T·H·E
OHIO
STATE
UNIVERSITY

UNIVERSITY MEDICAL CENTER

Cooperative Biomedical Research, Data Virtualization and Grid Computing

Joel Saltz
Chair, Biomedical Informatics
Professor Computer Information Science
The Ohio State University
Tony Pan
Research Scientist
The Ohio State University

Imaging, Medical Analysis and Grid Environments (IMAGE)

September 16 - 18 2003

Organisers: Malcolm Atkinson, Richard Ansorge, Richard Baldock, Dave Berry, Mike Brady, Vincent Breton, Frederica Darema, Mark Ellisman, Cecile Germain-Renaud, Derek Hill, Robert Hollebeek, Chris Johnson, Michael Knopp, Alan Rector, Joel Saltz, Chris Taylor, Bonnie Webber

Goals – what the world will look like? (IMAGE: Data Management Working Group)

- Identify, query, retrieve, carry out on-demand data product generation directed at collections of data from multiple sites/groups on a given topic
- reproduce each group's data analysis and carry out new analyses on all datasets
- Should be able to carry out entirely new analyses or to incrementally modify other scientist's data analyses
- Should not have to worry about physical location of data or processing
- Should have excellent tools available to examine data
- This should include a mechanism to authenticate potential users, control access to data and log identity of those accessing data.

Image/Data Application Areas

Processing Remotely-Sensed Data

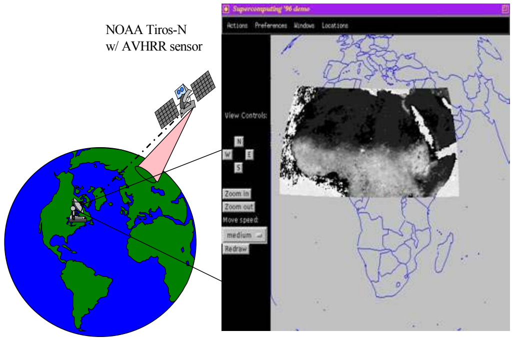

Satellite Data Processing

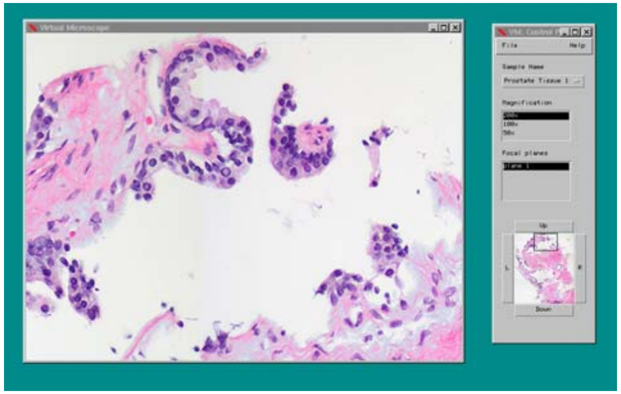

Digital Pathology

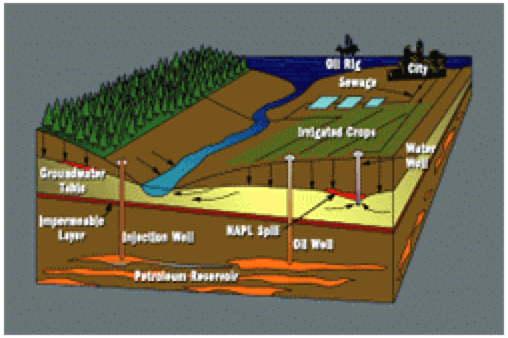

Managing Oilfields,
Contaminant Transport

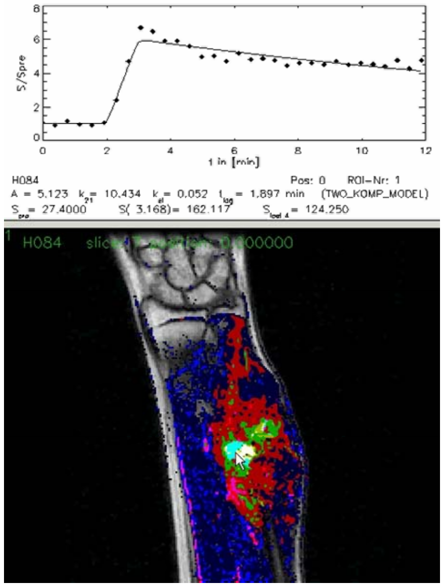

DCE-MRI Analysis

Imaging, Medical Analysis and Grid Environments (IMAGE)
September 16 - 18 2003
e-Science Institute,
15 South College Street, Edinburgh

Biomedical Image Analysis

Applications of Grids in Translational Research

- Multi-site therapy monitoring
  - Patients accrued at many sites
  - Quantify effects of treatments in a uniform manner across multiple sites
  - Carry out reproducible controlled studies
  - Collaborative studies where different researchers produce complementary data sets
- Make use of ensembles of datasets to develop excellent screening algorithms
- Specialized algorithms for molecular imaging

Multi-site therapy monitoring

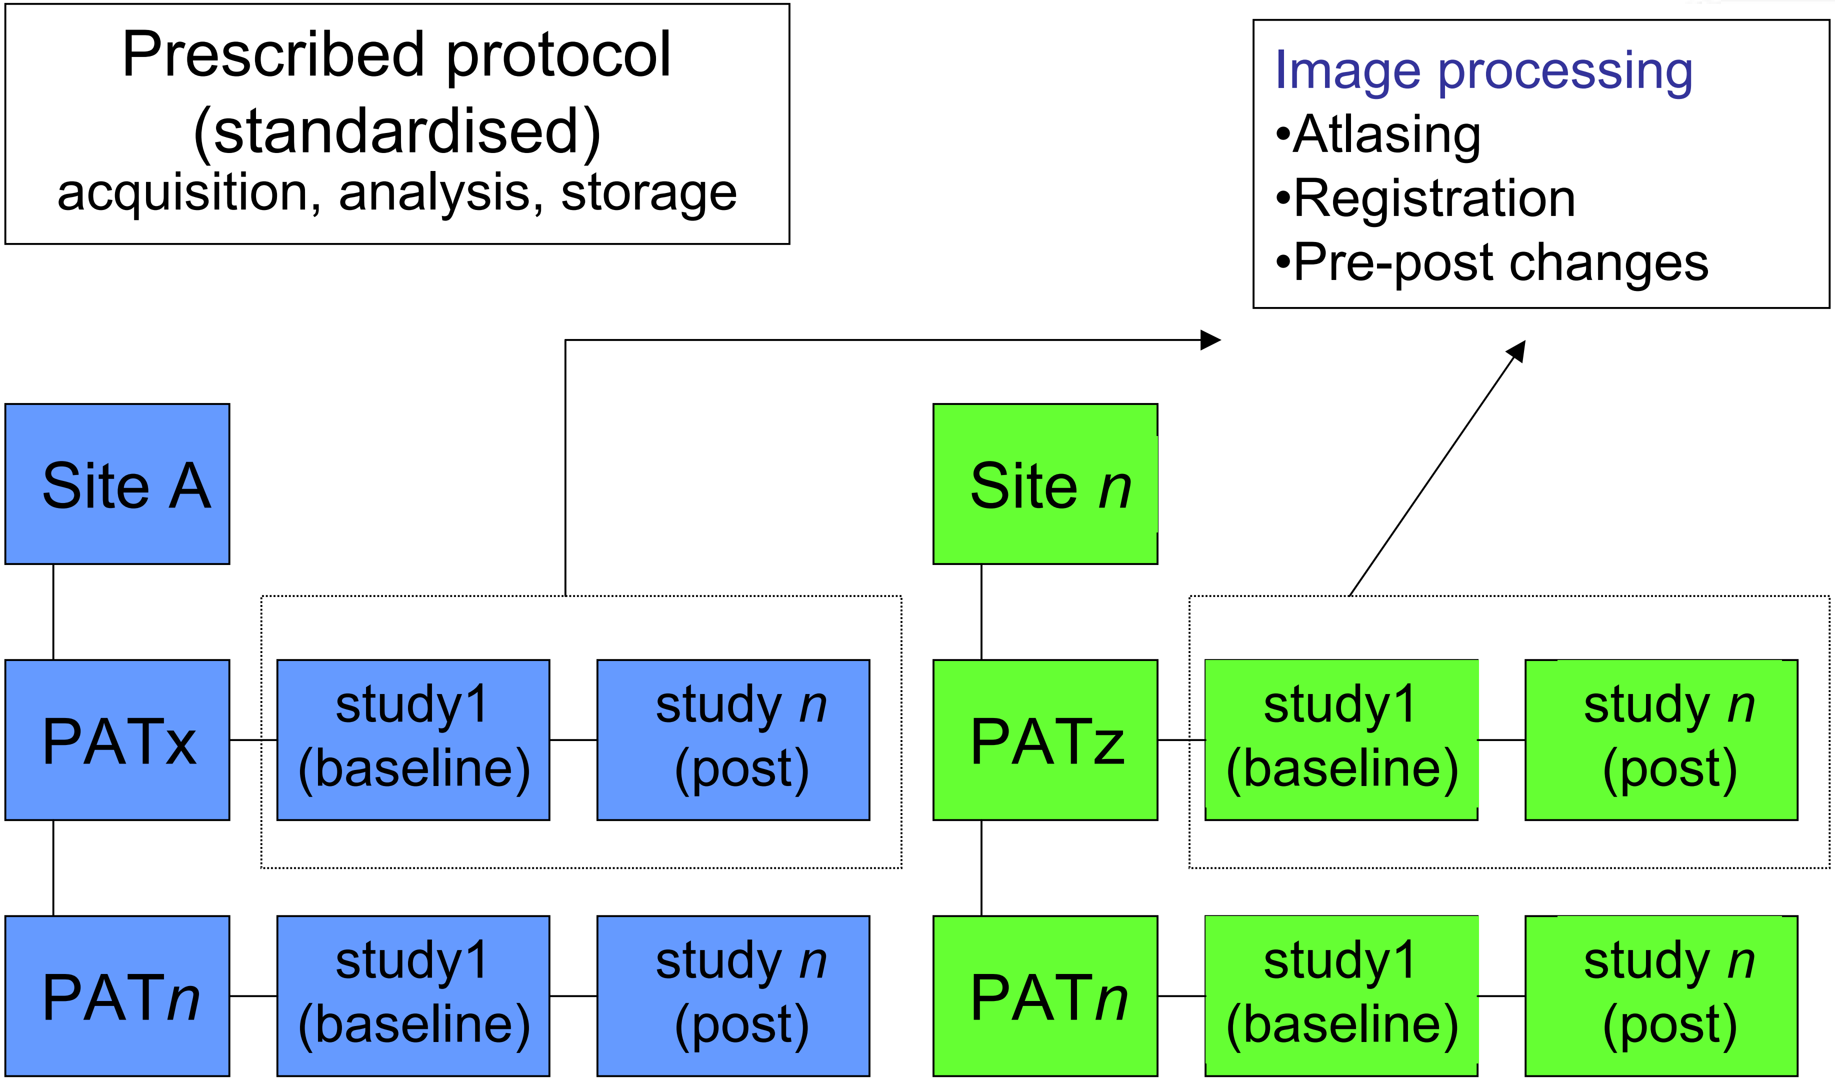

(Knopp)

