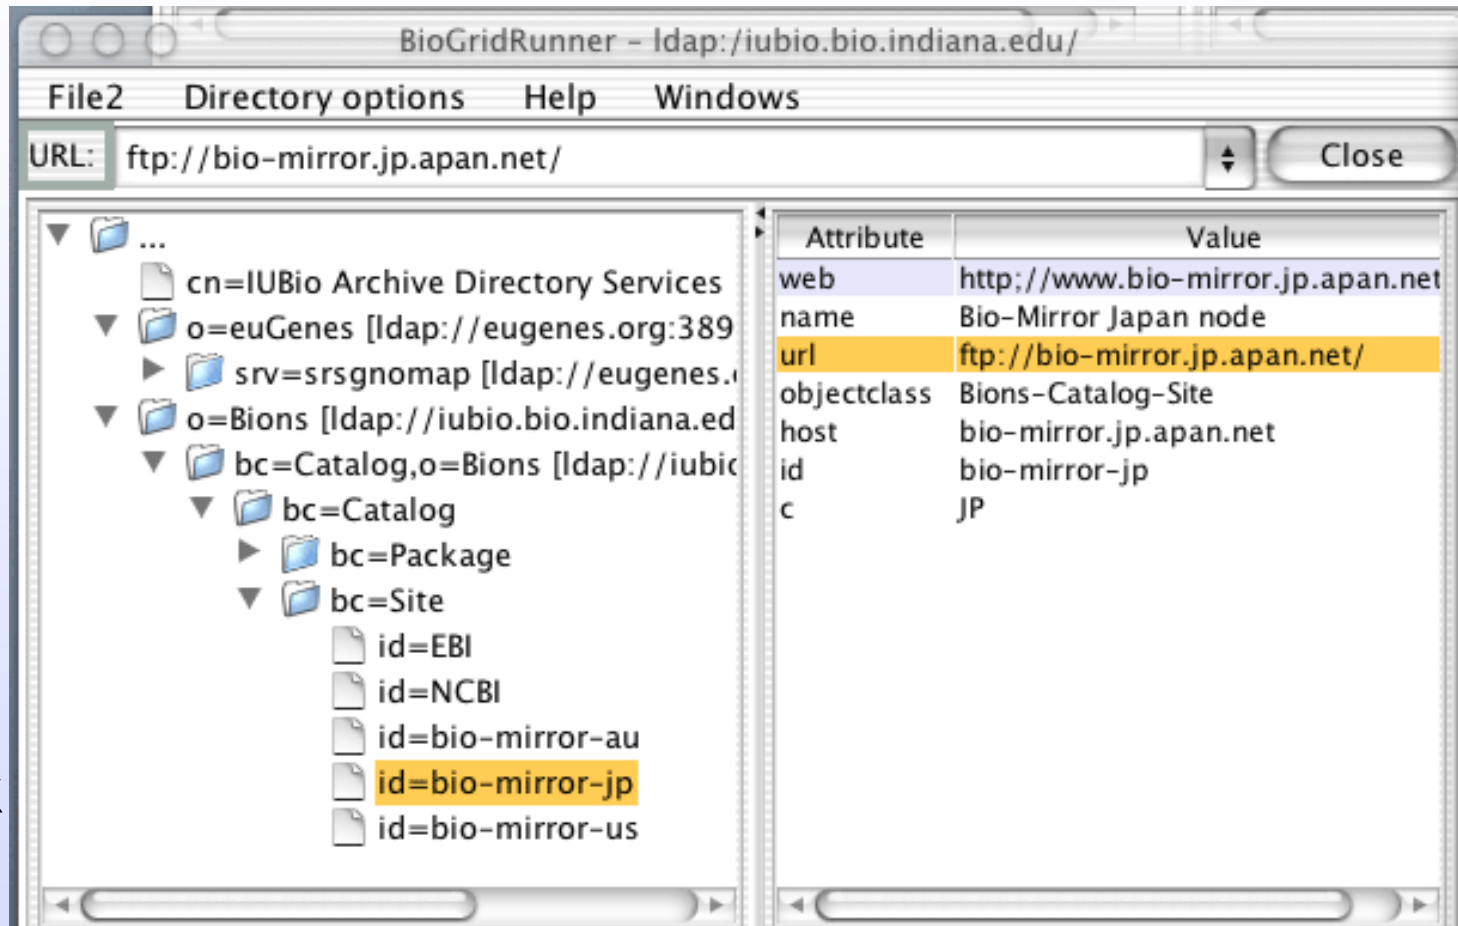
Automated
use

People use

Discovery

Search by
many
criteria

Retrieve bulk
subsets



Wrap up

- **Future of Bio-data on Grids**
 - High volume, complex, changing, distributed data
 - Computationally find and use this data
- **Best methods for Bio-data to Grids**
 - Efficient selection and transport to grid computers
 - LDAP works well ; Web-XML is usable ; Others?
- **Community needs and uses**
 - Common data descriptions, schema, ontologies
 - Simple, practical, flexible grid methods ; use existing dbs

See <http://iubio.bio.indiana.edu/biogrid/>

Towards an Infrastructure for Large-Scale Information Analysis, Visualization, Information Retrieval Research & Education



Katy Börner & Javed Mostafa
School of Library and Information Science
School of Informatics
Indiana University



10th Street & Jordan Avenue, Bloomington, IN 47405
katy@indiana.edu & jm@indiana.edu

I-Light Conference, December 4th, 2002

Database

- Oracle database of about 15 million records covering books, journals, proceedings, doctoral and master's theses, technical reports, patents, grants, and protein and gene data from cross-disciplinary research and specific knowledge domains.
- This database will be unique in its size and coverage. Most of the documents will be available in full text. Software that facilitates a continuous, automatic update of the database will be in place.

Note: ACM portal, CiteSeer, and PubMed currently provide access to about 361,400, 507,800, and 11 million records respectively.

Core Information

Primary Information - Cross-Disciplinary DLs:

- ISI's Journal Citation Reports
- Proceedings of the National Academy of Sciences
- Awarded NSF and NIH Grants
- Funding Opportunities published by the Community of Science

Core Information

Primary Information - Domain Specific DLs:

- Computer Science ACM Portal
- Networked Computer Science Technical Reference Library
- NEC's CiteSeer
- E-Print Archive
- Online Mendelian Inheritance in Man

Core Information

Secondary Information – Derivative Data:

Various types of pre-computed data will also be maintained in indexed form for fast access and manipulation:

- Term dictionaries extracted from corpora
- Term frequencies
- Term weights (calculated using different methods)
- Term vector distances
- Term clusters

Open Source Software Repository

Will provide access to:

- **Utility programs** (filtering, stemming, stop-word elimination, unique term extraction, automatic citation indexing, etc.),
- **Data analysis and dimensionality reduction** (e.g., Vector Space Model, Multidimensional Scaling, Pathfinder Networks, Latent Semantic Analysis, Clustering algorithms, etc.)
- **Visualization/interaction algorithms** (GRIDL GRaphical Interface for Digital Libraries, Treemap, Force Directed Placement, and Hyperbolic Tree, Self-Organizing Map, Fisheye Views, the Jazz Zooming Graphics Toolkit, etc.)

Web Services

The repository will implement various atomic web services, for example:

- Search <term> <corpus>
- Frequency <term> <corpus> ... <corpus>
- Weight <method><term> <corpus>

That will be used for fast manipulation of large corpora content and secondary data by other software (remote or local).

Software Interchange Format

- Will be based on metadata standards and ensure that algorithms can be combined in multiple ways (e.g., using different data mining algorithms with diverse information visualization algorithms).
- All Java-based algorithms can be run in stand-alone mode as an applet or application.
- A standardized software framework will interlink the network of different databases and services by a common communication protocol.

Infrastructure

- All services will run on IU's Sun E10000 Research System (Solar), a shared memory, multiprocessor system with 64 400MHz CPUs and 64GB memory.
- The infrastructure will directly support the research of the Information Processing Laboratory at IUB, the School of Informatics at IUB.

Application: Knowledge Domain Visualization

Domain Visualizations Can Help Answer Questions such as:

- What are the major research areas, experts, institutions, regions, nations, grants, publications, journals in xx research?
- Which areas are most insular?
- What are the main connections for each area?
- What is the relative speed of areas?
- Which areas are the most dynamic/static?
- What new research areas are evolving?
- Impact of xx research on other fields?
- How does funding influence the number and quality of publications?

Indicator-assisted evaluation and funding of research: Visualizing the influence of grants on the number and citation counts of research papers.
(Boyack & Börner, 2003)