Multi-site therapy monitoring

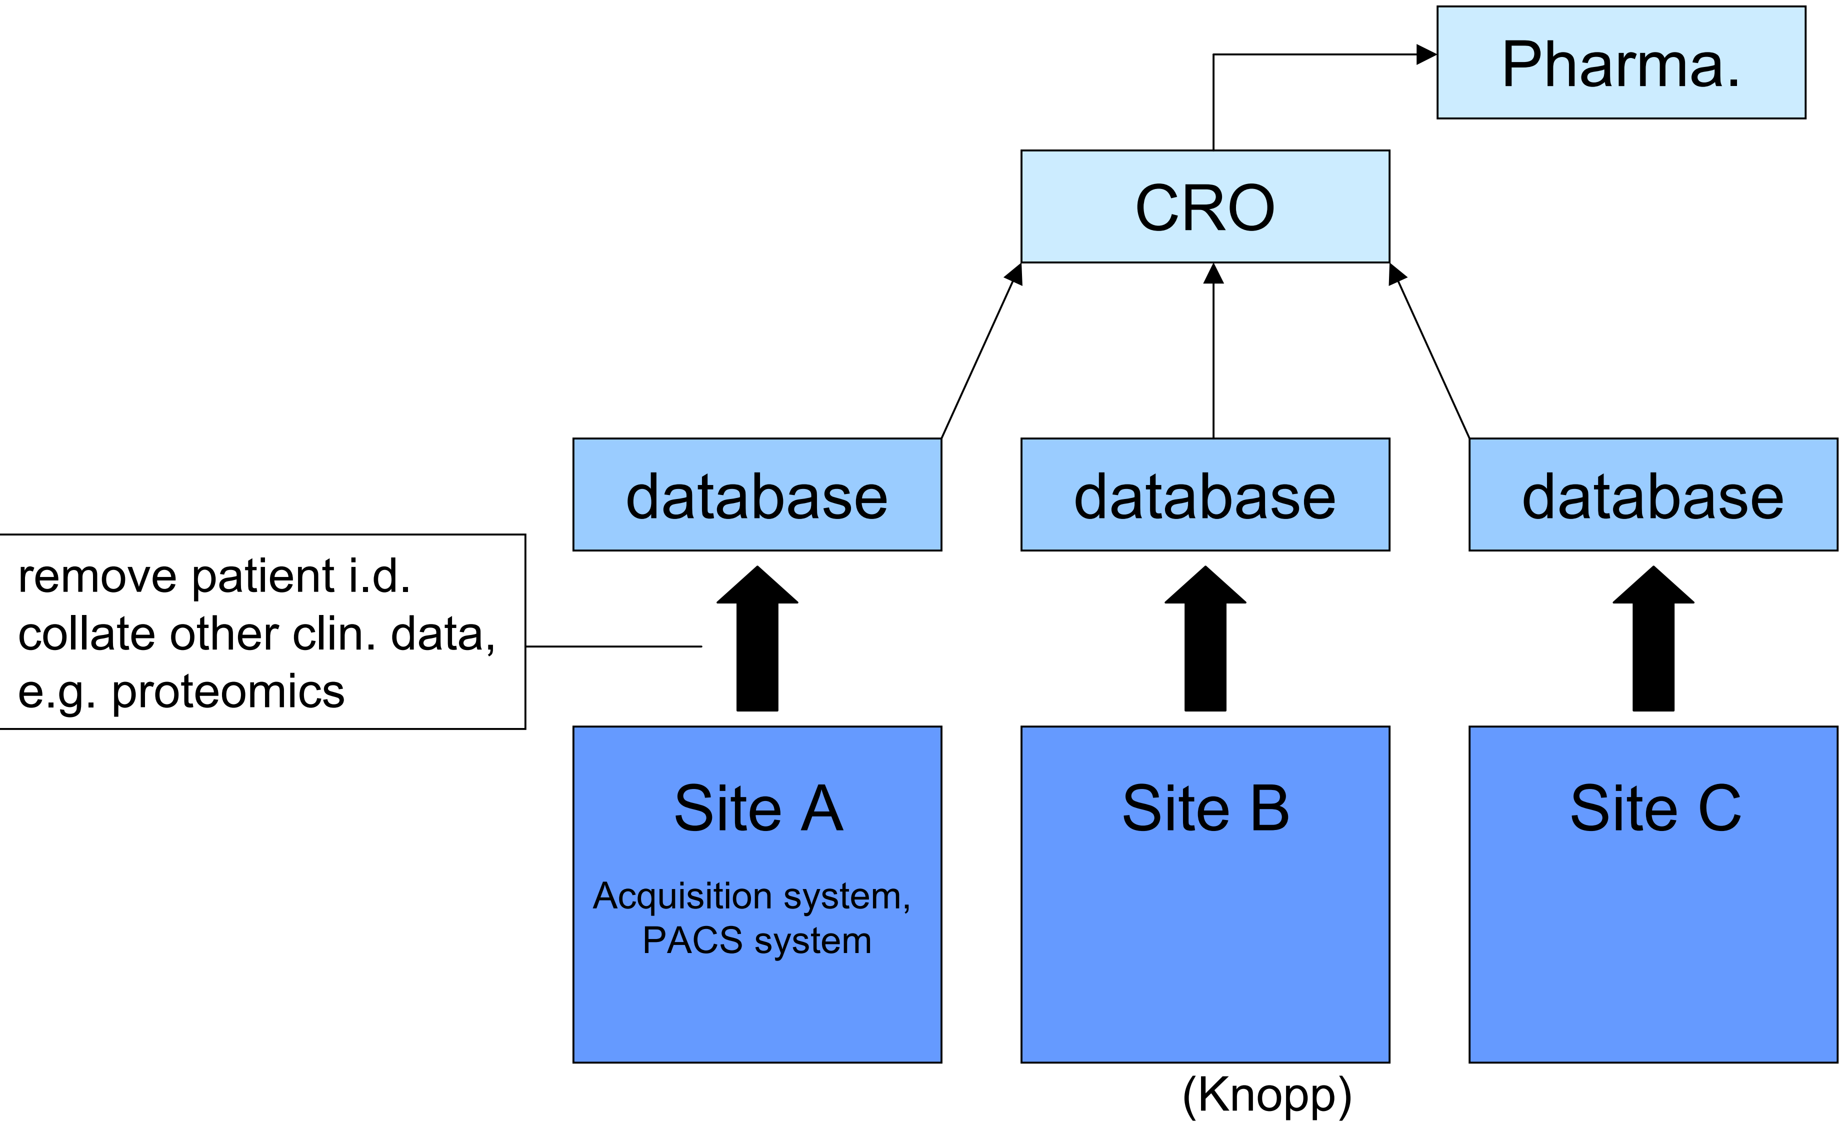

Image Processing

- How is the data stored?

raw Image (reconstructed)

- If we need to study image analysis strategies
- Sites do not necessarily have information analysis knowledge
- Potentially lots of data fMRI -1000s of images per patient

processed (e.g. fMRI, DCE parametric maps)

- Easily assess therapy induced changes
- Sites use prescribed analysis protocol
- Smaller data sets, more compatible with slower networks

(Knopp)

Image/Data Application Areas

Processing Remotely-Sensed Data

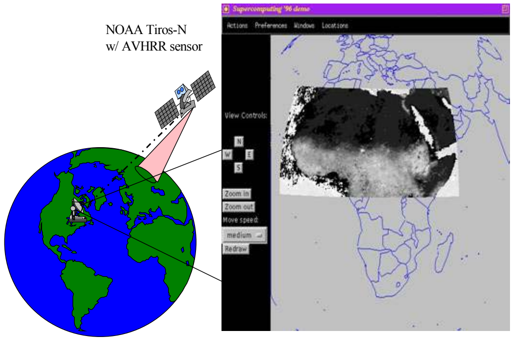

Satellite Data Processing

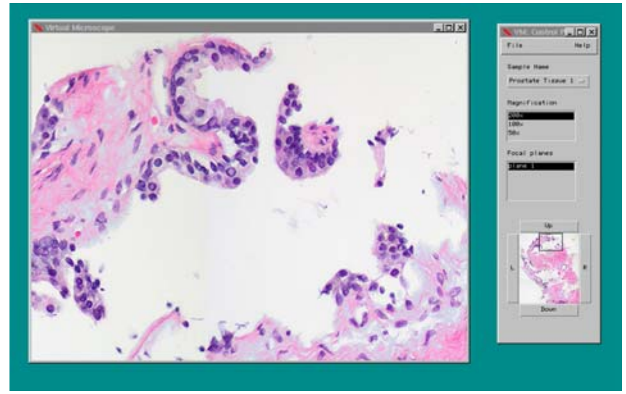

Digital Pathology

Imaging, Medical Analysis
and Grid Environments (IMAGE)
September 16 - 18 2003
e-Science Institute,
15 South College Street, Edinburgh

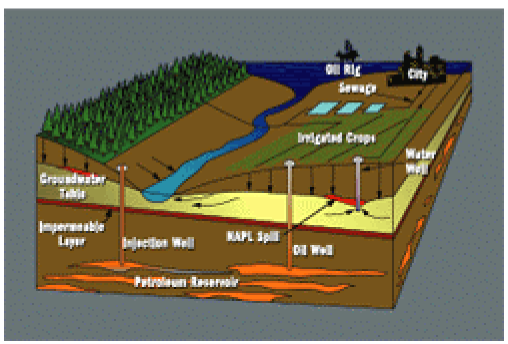

Managing Oilfields,
Contaminant Transport

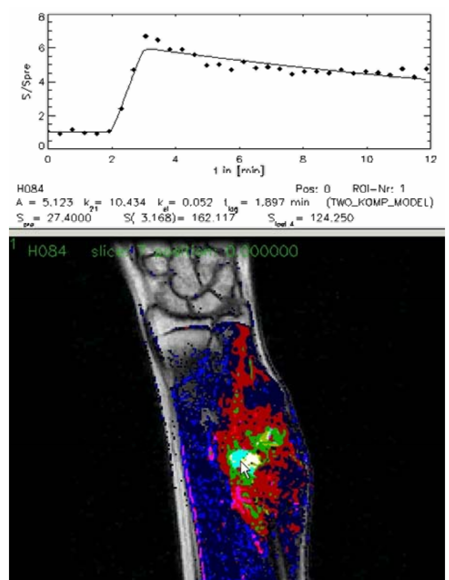

DCE-MRI Analysis

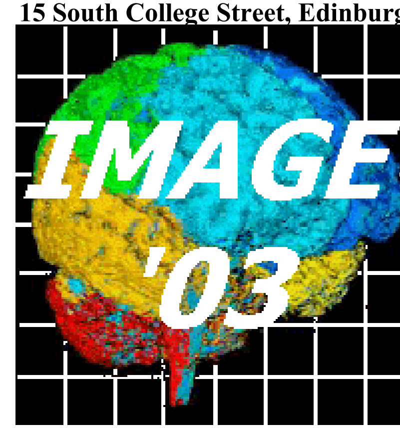

Biomedical Image Analysis

Metadata and query (IMAGE workshop)

- Deal with ever changing data models, changing classification schemes, ontologies (e.g. is it a plasma membrane protein or shuttling protein?)
- Need precise definitions of data transformations/filters to ensure reproducibility
- Want tools that are as easy to use as Google – ability to select data without presupposing relationships
- Separate "concrete or well defined" entities from "abstract" concepts
  - Should these be dealt with differently?
  - Not clear that there is consensus on what is defined although there are clear limiting examples

Computer Assisted Annotation (IMAGE Workshop)

Computer assisted annotation

- Information off of DICOM, tiff file (e.g. date)
- DICOM Information from manufacturer can be captured and translated
- What is image of?
- Need to capture metadata from grant application to experiment
- Capture and describe protocols
- Store and capture metadata in way one can reason (i.e. use ontology)
  - Potential ability to detect inconsistencies, contradictions etc
  - Semantic diff
- Feature detection should generate ontology tags
- Metadata should be signed in a way that identifies person or group

Computer Assisted Annotation (IMAGE Workshop)

- Microscopy – operating conditions of microscope, bandwidth of filters, etc. Each instrument type has different ways of sensing and delivering information. Phase contrast v.s. fluorescence.
- Correct interpretation of data requires description of physics
- Physics information needs to be captured
- End user does not know much of this so the process of encoding and utilizing physics information needs to be automated

OSU Research, Treatment Grid

- Distributed computers and databases
- Sets of interacting web/grid services
- Controlled vocabularies, metadata management
- Ubiquitous access to all clinical, laboratory, radiology, pathology, treatment data
- Services regularly scan patient information to evaluate interventions
- Services regularly aggregate and mine patient information to evaluate how to optimize treatment

Molecular data and OSU Testbed Effort

- Data sharing in OSU shared resource
  - Support for all data sharing in hundreds of research studies in OSU comprehensive cancer center
- State of Ohio BRTT
  - Integration of clinical, genotype, proteomic, histological, gene regulatory data in context of 4 translational research projects
  - $2M per year to fund development of bioinformatics data sharing infrastructure

Examples of data sources to be Integrated: Examples of data types that are generated or referenced by OSUCCC Shared Resources

.

Shared Resource	Example data types
Molecular Cytogenetics	Datasets from Karyotype analysis, data from SKY/FISH experiments, Karyotype images (when available), probe category, probe type, chromosomes to which results apply, and diagnosis results.
Analytical Cytometry	FACSCaliber output, raw data from DiVaOption system, sample identification, study id, description of the analysis process.
Genotyping and Sequencing	DNA sequence information for the sample, primer sequence, PCR sequencing information, sample identification, chromatic graphs.
Microarray	Output from Affymetrix gene expression and Custom microarray analyses, RNA integrity values, description of the process, DAT files (image data), CEL files, expression estimates from Affymetrix and dChip software, Genepix output, plate files describing the mapping between wells and probe ids, print files, well location on the microarray for the specimen under study.
Mouse Phenotyping	Digitized radiographic, gross, and histologic images, hematology characterization, attributes associated with complete clinical chemistry, description of all findings, the genotype of the mouse.
Real Time-PCR	Description of sample plate, raw and processed output files.
Tissue Procurement	Anonymized pathology report, age, gender, race, tissue procurement id, patient id (if consent form is available), consent form, virtual slide of the tissue, if available.
Leukemia Tissue Bank	Sample processing date, date sample taken from the patient, accession id, patient id, patient name and last name, specimen type, number of tubes, diagnosis info, protocol #, raw cell count, total cell count, number of vials per cell, percent viability, technician name.
Proteomics	1D and 2D gel images, sample information (PI name, analysis method, instrument name), diagnosis, protein expressions for spots.
Clinical Trials Office	***Protocol descriptions, investigator, title, status, approval processing ids, clinical data for patients on protocols, lab reports, adverse events, and trial outcomes***.

NCI CCC® A Comprehensive Cancer Center Designated by the National Cancer Institute

Center for Grid Enabled Image Analysis (NIH BISTI Program)

Biomedical Research	Imaging Research	Computer Science
Ensure success of biventricular pacing	Semi-gated cardiac imagery analysis	Machine vision On-demand large data analysis
Role of oncogenes in development	Multiple modality mouse placenta imaging, information synthesis, registration	Image analysis in ensembles of very high resolution 3-D imagery. Interactivity
Mechanism of ischemic cardiac injury	Synthesis of multimodal imaging, genotype, gene expression, proteomic data	Grid data management and query, Information integration involving multi-modal image, molecular data

Data Storage

- Clusters provide inexpensive and scalable storage
- $50K-$100K clusters at Ohio State, OSC, U. Maryland range from 16 to 50 processors, 7.5TB to 15TB IDE disk storage
- Data declustered to cluster nodes to promote parallel I/O
- Uses DataCutter and IP4G toolkits for data preprocessing and declustering
- R-tree indexing of declustered data

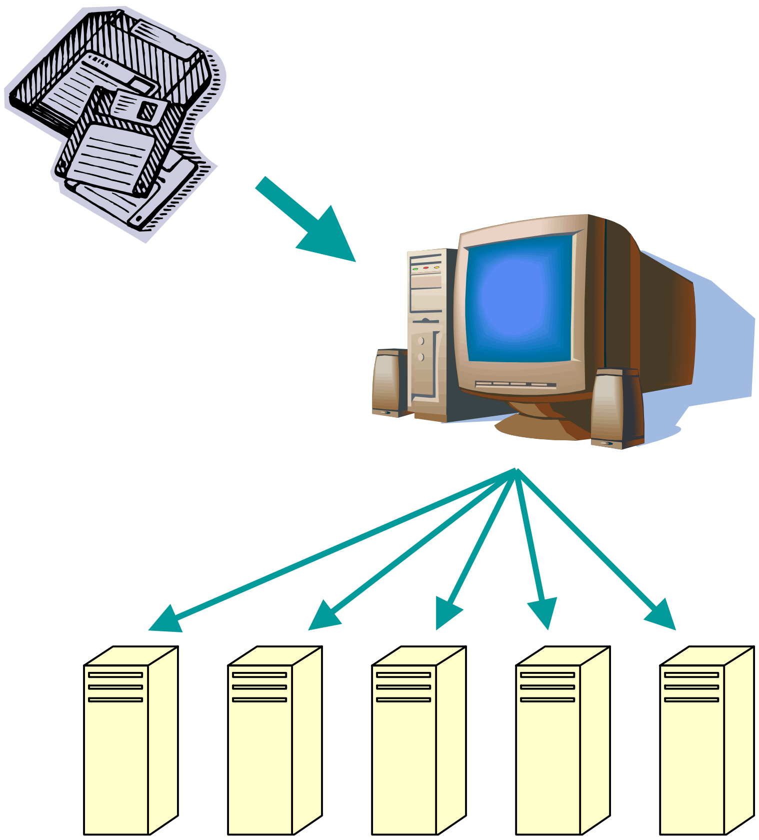

Ohio Supercomputing Center Mass Storage Testbed

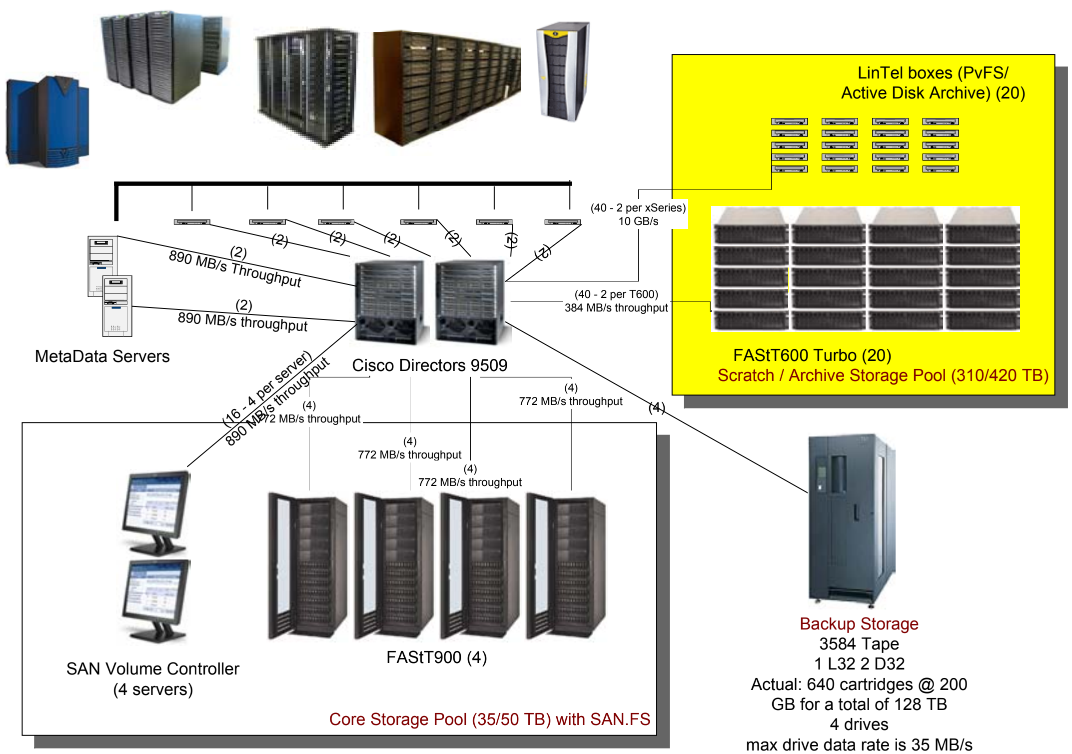

- **50 TB of performance storage**
  - home directories, project storage space, and long-term frequently accessed files.
- **420 TB of performance/capacity storage**
  - Active Disk Cache - compute jobs that require directly connected storage
  - parallel file systems, and scratch space.
  - Large temporary holding area
- **128 TB tape library**
  - Backups and long-term "offline" storage

IBM’s Storage Tank technology combined with TFN connections will allow large data sets to be seamlessly moved throughout the state with increased redundancy and seamless delivery.