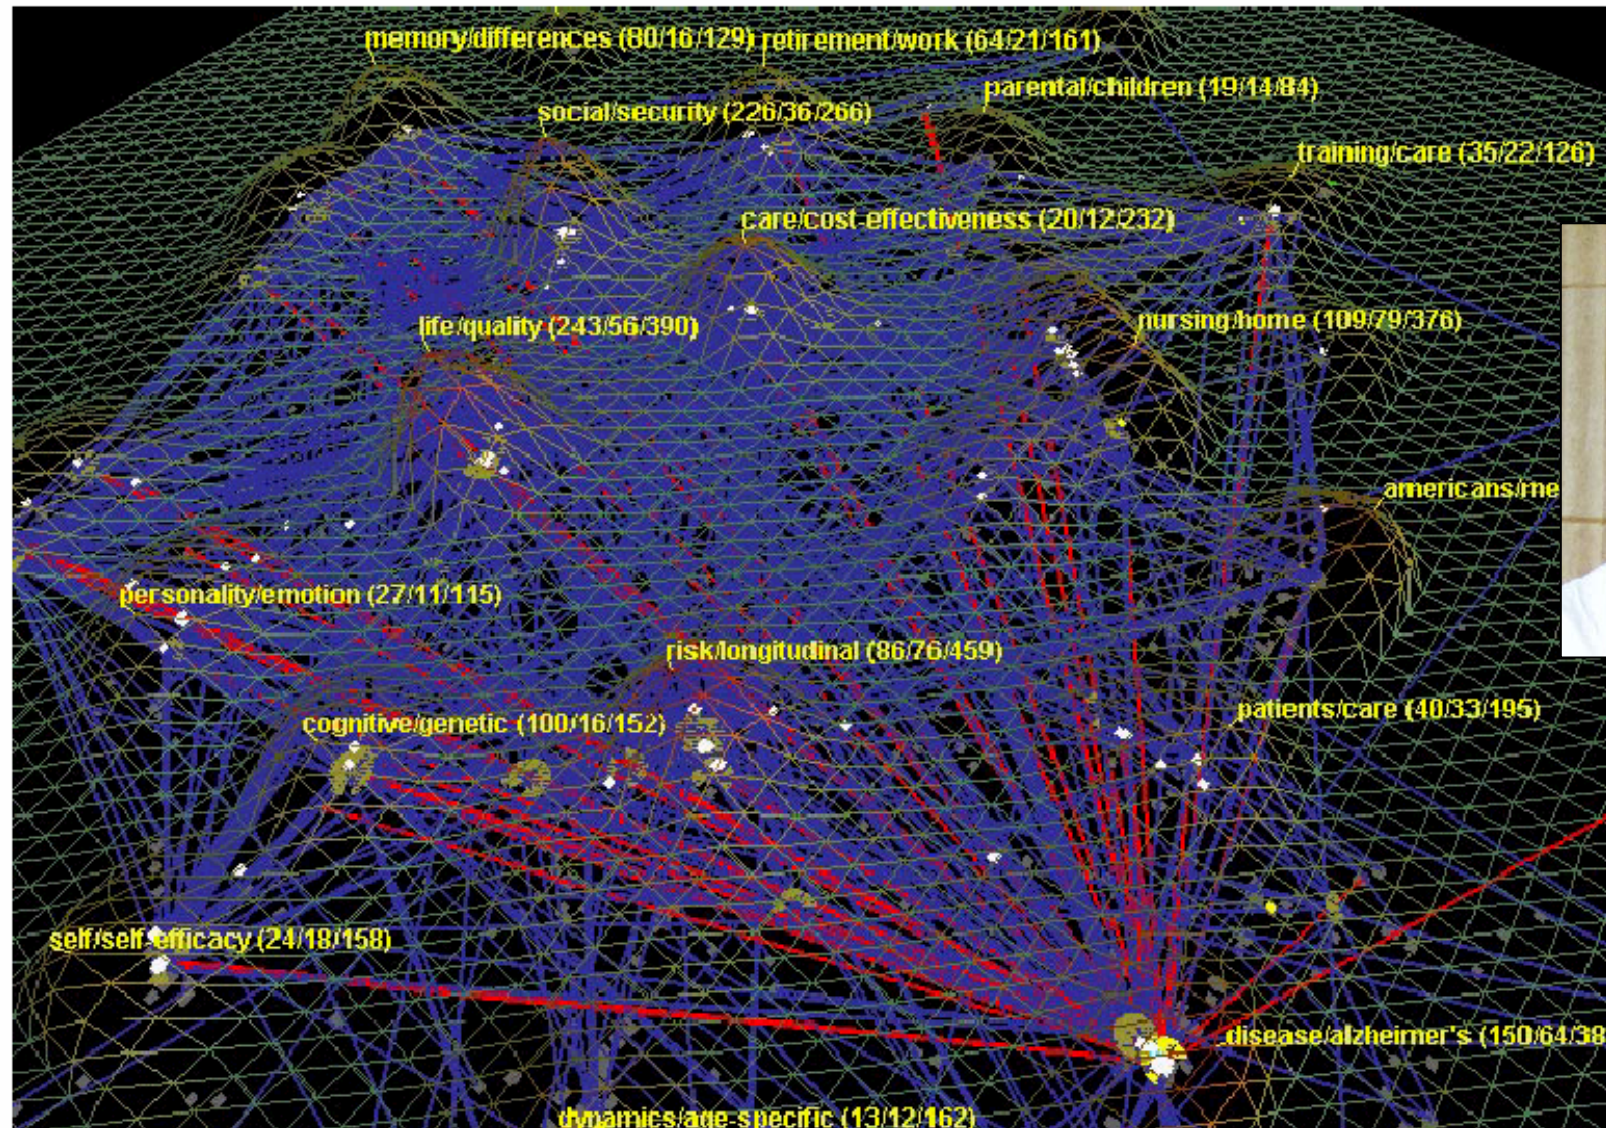



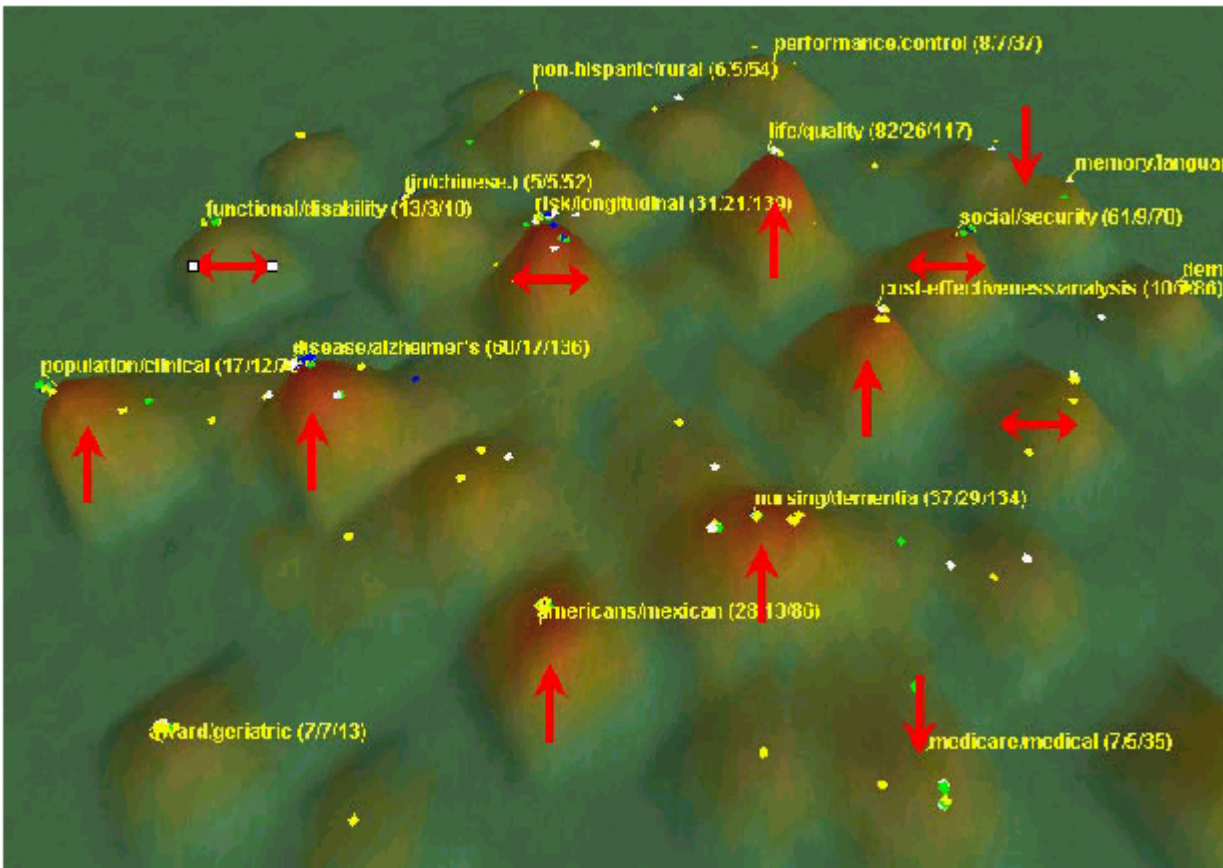
Figure 7. Author supplied linkage patterns from grants to publications with links highlighted in red for grant 01 P50 AG11715-01.

Figure 4. BSR map in two different time periods showing the impact of funding on the number of articles (size of peaks). Small grants are shown by light colored markers on the terrain. Large grants are shown by dark colored markers. Significant changes in the number of publications over time by peak are shown by red arrows.

(Boyack & Börner, 2003)



1999-2000



Application: Information Retrieval

In addition, an NSF funded project on

- Digital Libraries (see <http://sifter.indiana.edu>)



and an NSF ITR project on

- Bioninformatics (see the BioSIFTER site <http://xtasy.slis.indiana.edu/biosifter/>)

that involves computer scientists, biologists, and information scientists from both IUPUI and IUB campuses will utilize this infrastructure.

Unique Features

- Uniform, modular, open architecture
- Scalability to handle GB size data sets
- Parallel computing infrastructure
- Usage of XML-based, OAI derived communication protocols for easy integration of new databases and services as well as the serialization of software packages
- Detailed documentation of data and code but also links to related publications
- Online GUI supporting the request and navigation of diverse information processing jobs for teaching and research purposes.

Related Publications

References

Baumgartner, J. and Börner, K., Towards an XML Toolkit for a Software Repository Supporting Information Visualization Education. in *IEEE Information Visualization Conference*, (Boston, MA, 2002).