Services

- Filter-stream based distributed execution middleware (**DataCutter, STORM**)
- Grid based data virtualization, query, on demand data product generation (**STORM, Active ProxyG, Mako**)
  - Supports distributed storage of XML schemas through virtualized databases, file systems
- Distributed metadata management (**Mobius Global Model Exchange**)
  - Track metadata associated with workflows, input image datasets, checkpointed intermediate results

Biomedical Informatics Research Network Testbed

- Goal
  - Remote access, processing of subsets of large, distributed virtual slides
- DataCutter, IP4G
  - Indexing, querying, caching and subsetting
  - Image processing by custom routines, VTK and ITK layered on DataCutter.
  - Use of heterogeneous, distributed clusters for data processing.

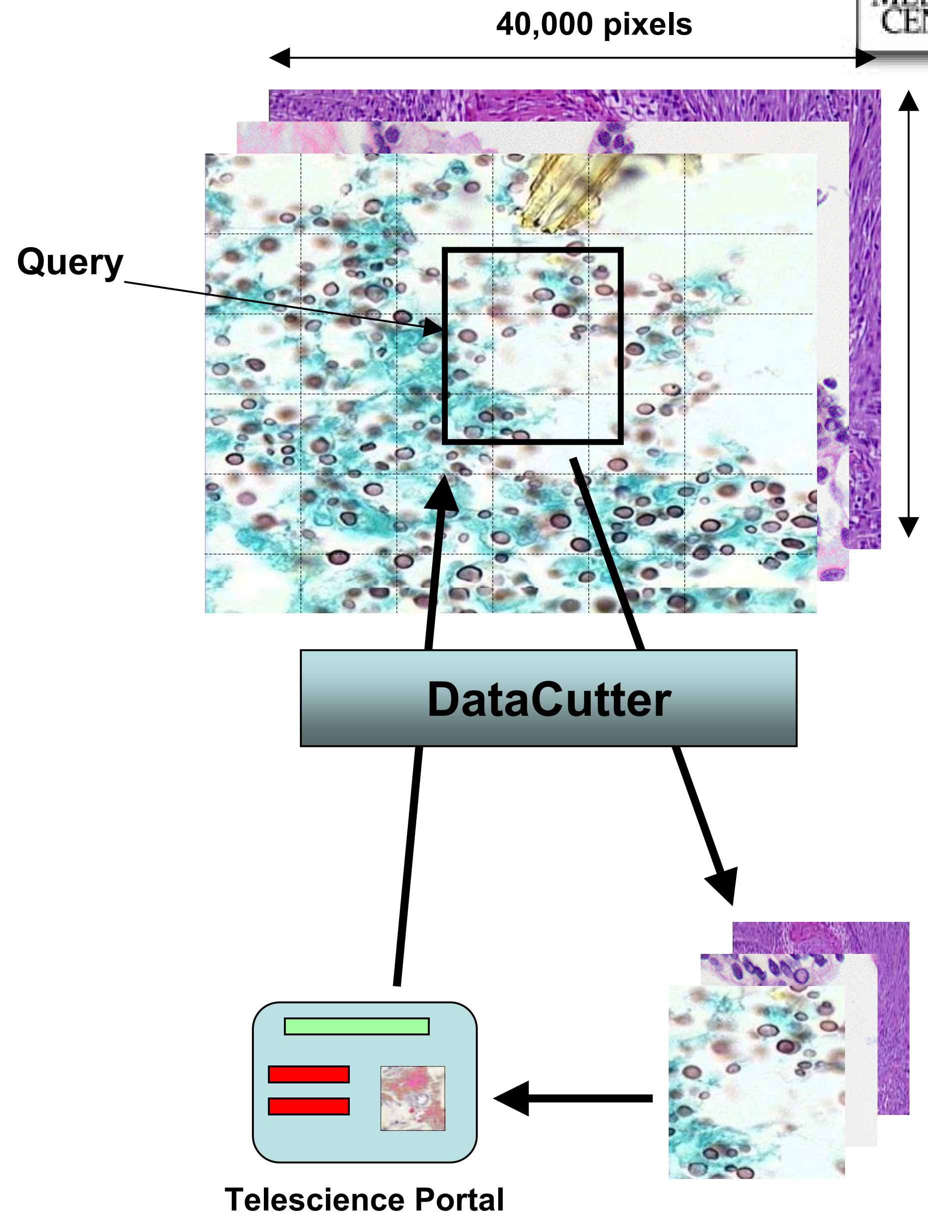

Telescience Portal and DataCutter

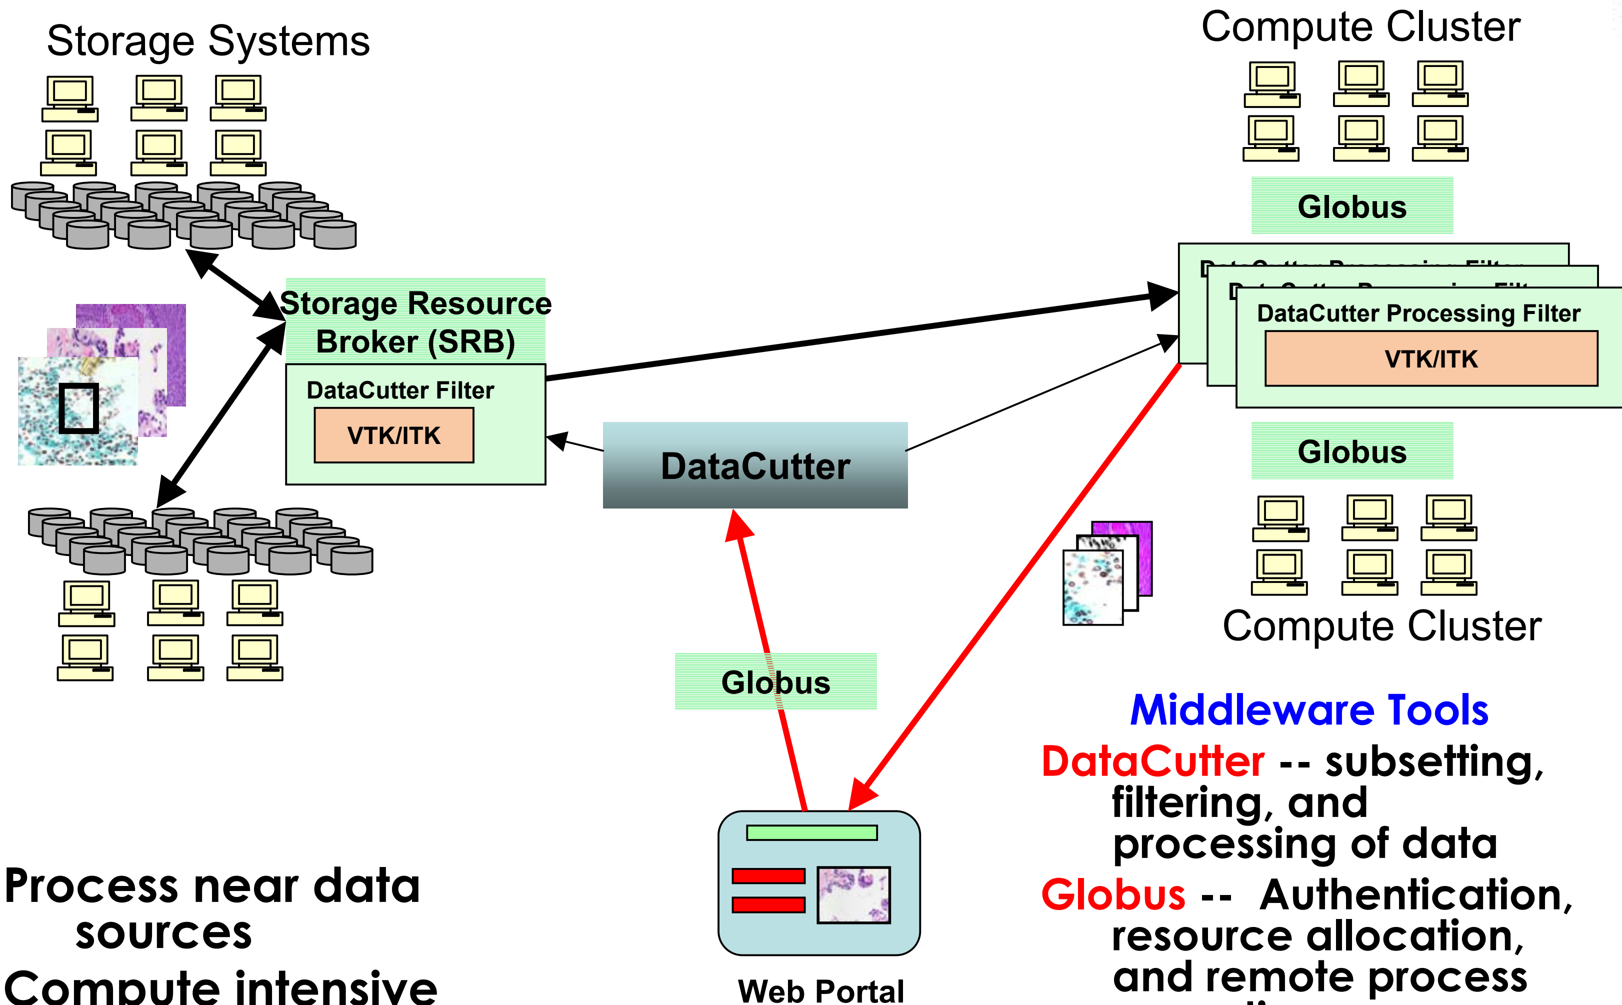

Process near data sources

Compute intensive operations on compute clusters.

Middleware Tools

DataCutter -- subsetting, filtering, and processing of data

Globus -- Authentication, resource allocation, and remote process execution.

SRB -- file I/O to different storage systems.

Virtual Slide Cooperative Study Support

- Children's Oncology Group, CALGB Cooperative Studies
- 30 slides/day – 30 GB/day compressed, 300GB/day uncompressed
- Remote review of slides
- Computer assisted tumor grading
- Tissue Microarray support
- **CALGB began November 2003, Childrens Oncology in May 2004**

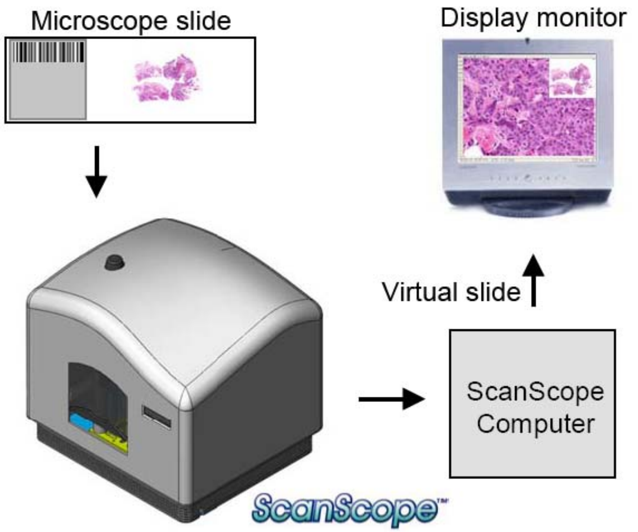

The ScanScope T108 Slide Scanning System. An entire microscope slide is rapidly scanned by the ScanScope®, creating a virtual slide that is viewed on a display monitor. The ScanScope® computer controls the ScanScope® using Aperio's console software.

Distributed, Federated and Integrated

- Consortia Group Support
  - Virtual slides
  - OSC 420 on-line Terabyte storage system
  - CALGB & COG
- OSU's Virtual Placenta Project
  - Embryonic development and gene expression
- BIRN: OSU/UCSD multi-photon project
  - Detailed microanatomic description of gene expression in brain
  - Search for rare mitoses

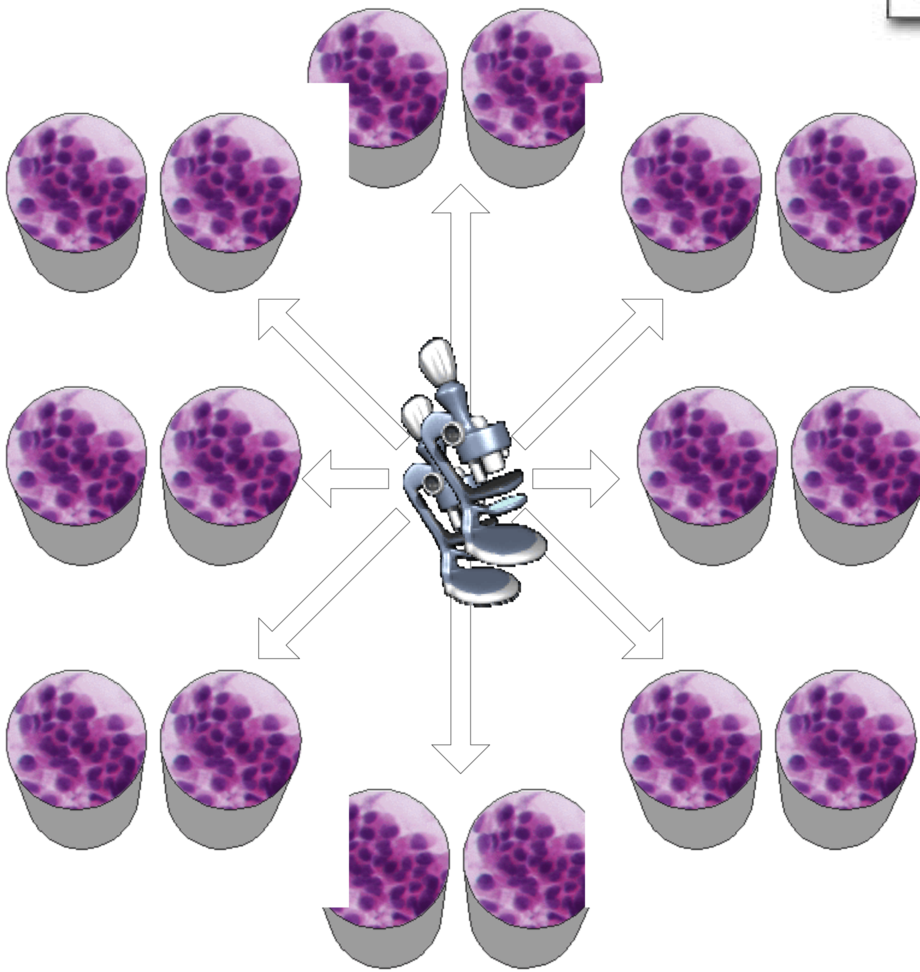

Case Study: Microscopy Image Reconstruction

- 3-D reconstruction and registration of virtual slides
- Quantitative characterization of tissue on 2-D/3-D slides, microCT
  - Cell morphology, structural attributes, anotomical description of gene expression
  - Quantitation of expressed proteins

3D Reconstruction Motivation

- Correlate feature changes across multiple slides
- Allow visualization of 3D structural geometry
- Extract additional information from volumetric data

- A registration problem

Challenges

- Large Image Size:
  - Mouse placenta image: 200X mag, 14000 x 14000 RGB, 570MB
  - Neuroblastoma image: 200x mag, 39000 x 49000 RGB, 5.3GB
- Large number of serial slices
  - Placenta : 100 to 500 slides
- High Magnification
  - 400X mag generates 4 times more data
- Scanner proprietary image compression format
  - Collaborate with scanner maker to get codec access
  - Currently using uncompressed slide stripes

Technical Considerations

- Large number of slices
  - As much automation as possible
- Large Dataset
  - Distributed storage
  - Out of core storage
- Partitioned datasets
  - Stripes on disk
- Image non-uniformity
  - Color variation across stripes
- Feature-based registration
  - Point-Point correspondence
    - User-interactive
    - Need a few good points
  - Iterative closest points
    - Unsupervised
    - Need good feature extraction
- Intensity-based registration
  - Mutual Information Registration
    - Unsupervised
    - Require complete images

Unregistered Images

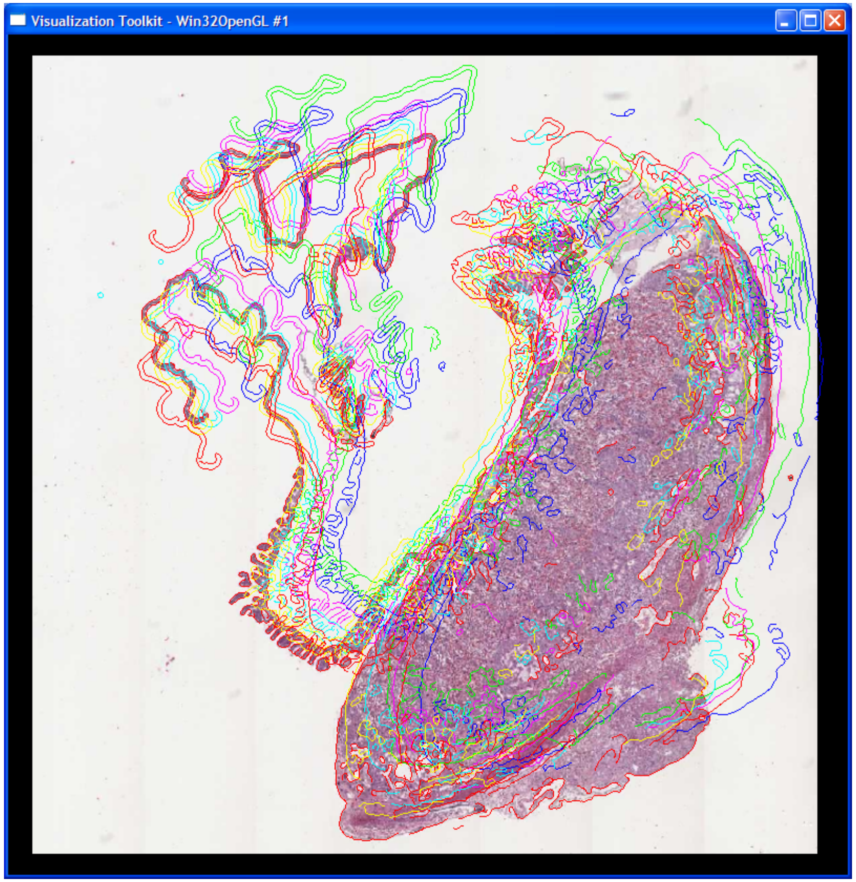

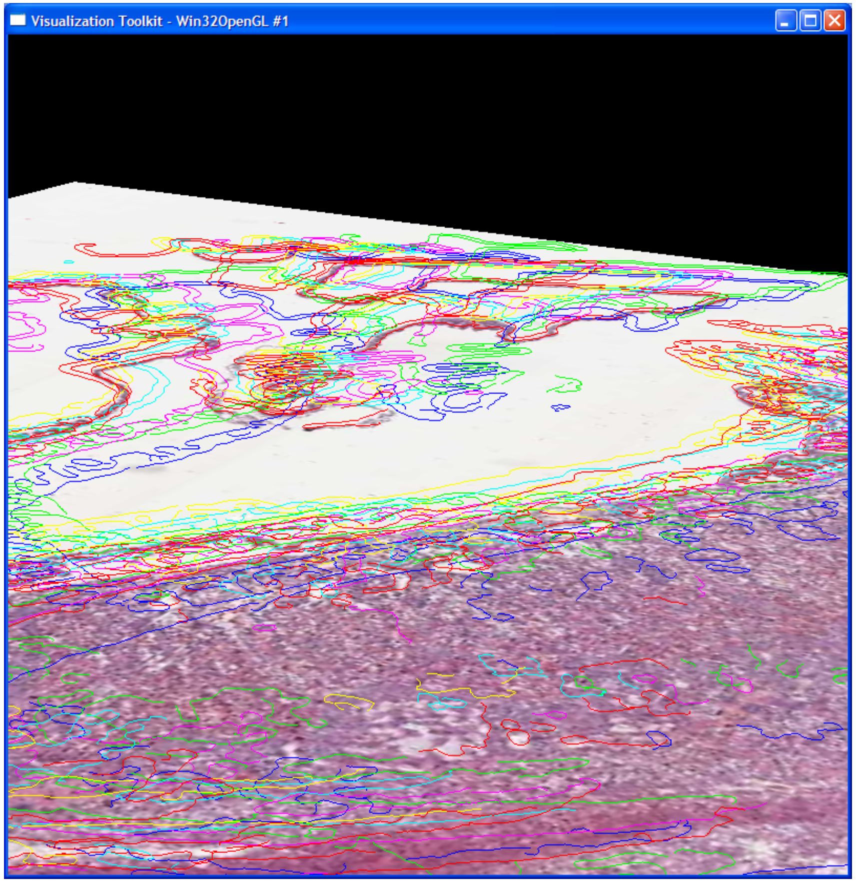

Registered Images

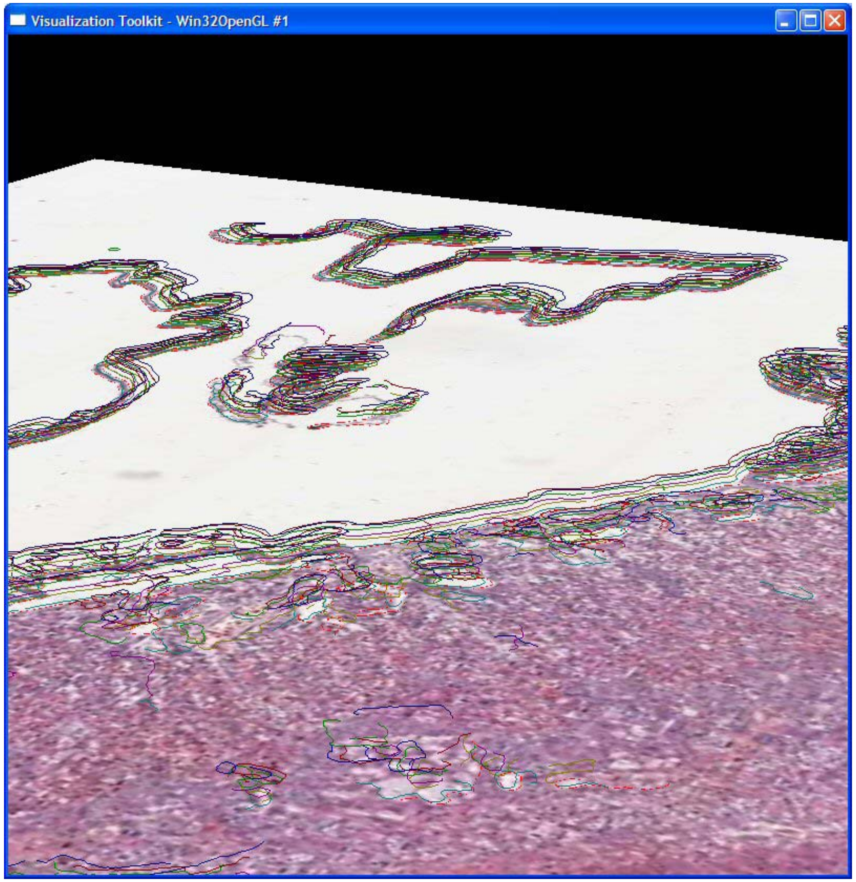

Case Study: Visible Human Dataset

- Data Set
  - Synthetic multi-resolution dataset motivated by NIH Visible Human dataset. *Synthetic dataset consists of 5 datasets at different spatial range and different resolution and the total size of the data is more than 1TB.*
- Declustering: Hilbert-curve based partition onto disks on a 20-node linux cluster.
- Indexing: We build a distributed hierarchical R-tree index based on the original physical coordinates.
- Client: Parallel renderer developed by Ohio Supercomputer Center, run on 8 processors

Visualization Application

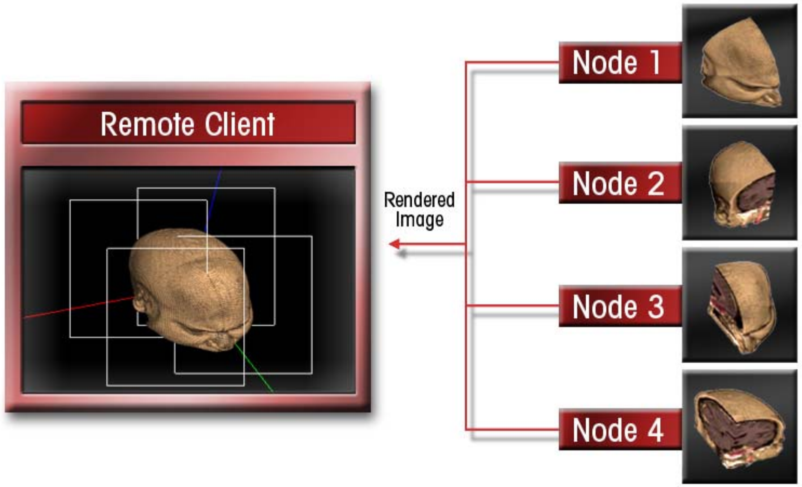

Rendering with MPI renderer

Ohio Supercomputer Center
Parallel Renderer

- MPI (MPICH-GM)
- OpenGL, hardware accelerated

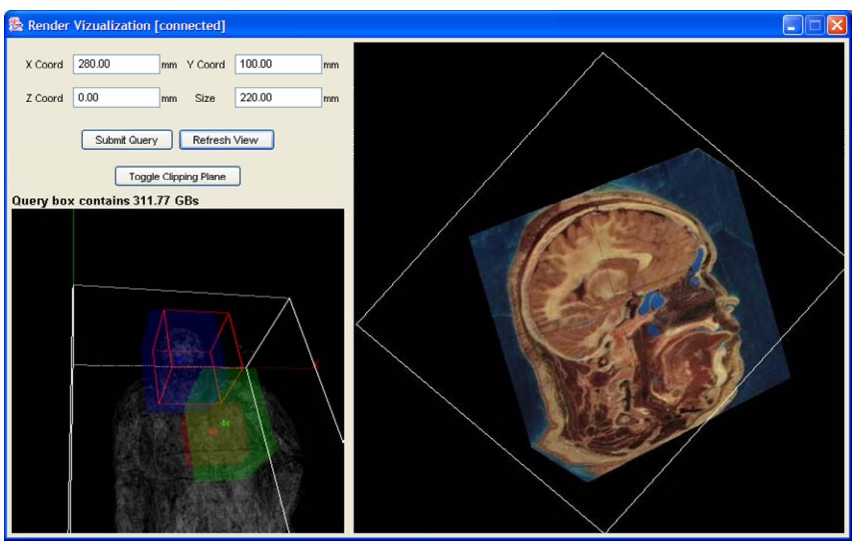

Remote Display

- Lightweight Java Client
- Connect to OSC renderer front-end
- Allow user query and interactive manipulation