Börner, K. and Zhou, Y., A Software Repository for Education and Research in Information Visualization. in *Information Visualisation Conference*, (London, England, 2001), IEEE Press, 257-262.

Mostafa, J., & Lam, W. Automatic Classification Using Supervised Learning in a Medical Document Filtering Application. *Information Processing & Management*, 36(3), 415-444, 2000.

Zhang, J., Mostafa, J., & Tripathy, H. Information Retrieval by Semantic Analysis and Visualization of the Concept Space of D-Lib® Magazine, 8(10), 2002.

<http://www.dlib.org/dlib/october02/10contents.html>

Additional Resources

Web Resources

InfoVis Lab: <http://ella.slis.indiana.edu/~katy/ivlab/>



Software Repository:
<http://ella.slis.indiana.edu/~katy/L697/code/>



LAIR Lab: <http://lair.indiana.edu>

High Performance Computing Resources at Indiana University

Mary Papakhian

Manager, Research and Technical Services

Dave Hart

Manager, High Performance Computing Support

University Information Technology Services

Sun E10000

- Acquired 4/00
- Shared memory architecture
- ~52 GFLOPS
- 64 400MHz cpus, 64GB memory
- > 1TB external disk
- Solaris 2.8



Sun E10000


- Fortran, C and C++ Compilers with MPI, OpenMP, TotalView and Prism debuggers
- SAS, Splus, SPSS, Matlab, Mathematica
- Bioinformatics software: GCG, SeqWeb, NCBI Toolkit, BioPerl, QTL Cart and Seals, Wublast
- Oracle databases
- Mix of interactive, serial and parallel jobs

IBM Research SP

- Acquired 9/96
- Expanded in 1998, 1999, 2000,2001,2002 with IBM SUR grants and Lilly Endowment
- Geographically distributed at IUB and IUPUI
- Ranked 112th in world Top500 supercomputer list
- 632 cpus, ~1 TeraFLOPS
- Distributed memory system with shared memory nodes
- AIX 5.1