Multiscale Pipeline

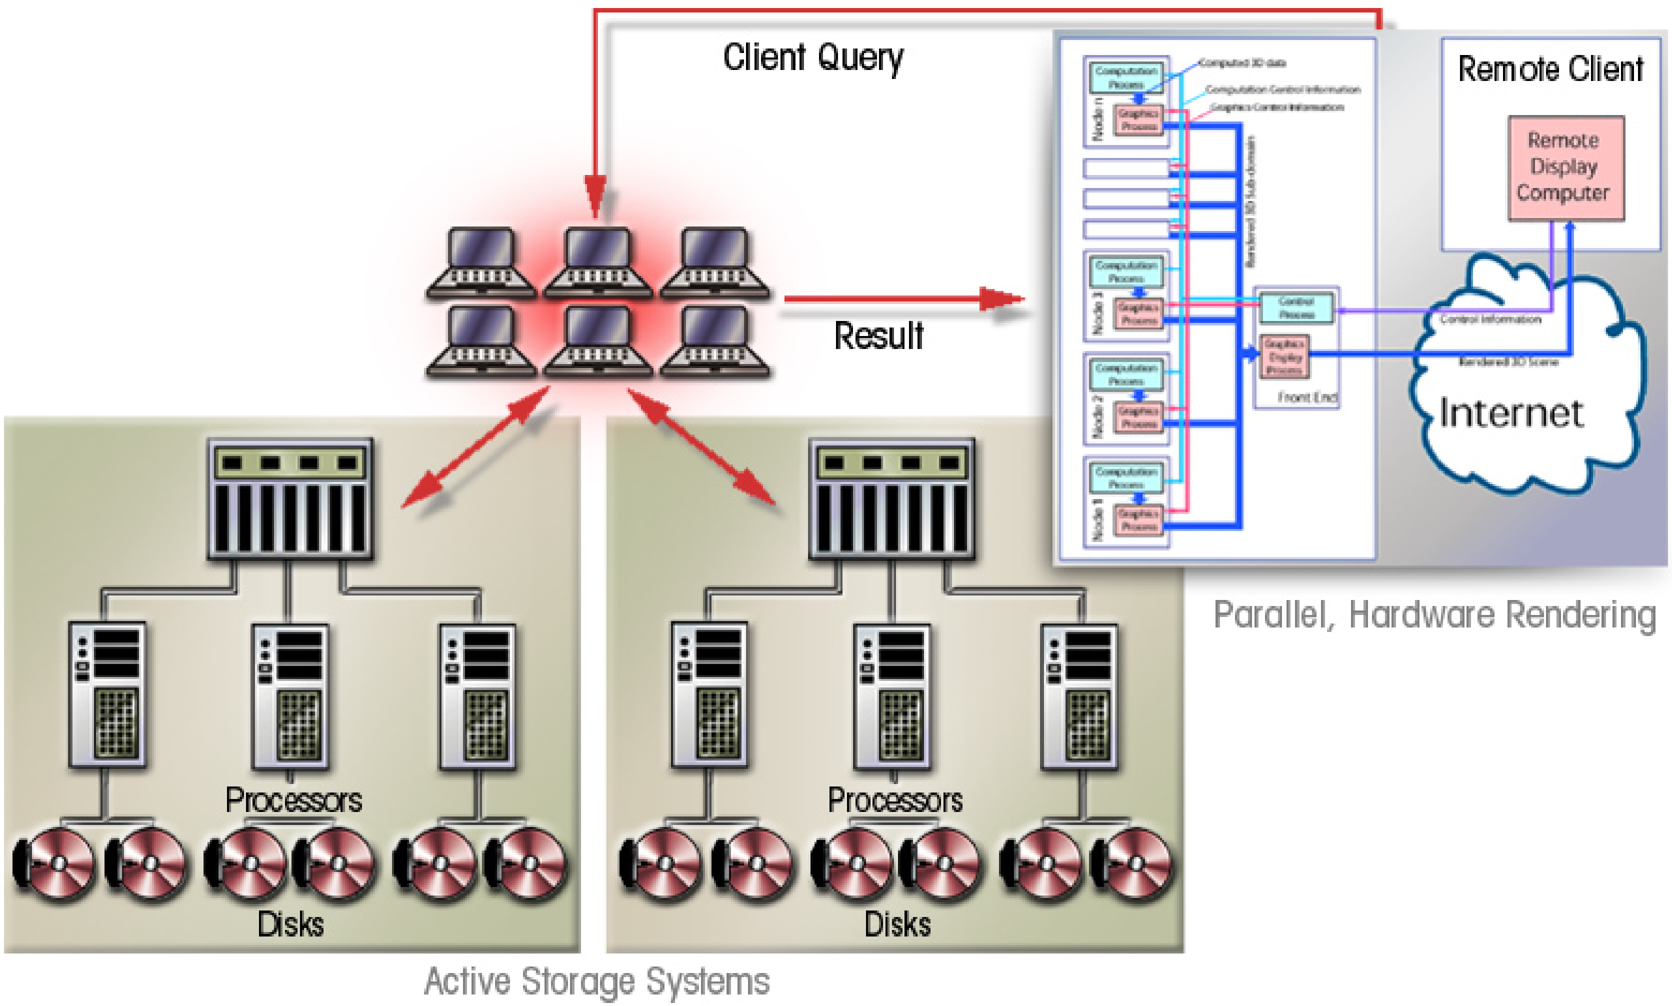

Questions and Answers

Ohio State Grid Middleware and Data Virtualization

DataCutter

- Flow control between components
- Schedulers place filters on grid processors (scheduler API)
- Stream based communication
- Data aggregation implemented as a component
- NPACkage, NMI
- www.datacutter.org

Combined Data/Task Parallelism

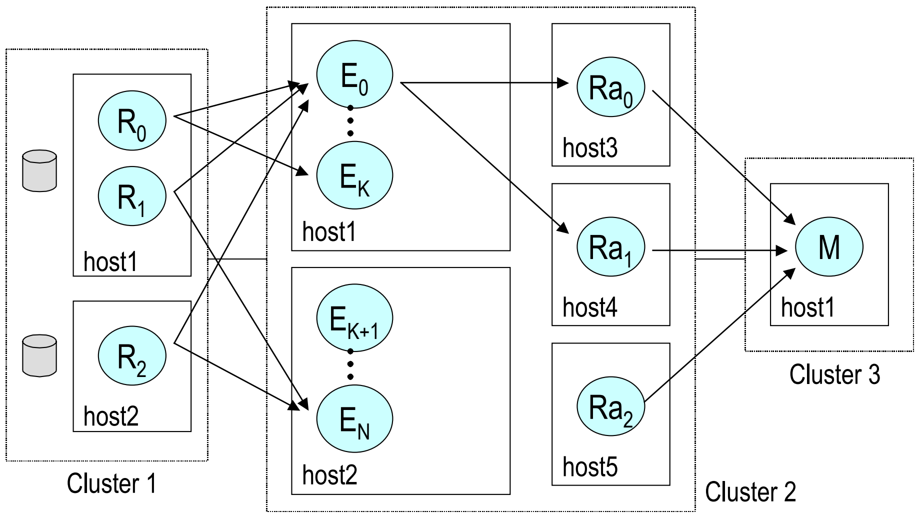

DataCutter: Support for Demand Driven Workflows

- Many data analysis queries have multiple stages
- Decompose into parallel components
- Strategically place components
- Create GT4 compliant services

Virtual Microscope

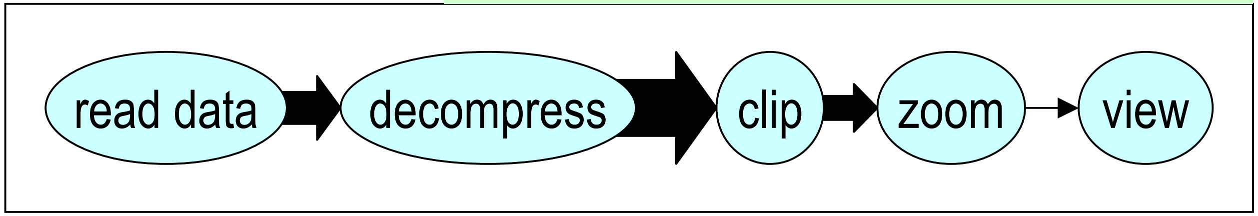

Iso-surface Rendering

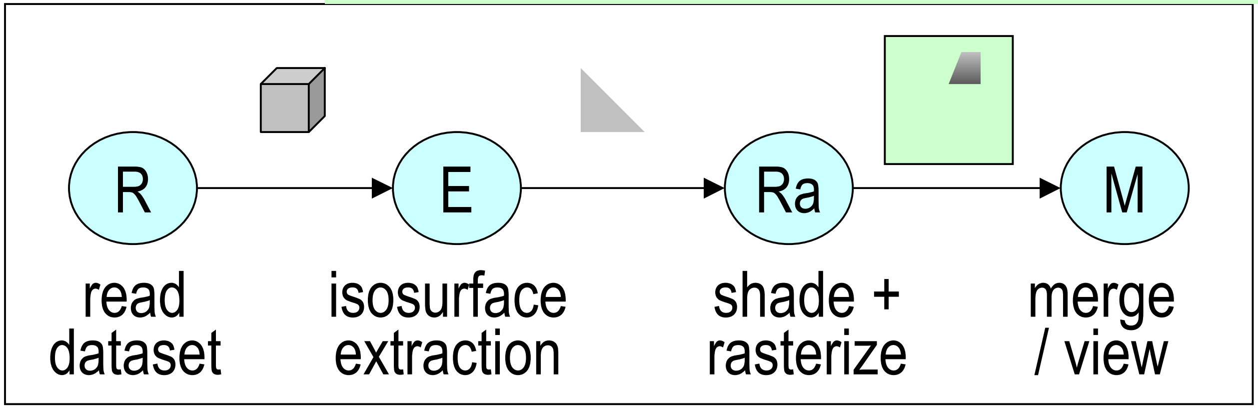

Image Processing For the Grid (IP4G)

- Based on DataCutter
- XML work flow description and component layout
- Serialization to and from DataCutter, and VTK/ITK filter initialization.
- User application models the image analysis process in XML or C++, and invokes the analysis.

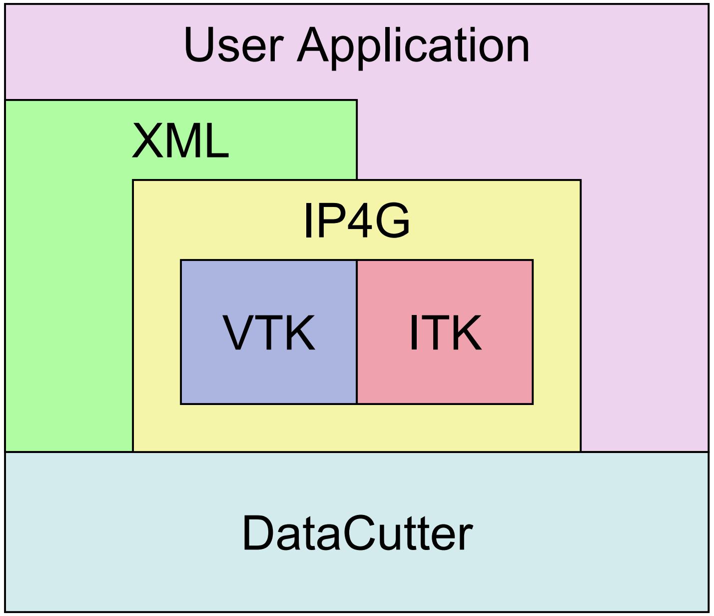

Automatic Data Virtualization

Scientific and engineering applications require interactive exploration and analysis of datasets.

Applications developers generally prefer storing data in files

Support high level queries on multi-dimensional distributed datasets

Many possible data abstractions, query interfaces

- Grid virtualized object relational database
- Grid virtualized objects with user defined methods invoked to access and process data

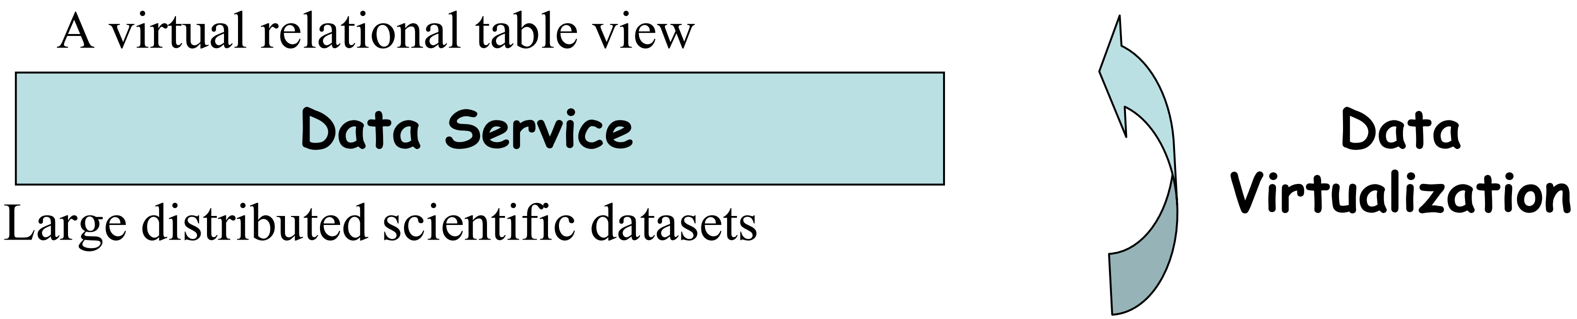

STORM Query Planning

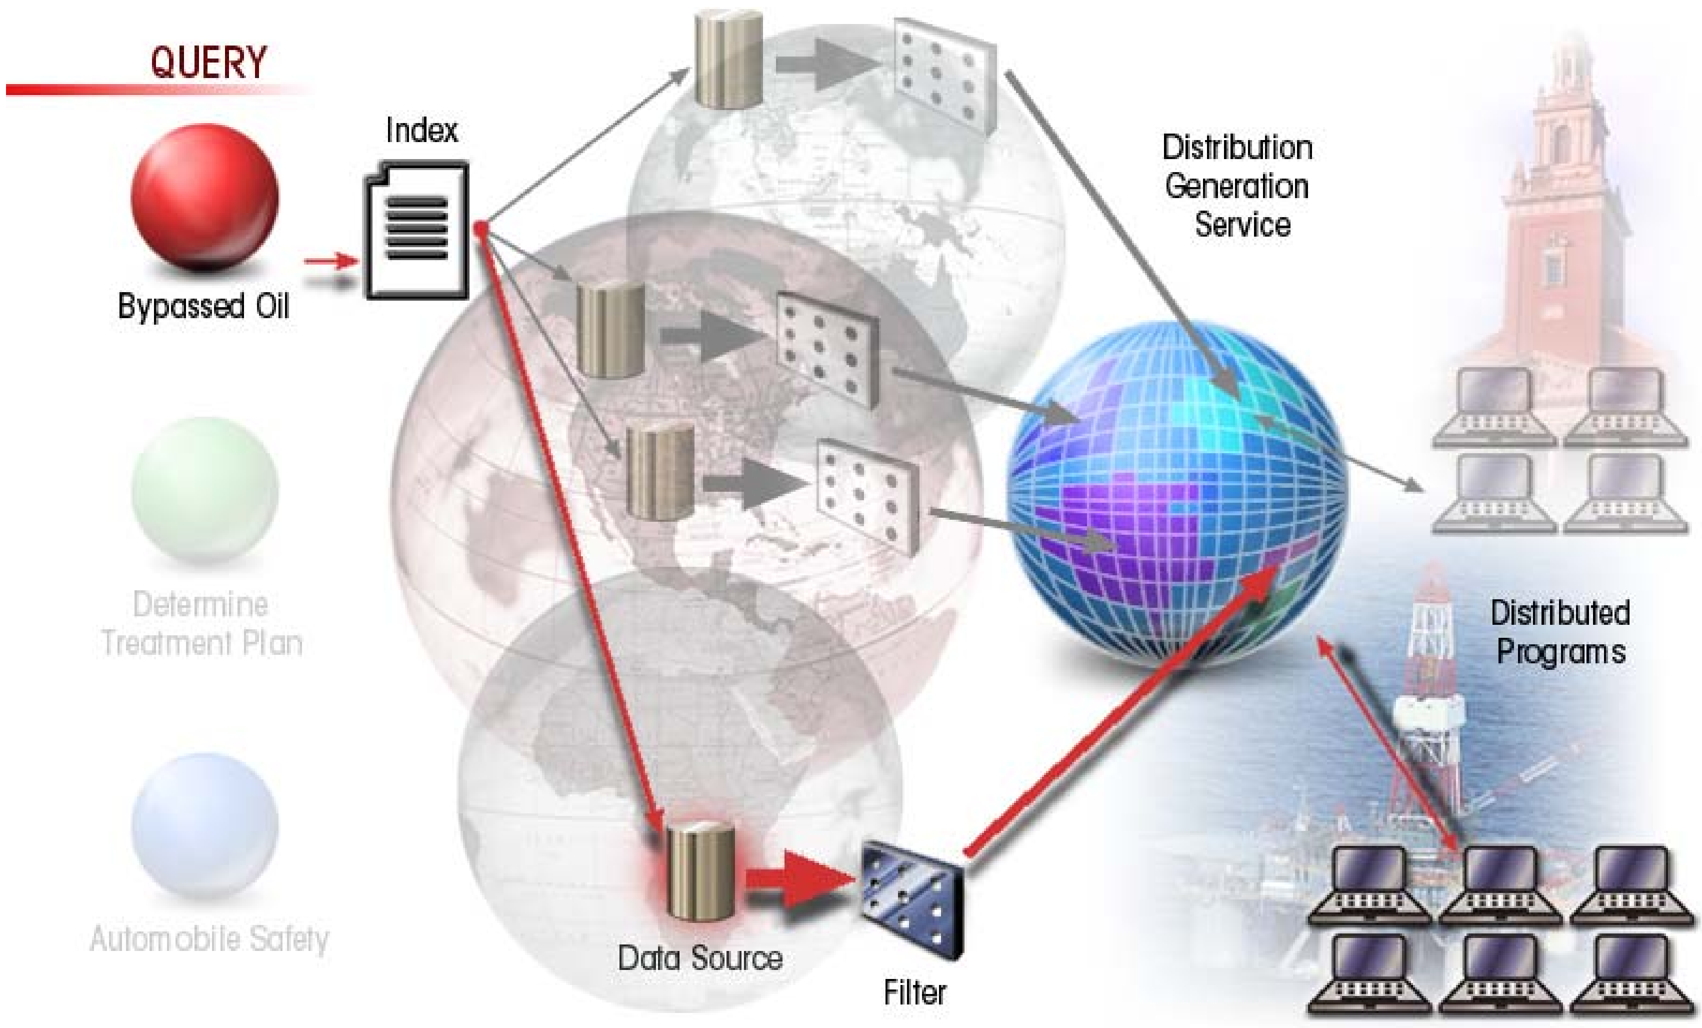

STORM Query Execution

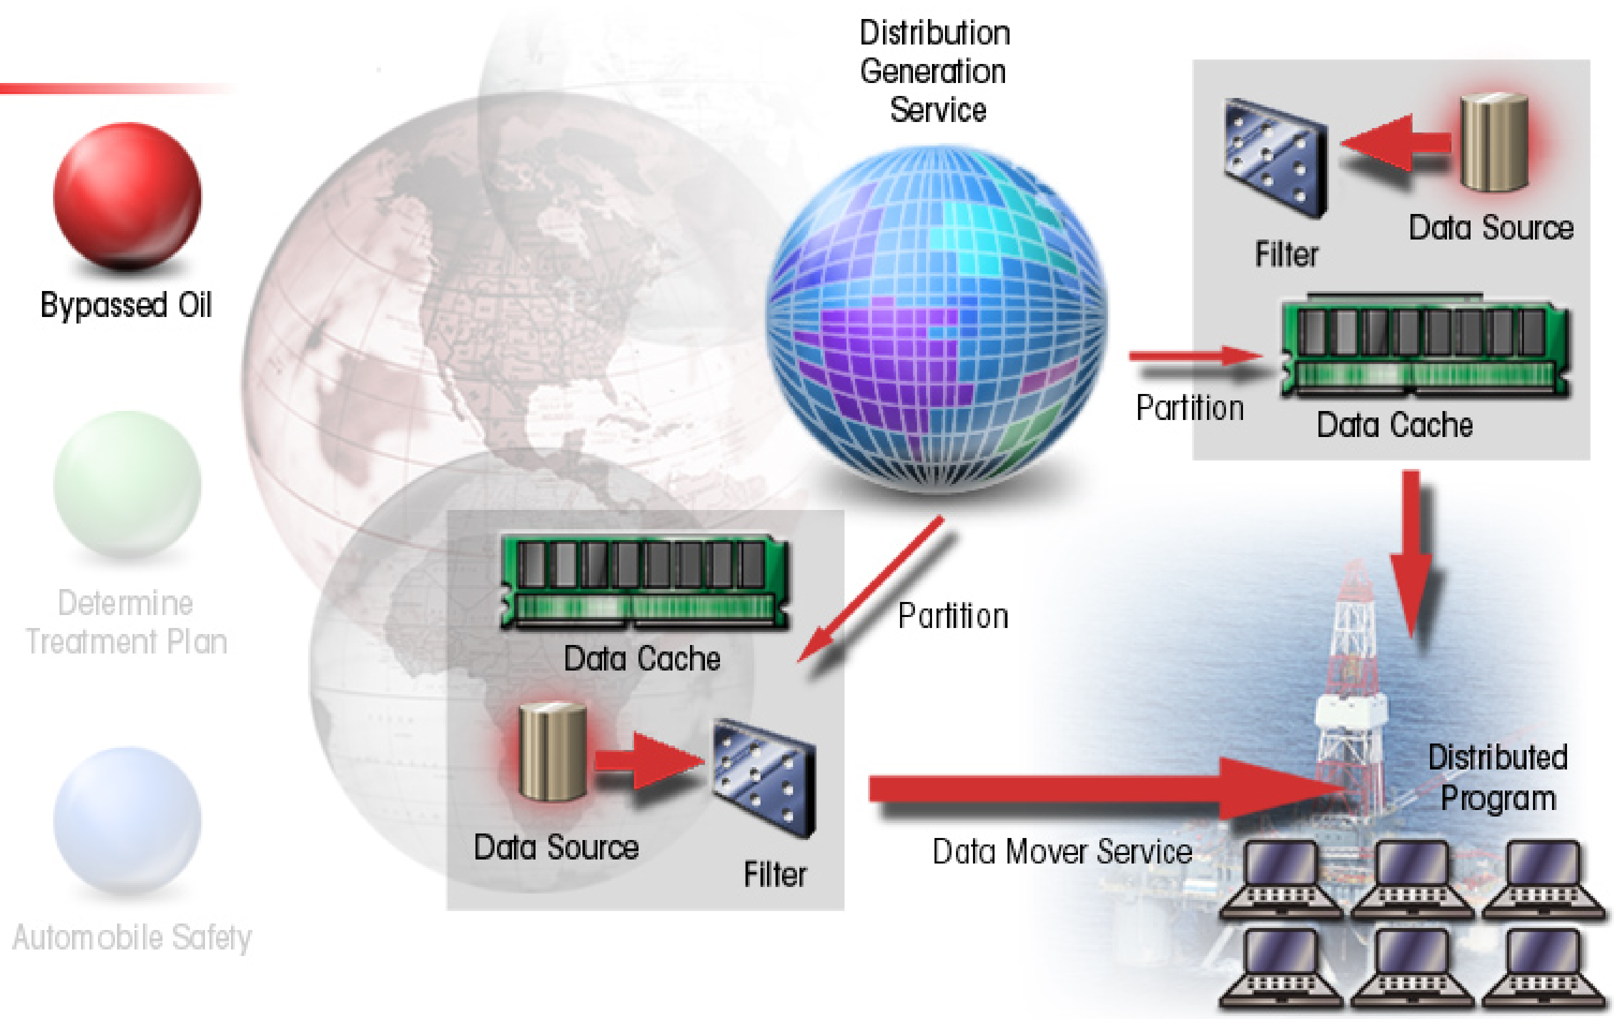

Data Returned as stream of tuples:

- Stand-alone client
- MPI program
  - MPI program provides partitioning function
  - Partitioning service generates mapping
  - Data mover sends data to appropriate MPI process
- Single or replicated copies of a DataCutter filter group

System Architecture

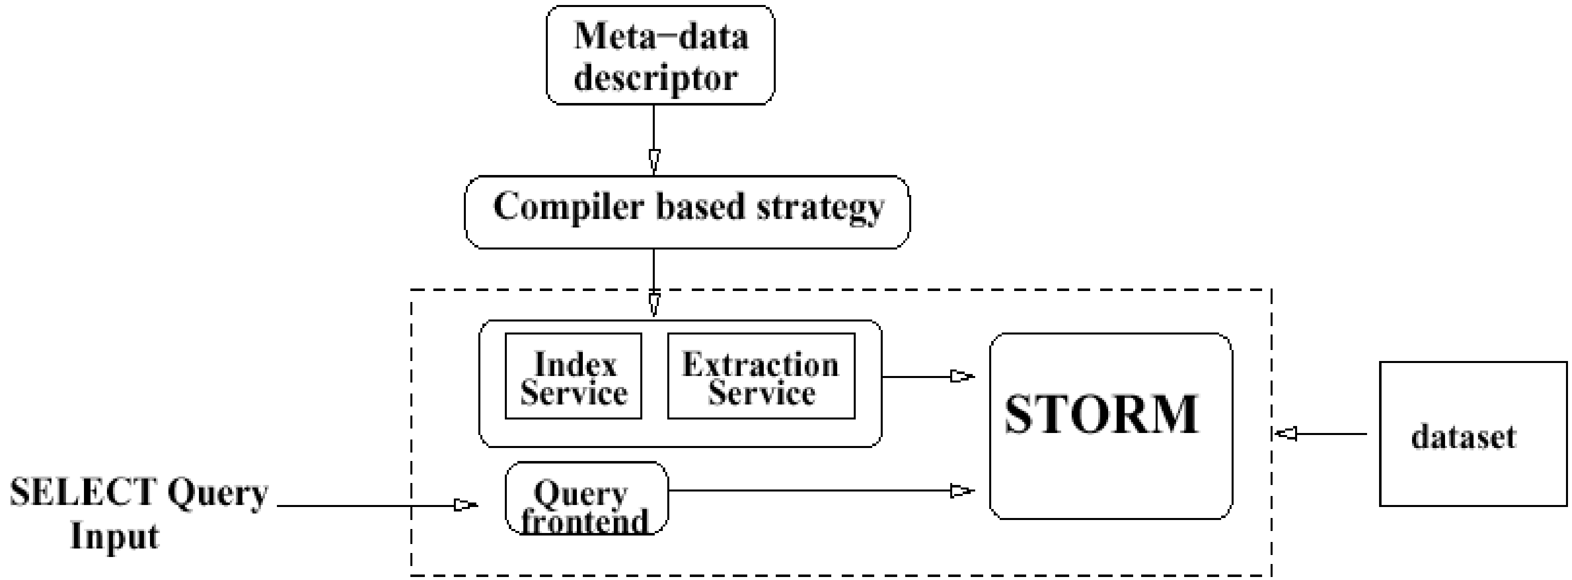

SELECT * FROM IPARS
WHERE RID in (0,6,26,27) AND TIME>1000 AND TIME<1100
AND SOIL>0.7 AND SPEED(OILVX, OILVY, OILVZ)<30.0;

Common operations: Subsetting, filtering, user defined filtering

Comparison with hand written codes

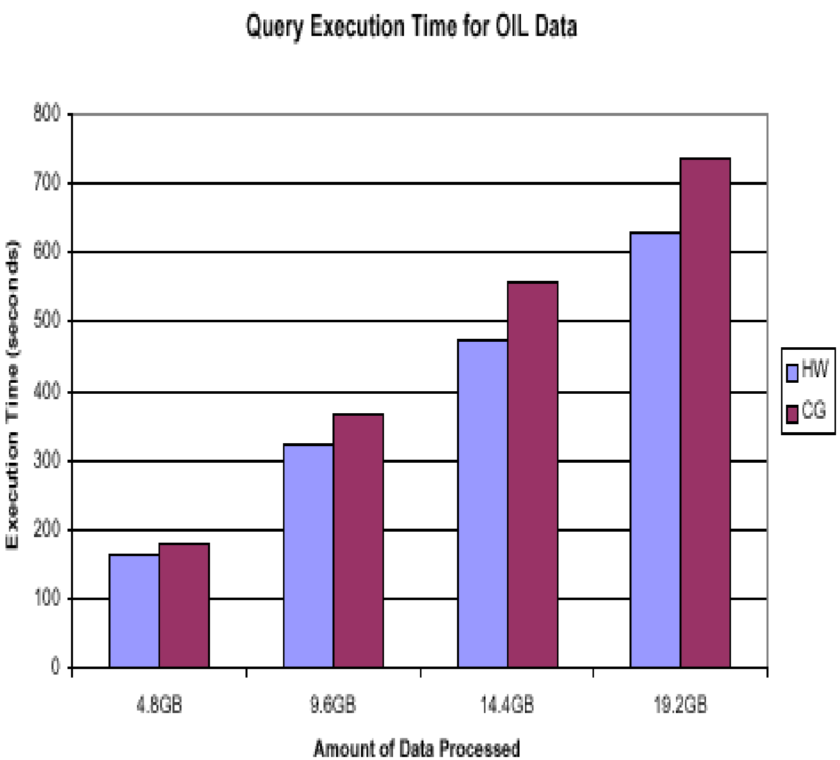

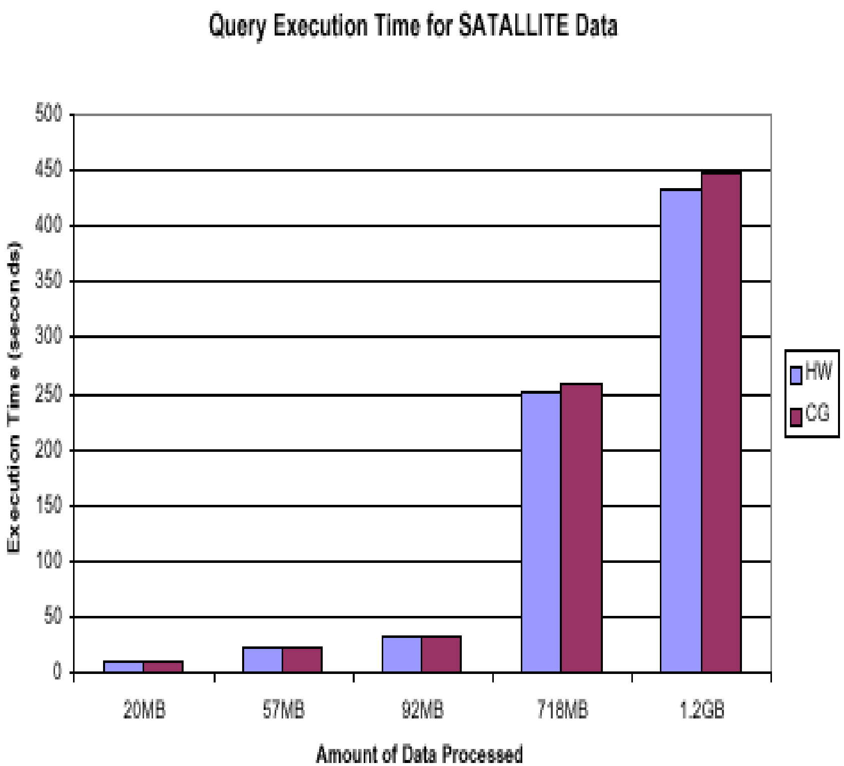

Dataset stored on 16 nodes.
Performance difference is within 17%,
With an average difference of 14%.

Dataset stored on a single node.
Performance difference is within 4%.