Aries Complex



4 cpus
2-4 GB 4 cpus
3 GB
Power3+ thin Power3+ wide

Orion Complex



16 cpus
8 GB 4 cpus
1 GB 16 cpus
32 GB
Power3+ high PowerPC Power4

IBM Research SP

- Wealth of software including SAS, SPSS, S-Plus, Mathematica, Matlab, Maple, Gaussian, GIS, scientific/numerical libraries, Oracle and DB2, and more
- ~800 users from ~90 departments on IUB, IUPUI, Fort Wayne, Purdue campuses
- ~5TB disk
- Mix of interactive, serial, parallel jobs

AVIDD Linux Clusters

- Analysis and Visualization of Instrument-Driven Data
- Preemptive scheduling for realtime jobs
- 2 200-cpu IA32 clusters at IUB and IUPUI
- ~5TB storage via GPFS
- 36-cpu IA64 cluster with ~1TB storage at IUB
- 16-cpu cluster at IUN
- RedHat Linux 7.3

Research and Technical Services

For more information, and to apply for accounts

<http://www.indiana.edu/~rats>

Send questions to rats@indiana.edu

High Performance Computing Support

- The HPC support group provides programming support for IU researchers, for example answering questions about using software libraries or using more than one processor at a time, to gain the advantages such systems can offer.
- Our group consists of highly qualified professional computational scientists, who seek collaborative relationships with IU research programs.
- Send questions to hpc@iu.edu.

Acknowledgements

- This material is based in part upon work supported by the National Science Foundation under Grants No. 0116050 and CDA-9601632. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).
- This research was supported in part by the Indiana Genomics Initiative (INGEN). The Indiana Genomics Initiative (INGEN) of Indiana University is supported in part by Lilly Endowment Inc.
- This work was supported in part by Shared University Research grants from IBM, Inc. to Indiana University.
- This work has benefited from IU's relationship with Sun as a Sun Center of Excellence.

Research Data Storage Services at IU

Anurag Shankar

Manager, Distributed Storage Services
University Information Technology Services
Indiana University

<http://storage.iu.edu/presentations/ilight.ppt>

Acknowledgements: IU's massive data storage effort received funding in part from NSF grant no. 0116050, a Shared University Research (SUR) grant from IBM, Inc., the Indiana Genomics Initiative (supported in part by Lilly Endowment, Inc.

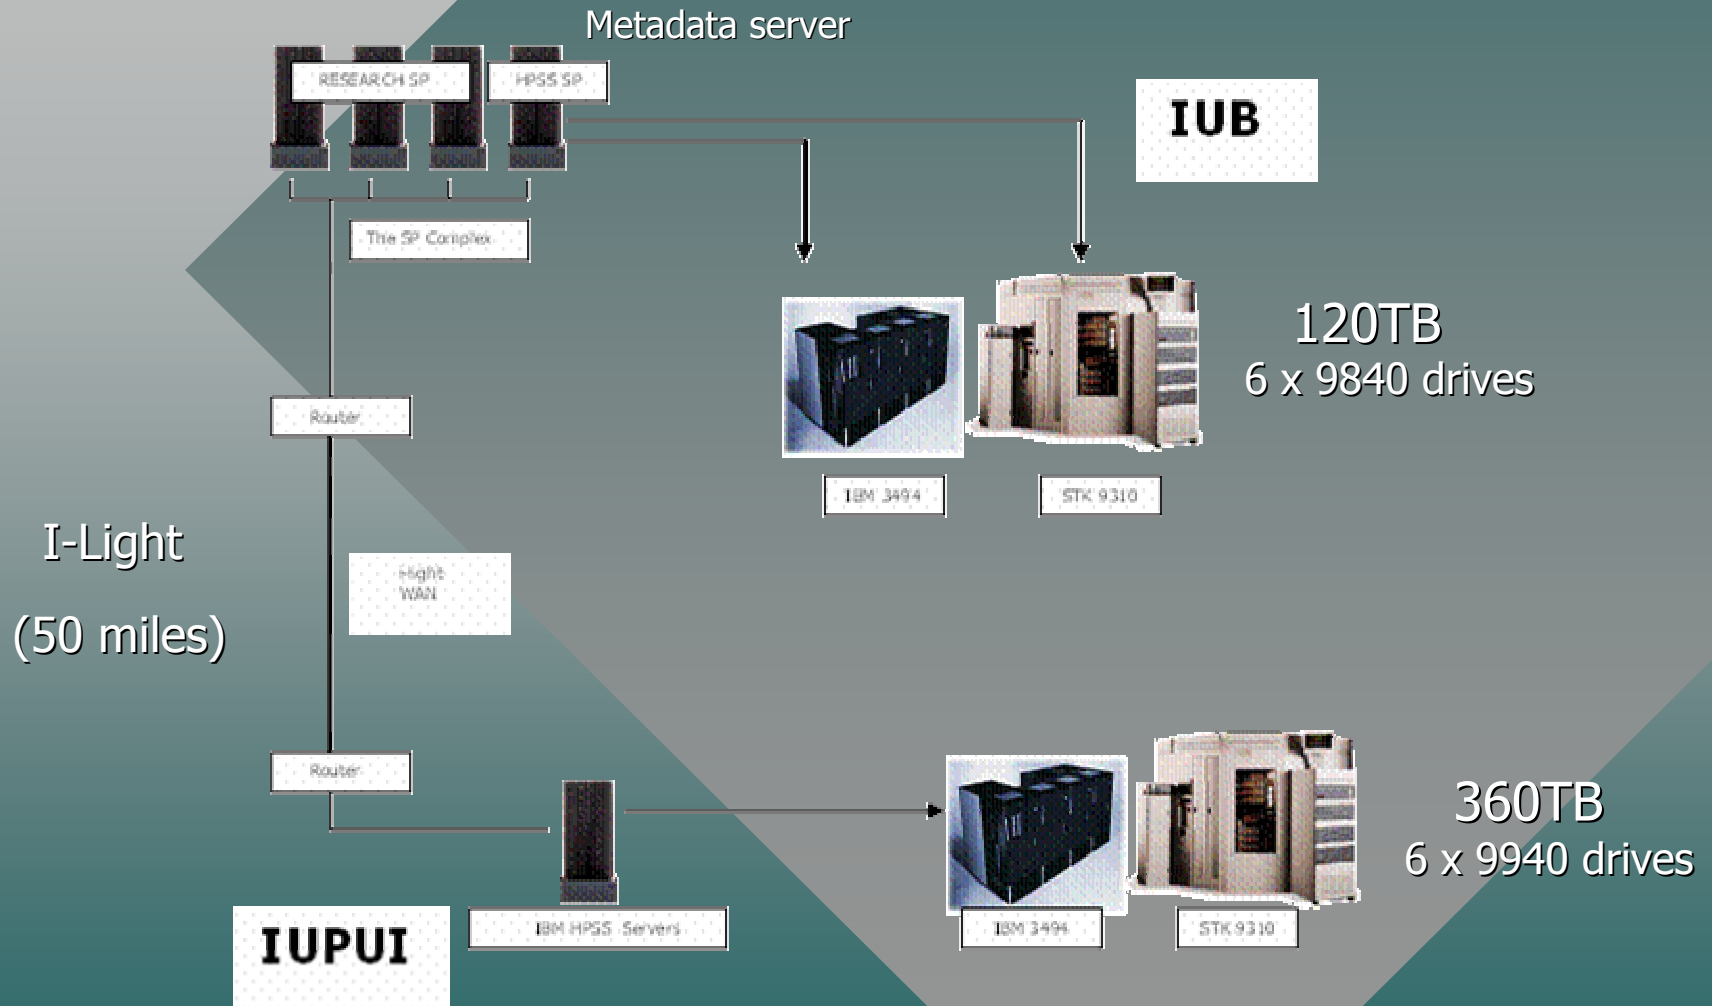
Research Data Storage @ IU

- IU is **unique** among academic institutions in making a massive data storage infrastructure available to **any user** on the **eight IU campuses**.
- IU's **Massive Data Storage System (MDSS)** service uses the High Performance Storage System (HPSS) software and IBM hardware to provide **500GB** of **free, near-line** storage.
- The **MDSS** utilizes **I-Light** to provide **disaster tolerance** (by making two tape copies automatically, one at IUB and the other at IUPUI).
- **General disk storage** (up to a few GB) provided via the **Common File System (CFS)** service.

The MDSS and CFS

- The MDSS provides massive storage capacity by using a modest amount of disk cache (2TB) in front of massive tape libraries (500TB). (Tapes are more cost effective to store data.)
- Unused data are migrated from disk to tape seamlessly. Later retrievals cost up to two minutes while the tape containing data is located/mounted.
- The MDSS is ideal for long-term, archival storage.
- The CFS provides easy, web accessible, disk-based storage from a desktop anywhere.

Distributed MDSS between Bloomington and Indianapolis



For More Information

- Visit the Distributed Data Storage @ IU web site: <http://storage.iu.edu/>.
- Search for the keywords “mdss” or “cfs” in the IU **Knowledge Base** <http://kb.iu.edu/>.
- Email store-admin@iu.edu.

Acknowledgments

- This material is based in part upon work supported by the National Science Foundation under Grants No. 0116050 and CDA-9601632. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).
- This research was supported in part by the Indiana Genomics Initiative (INGEN). The Indiana Genomics Initiative (INGEN) of Indiana University is supported in part by Lilly Endowment Inc.
- This work was supported in part by Shared University Research grants from IBM, Inc. to Indiana University.

Resources Accessible Through I-Light and How You Can Use Them to Enhance Your Research

Rick Westerman - Genomics Initiative
Dave Seaman, Dale Talcott, Bill Whitson - ITaP

I-Light Applications Workshop
December 4, 2002

“As we build a more global infrastructure, researchers will be able to choose from resources around the world to increase their throughput.”

John Shalf, LBL, leader of the winning team in the SC2002 HPC Bandwidth Challenge

- HPC systems
 - IU, IUPUI, Purdue
 - Alliance, NPACI, NCAR, etc.
- Data
 - Biodata
- Grid
 - Grid Testbed using recycled machines

- Providing student access to local and remote resources for rendering
- Using network infrastructure and clusters of recycled PCs
- Use for parallel programming and access to high-performance applications planned

Recycled Clusters

- Approx. 400 machines (450 MHz PII) in use
- 200 more (550 MHz PIII) soon
- 256 MB memory, 9 GB disk, RedHat 7.3
- 100 Mbit ethernet
- 48-port switches




Grid Testbed

- Use recycled cluster to test and develop Grid infrastructure



- Will explore Globus Toolkit and Cactus initially

- Remote Visualization
- Access Grid and Videoconferencing
- IP Telephony




The biodata disk

- Contains data files for BLAST, HMMer, etc.
- Updated weekly
- Accessible via NFS and SAMBA
- Databases include:
 - nt (all of Genbank, EMBL, DBJ)
 - mouse_est
 - rice_est
 - arabidopsis genome
 - nr (translated Genbank, et.al.)
 - chromosome (all sequenced ones)
 - ... and more