Components of Meta-data Descriptor

- Describe attributes, location of files, layout of data in files, indices

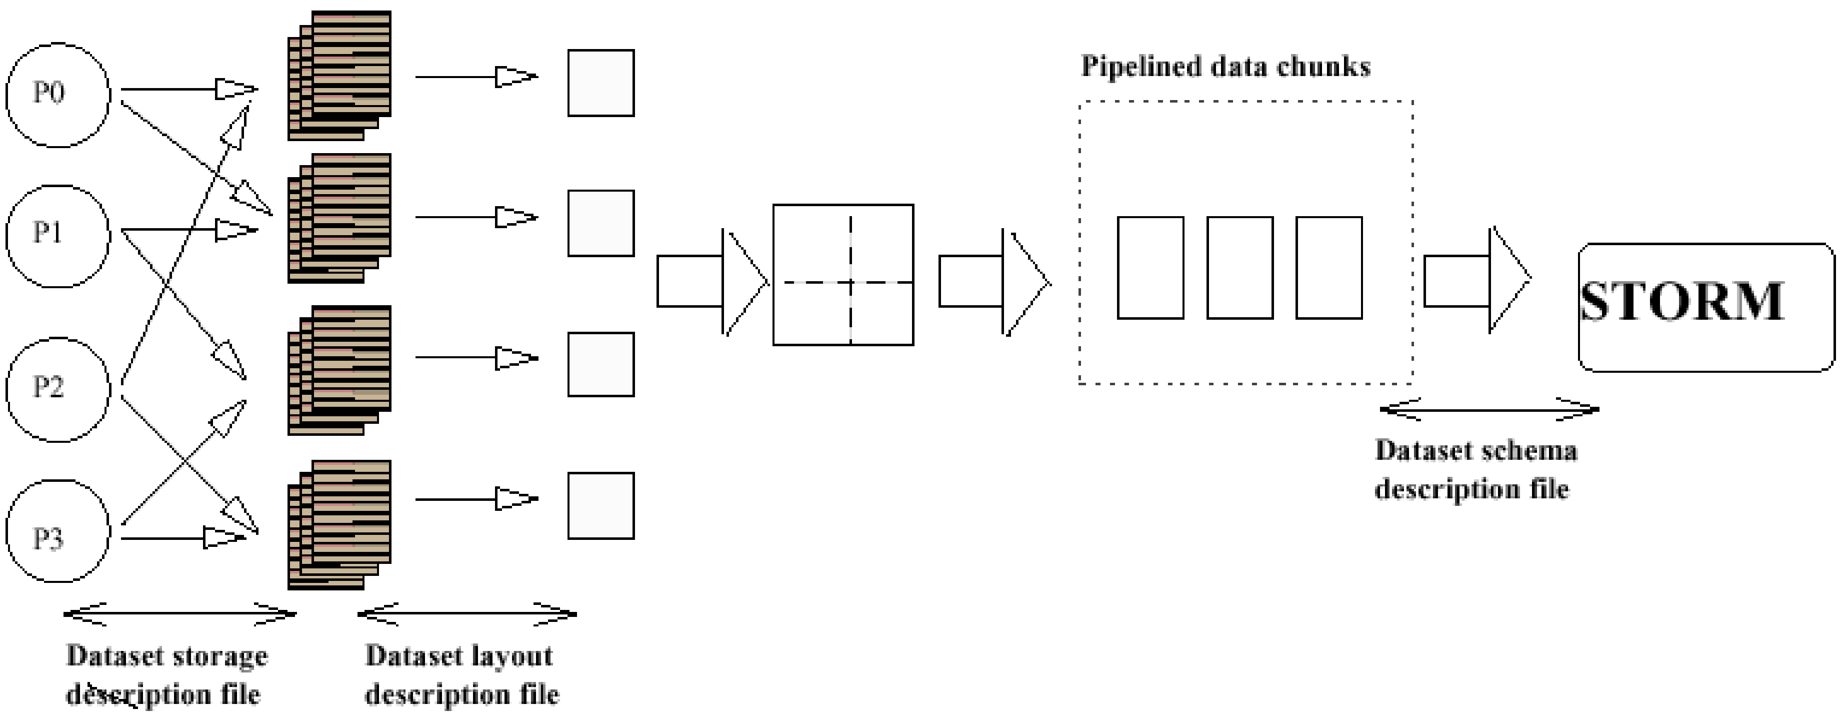

Multi-Query Optimization:Active Proxy G

- Goal: minimize the total cost of processing a series of queries by creating an *optimized access plan* for the entire sequence [Kang, Dietz, and Bhargava]
- Approach: minimize the total cost of processing a series of queries through data and computation reuse
- [IPDPS2002,SC2002,ICS02]

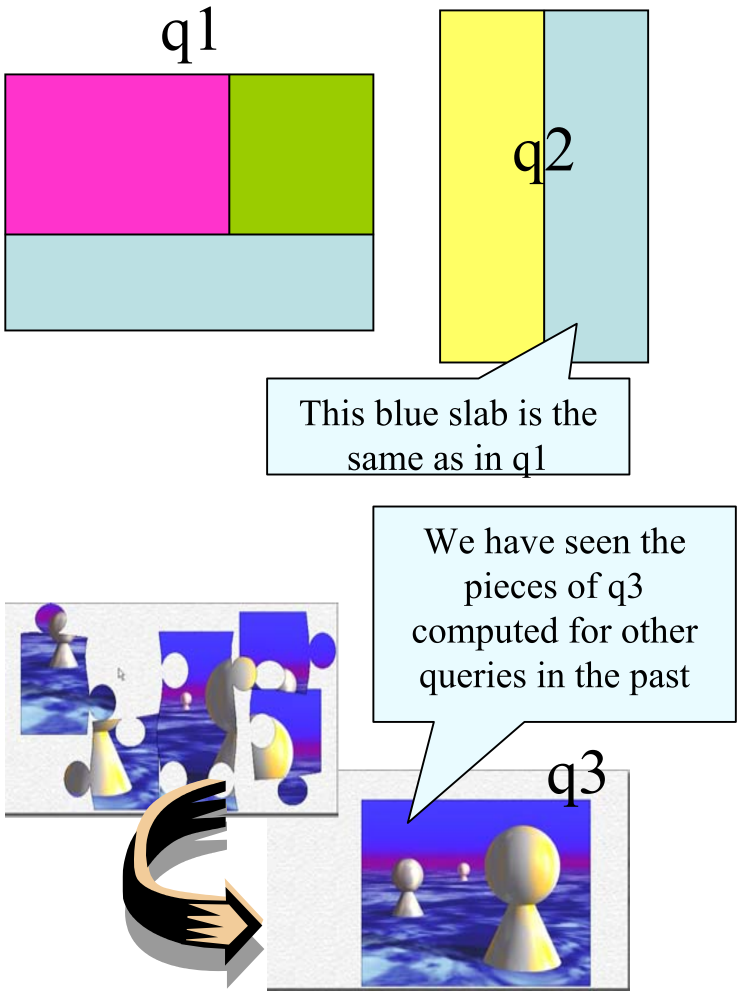

Query Optimization: query box break-down

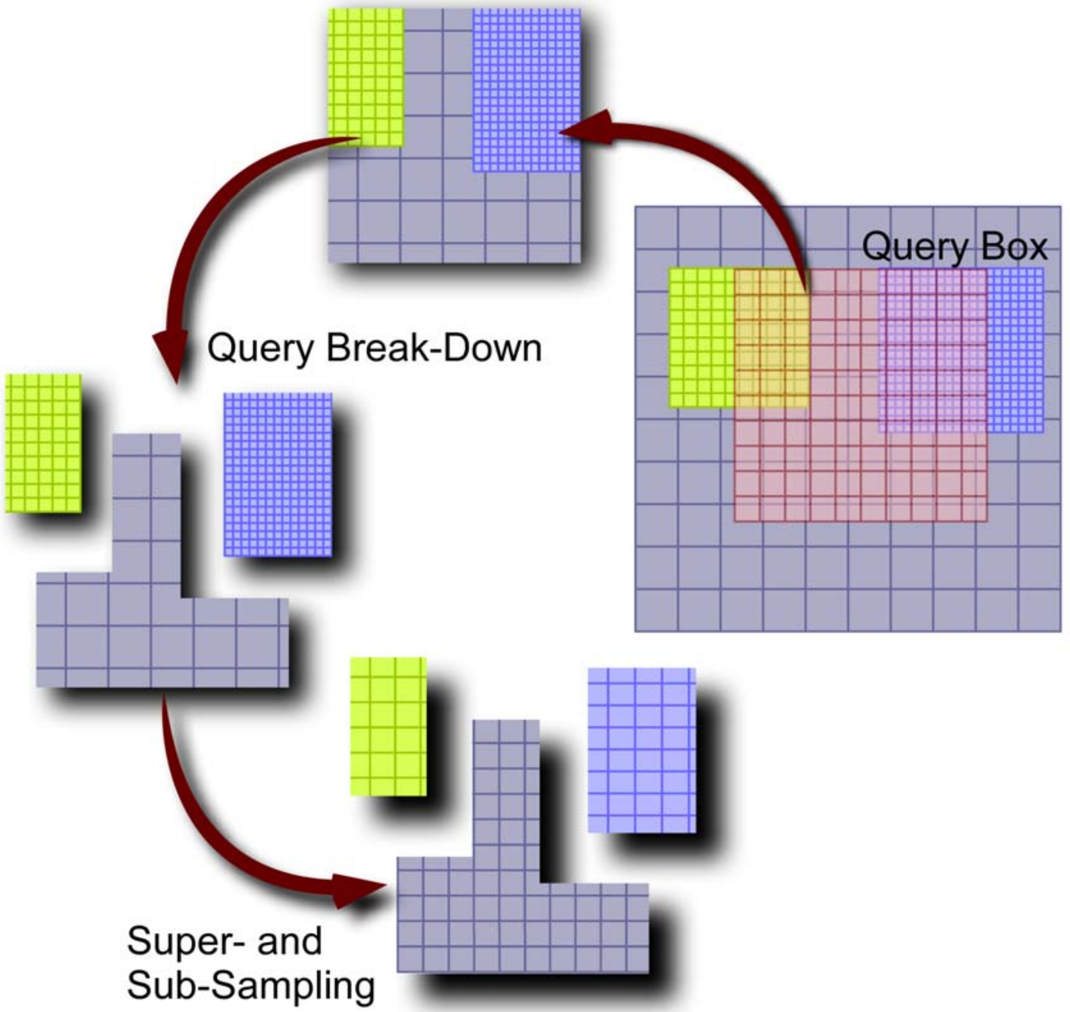

What does it buy? (Digital microscopy)

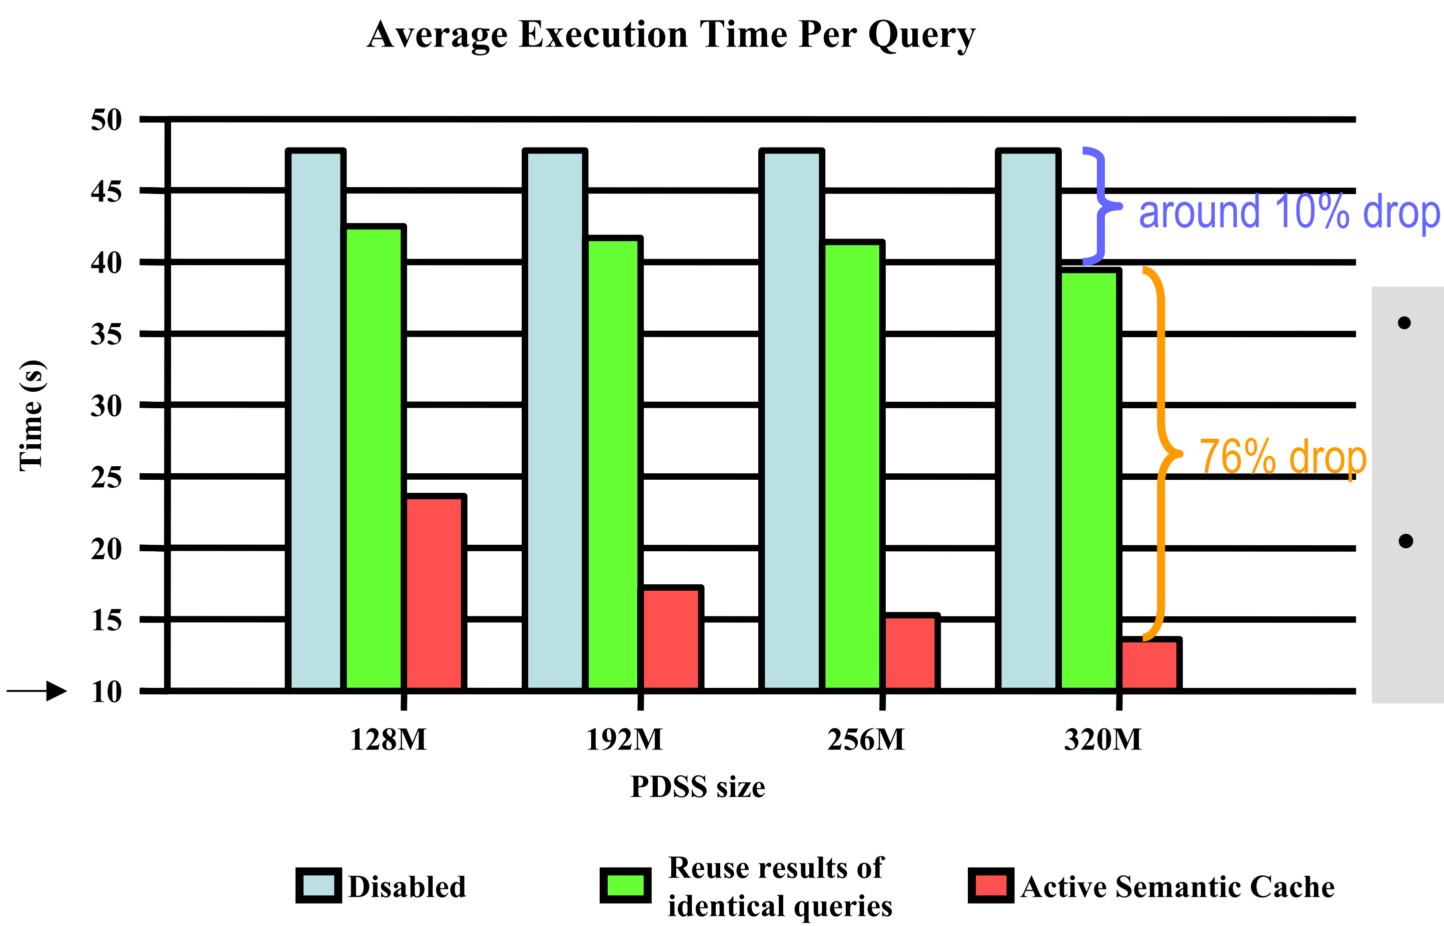

- 12 clients
  - 4 x 16 VR queries
  - 8 x 32 VM queries
- 4 processors (up to 4 queries simultaneously)

Active ProxyG: Functional Components

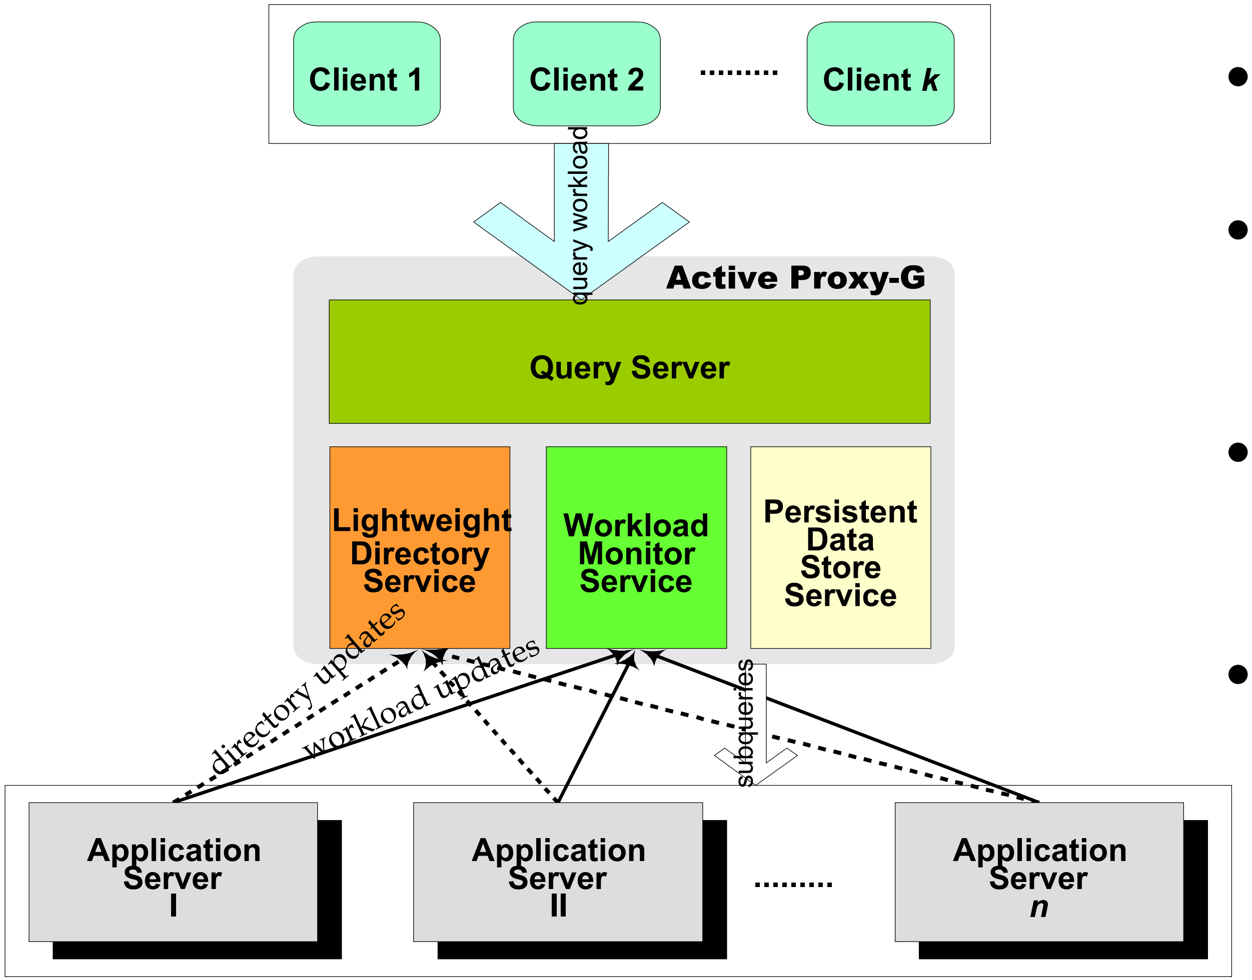

- Query Server
- Lightweight Directory Service
- Workload Monitor Service
- Persistent Data Store Service

Mobius

- Middleware system that provides support for management of metadata definitions (defined as XML schemas) and efficient storage and retrieval of data instances in a distributed environment.
- Mechanism for data driven applications to cache, share, and asynchronously communicate data in a distributed environment
- Grid based distributed, searchable, and shareable persistent storage
- Infrastructure for grid coordination language

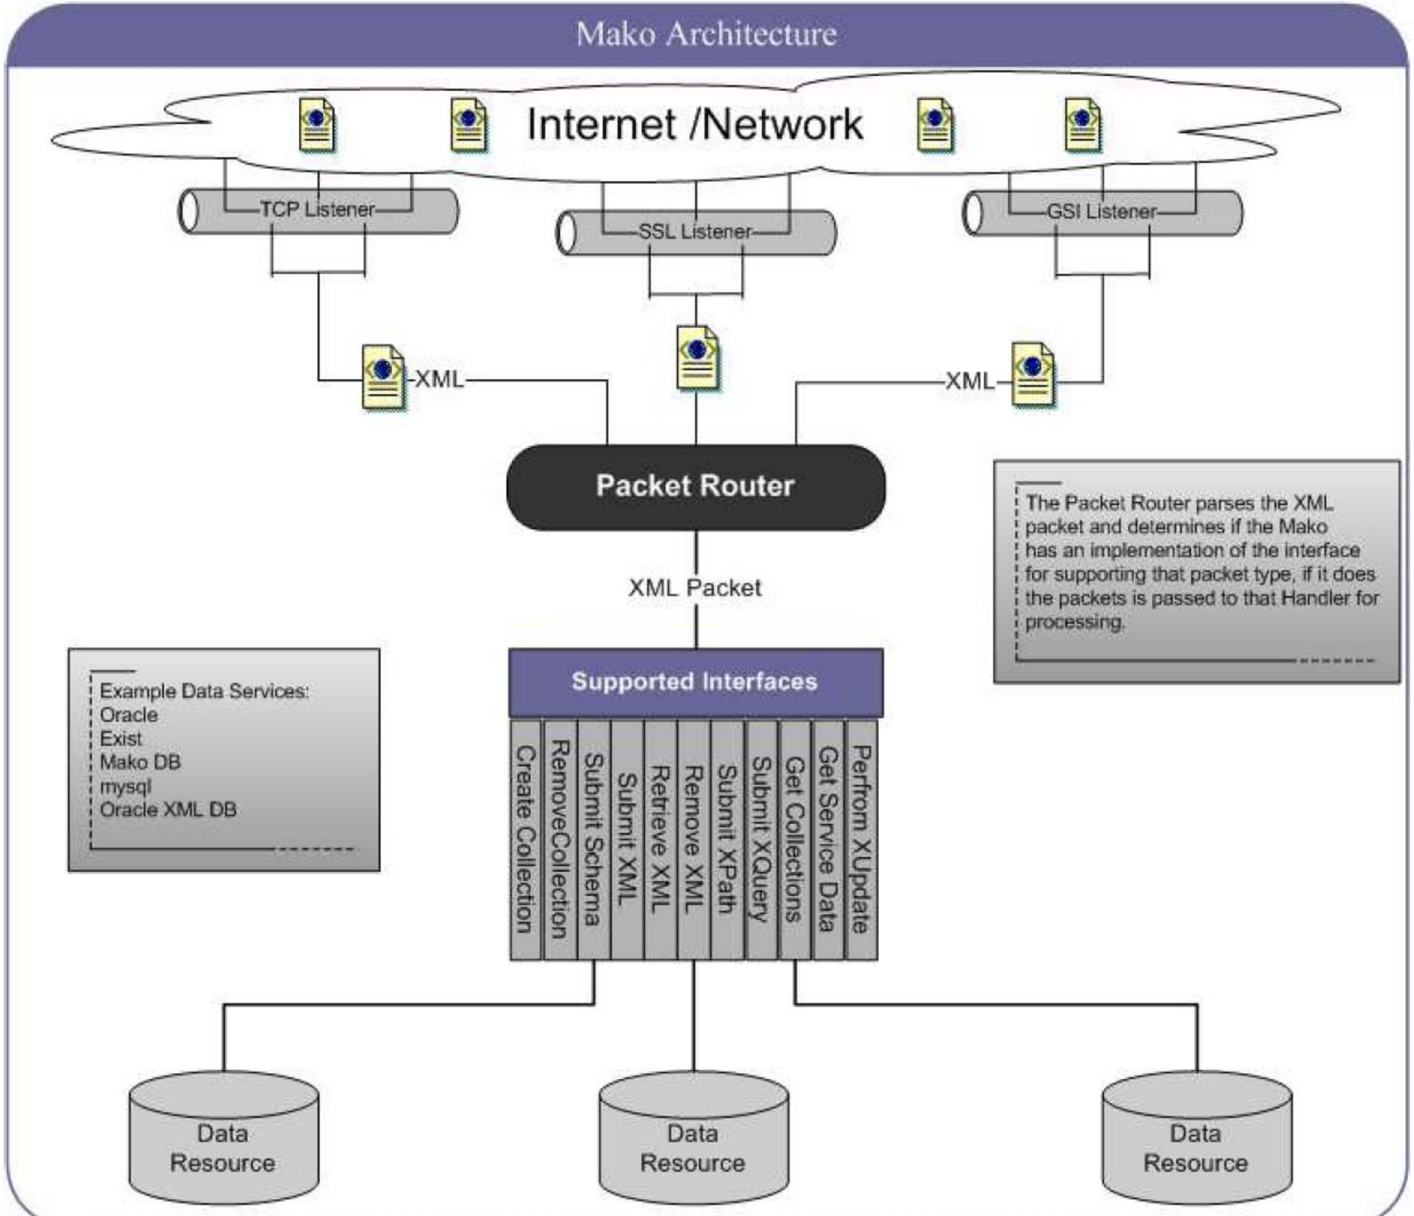

Mako Architecture
Internet /Network
TCP Listener
SSL Listener
GSI Listener
XML
XML
Packet Router
XML Packet
The Packet Router parses the XML packet and determines if the Mako has an implementation of the interface for supporting that packet type, if it does the packets is passed to that Handler for processing.
Example Data Services:
Oracle
Exist
Mako DB
mysql
Oracle XML DB
Supported Interfaces
Create Collection
RemoveCollection
Submit Schema
Submit XML
Retrieve XML
Remove XML
Submit XPath
Submit XQuery
Get Collections
Get Service Data
Perfrom XUpdate
Data Resource
Data Resource
Data Resource

Image Processing Pipeline with Checkpointing

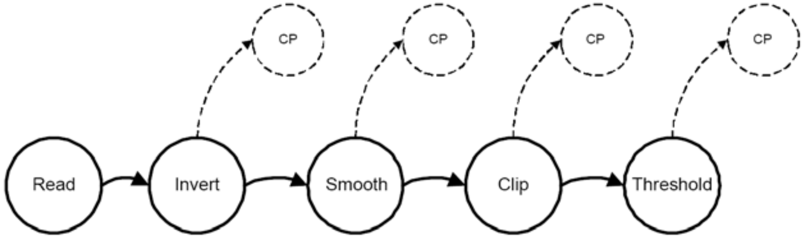

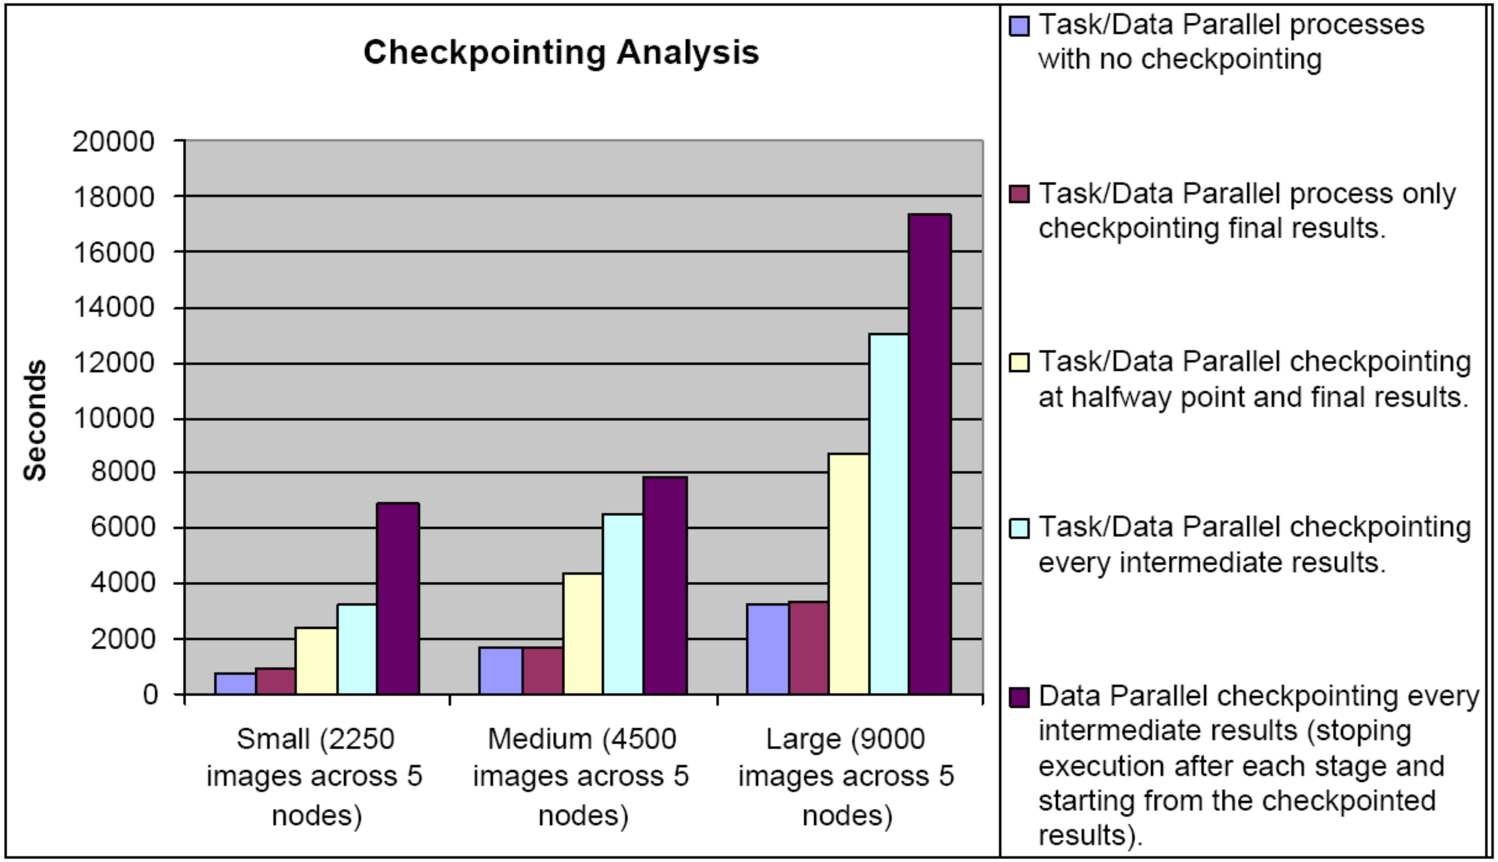

Active Projects/Funding

- National Science Foundation National Middleware Infrastructure
- National Science Foundation ITR on the Instrumented Oilfield (Dynamic Data Driven Application Systems)
- National Science Foundation NGS: An Integrated Middleware and Language/Compiler for Data Intensive Applications in Grid Environment
- Center for Grid Enabled Image Biomedical Image Analysis (NIH,NBIB, NIGMS)
- Biomedical Research Technology Transfer Partnership Award, Biomedical Informatics Synthesis Platform (State of Ohio)
- Department of Energy: Data Cutter: Software Support for Generating Data Products from Very Large Datasets
- NCI: Overcoming Barriers to Clinical Trial Accrual
- OSU Cancer Center Shared Resource

Related Work

- GGF
- Grid Middleware: Globus, Network Weather Service, GridSolve, Storage Resource Broker, CACTUS, CONDOR
- Common Component Architecture
- Query, indexing very large databases: Jim Gray – Microsoft; keyhole.com
- Close relationship to much viz work

Multiscale Laboratory Research Group

Ohio State University

- Joel Saltz
- Gagan Agrawal
- Umit Catalyurek
- Dan Cowden
- Mike Gray
- Tahsin Kurc
- Shannon Hastings
- Steve Langella
- Scott Oster
- Tony Pan
- Sivaramakrishnan (K2)
- Michael Zhang

The Ohio Supercomputer Center

- Don Stredney
- Dennis Sessanna
- Jason Bryan

University of Maryland

- Alan Sussman
- Henrique Andrade
- Christian Hansen

Center on Grid Enabled Image Processing

- Joel Saltz
- Michael Caligiuri
- Charis Eng
- Mike Knopp
- DK Panda
- Steve Qualman
- Jay Zweier