Pentagon Crash Simulation

- Simulations done using Purdue, IUPUI systems
- Involved interactive work plus movement of large amounts of data between sites



- High bandwidth and low latency needed for HPC applications
- Latency to IU and IUPUI
 - Under 5 msec round-trip.
 - » This means it is perfectly reasonable to share filesystems via NFS.
 - » When combining the IBM SPs at the three sites, the systems need to exchange status and update information. Over I-Light, this takes place as promptly as if the systems were in the same building.
 - This is still not good enough to support communication-intensive applications. Those need latencies below 40 usec (100 times as quick). But 5 msec is easily fast enough for loosely-coupled computation.
- Latency to Internet2 sites
 - Even San Diego Supercomputer Center is only 75 msec away.
 - Use of vi, emacs, or other screen-based application is fine. So, files can be edited remotely, rather than needing to copy them back and forth for editing.
 - X-Windows application response is tolerable, making remote visualization practical.

- 30+ events since Purdue's AG node went on-line 2/2002
 - Multi-day conferences and workshops
 - 1-3 hour seminars
 - Multi-day training sessions
 - Collaborative work sessions and meetings w/ other universities, NSF, Motorola, ...



“For college students, the Internet just works. It’s like turning on the tap and getting water or turning on the TV.”

Steve Jones, chair of the communications department at the University of Illinois, Chicago

I-Light removes barriers and enables researchers to think the same way about accessing remote resources...

For more information about these
projects, please see

<http://www.itap.purdue.edu/>

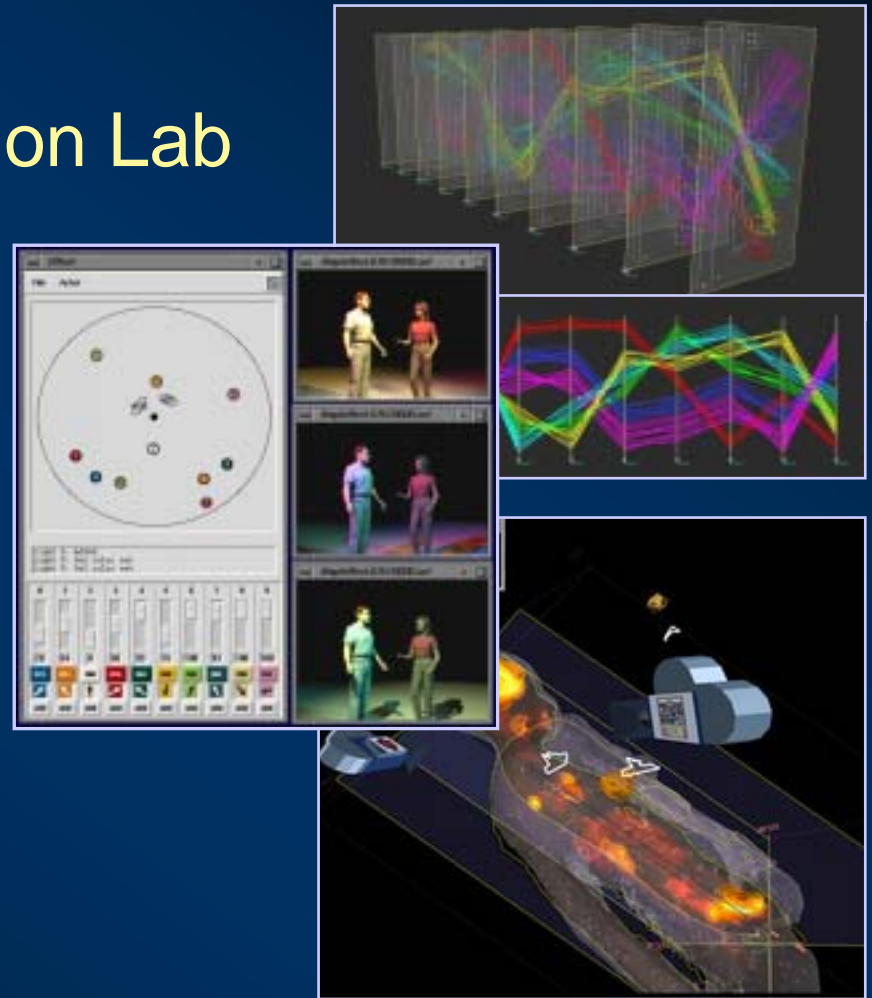
or contact

racs-info@purdue.edu

Network-based Visualization & Visual Collaboration Methods

UITS Advanced Visualization Lab

- Technology Areas:
 - Visualization
 - Virtual Reality
 - Visual Collaboration
- Support/Application Areas:
 - Research
 - Education
 - Creative Activities




Objectives

- Overview of applications and benefits of networking to...
 - Visualization (human <-> remote computer interaction)
 - Collaboration (human <-> remote human interaction)
- Many diverse methods
 - Visualization is a highly interactive (vs. batch) process
 - Must be customized to the goals of the users and application
 - Need flexibility to integrate user and visualization technologies with other IT technologies
- Nearly everything has been done at least once
 - Provide a point to jump-start your project



Network-based Visualization

- Human – remote computer interaction
- In our age of powerful, inexpensive desktop systems, we must identify significant advantages of central systems over local systems:
 - Data size – very large working memory or storage requirements
 - Data management – improved organization, integration, and/or security
 - Extensive computation - for visualization rendering and/or co-processing
 - Specialized rendering hardware – faster or better quality rendering
 - Specialized software – limited licensing or platform availability



Network-based Visualization

How are graphics sent across the network?

- Example: Molecular Visualization

- Data: (e.g. data server)




- `atom 1 (type, location), atom 2,...`
 - requires local rendering

- Graphics Commands (e.g. remote X, GLX)

- `color(red), sphere(radius 0.57), ...`
 - rendering/interaction load shared between systems

- Rendered Images (visual serving, VNC, etc.)

- `image 1, image 2, ...,`
 - rendering load on server, good for lightweight clients

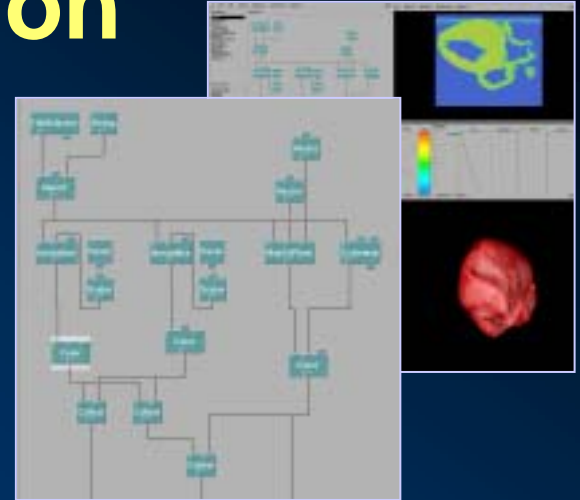


- Must compare data complexity, visual quality, image resolution, and bandwidth availability for your app

Network-based Visualization


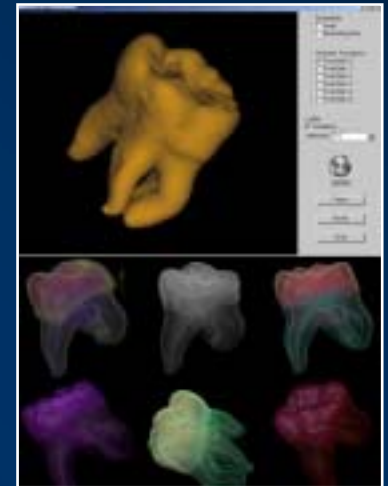
Issues

- Networking: latency, sustained throughput, last mile problem
- Resource allocation: scheduling and balancing
- Compression: methods and loss tolerance
- Development requirements: existing apps/data vs. custom apps



Technologies & Examples


- Modular vis environments: OpenDX, Iris Explorer
- Server-based rendering: VolumePro, Ray-tracing cluster
- Visual streaming & Application sharing: VNC, VizServer
- Remote data access: ODBC, Web protocols, etc.



Network-based Collaboration


Human-human interfacing

- a.k.a. Tele-collaboration
- Enables sharing of information, knowledge, and insight
- May be expert-to-novice or peer-to-peer



Issues:

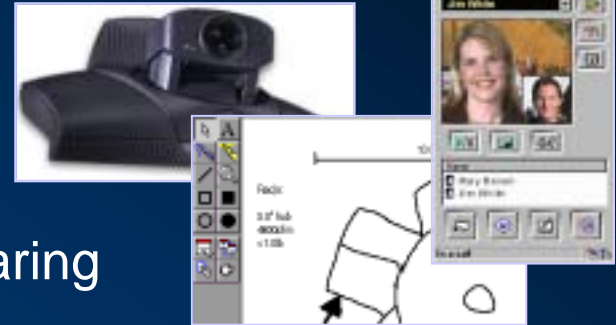
- Data delivery requirements: loss tolerance, latency, etc.
- Topology: point-to-point vs. multicast
- Synchronicity: synchronous vs. asynchronous
- Control: flow-of-control, restrictions, etc.
- Development: existing app and data or custom app?



Network-based Collaboration


What's sent?

- Audio and video
- Documents, whiteboards
- Images: for image-based application sharing
- Data: may be static or changing
- Control data: viewpoints, avatar data, vis parameters



Technologies

- Web-based file-sharing & data repositories
- Web-based 3D collaboration: ActiveWorlds, etc.
- Teleconferencing: Polycom, Access Grid, NetMeeting
- Application Sharing: VNC, NetMeeting
- Simple collaborative applications: desktop-to-desktop viewpoint & parameter sharing
- Tele-immersion: multi-point VR interaction



Summary

Questions:

- Does this stuff really work?
 - Yes. But it's not always “out-of-the-box” ready.
- Are people using it?
 - Yes. But it depends on motivation (and pain) of users.
 - Broader availability of tools and systems will help.
- How do I get started?
 - Talk to AVL staff (avl@avl.iu.edu, www.avl.iu.edu)
 - Start simple and expand as needed.



Acknowledgements:

- AVL staff, clients, & OVPIT
- This material is based in part upon work supported by the National Science Foundation under Grants No. 0116050 and CDA-9601632. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).
- This research was supported in part by the Indiana Genomics Initiative (INGEN). The Indiana Genomics Initiative (INGEN) of Indiana University is supported in part by Lilly Endowment Inc.

