Large Irregular Datasets and the Computational Grid

Joel Saltz

University of Maryland College Park

Computer Science Department

Johns Hopkins Medical Institutions

Pathology Department

Computational grids and Irregular Multidimensional Datasets

- Spatial/multidimensional multi-scale, multi-resolution
- Applications select portions of one or more datasets
- Selection of data subset makes use of spatial index (e.g. R tree or quad tree)
- Data not used “as-is”, generally preprocessing is needed
- ***Databases to carry out spatial queries and data aggregation in irregular multidimensional datasets***

Querying Irregular Multidimensional Datasets

- Irregular datasets
  - Think of *disk based* unstructured meshes, data structures used in adaptive multiple grid calculations
    - indexed by spatial location (e.g. position on earth, position of microscope stage)
  - Spatial query used to specify iterator
    - computation carried out on data obtained from spatial query
    - computation aggregates data - resulting data product size significantly smaller than results of range query

Application Scenarios

- Ad-hoc queries or data products from satellite sensor data
- Sensor data, fluid dynamics and chemistry codes to predict condition of waterways (e.g. Chesapeake bay simulation) and to carry out reservoir simulation
- Browse or analyze (multiresolution) digitized slides from high power light or electron microscopy (1-50Gbytes per digitized slide - 1000’s of slides per day per hospital)
- Predict materials properties using electron microscope computerized tomography sensor data
- Post-processing, analysis and exploration of data generated by large scientific simulations

Processing Remotely Sensed Data

NOAA Tiros-N
w/ AVHRR sensor

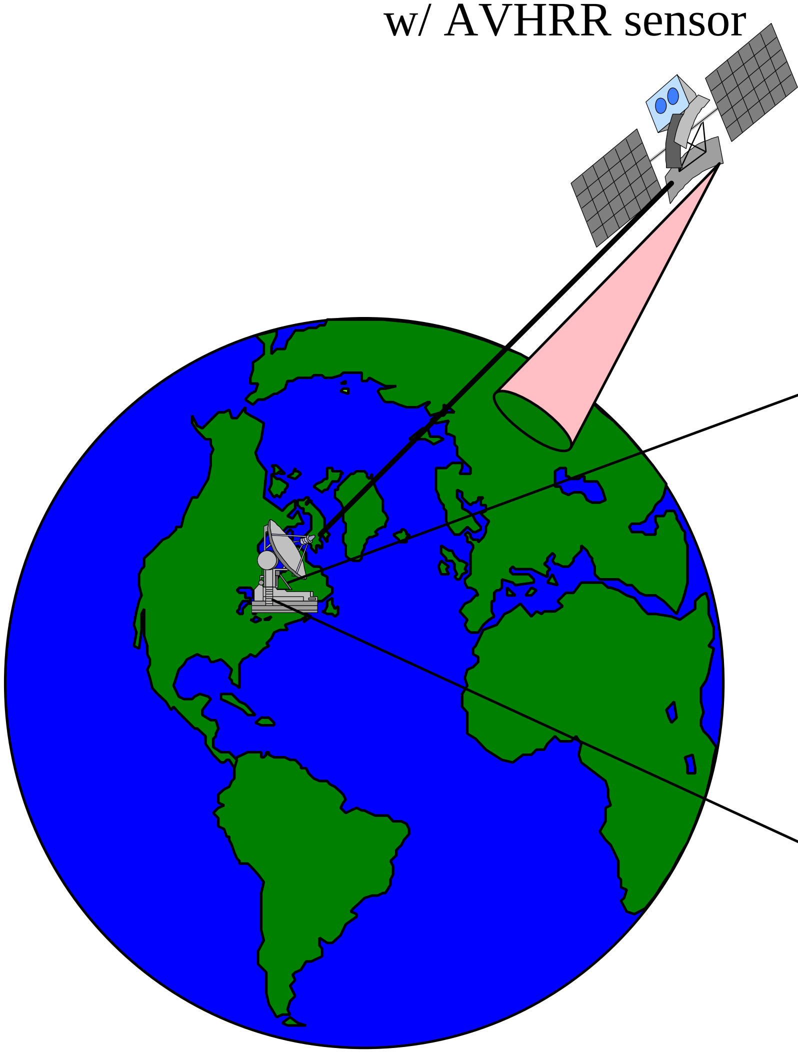

AVHRR Level 1 Data

- As the TIROS-N satellite orbits, the *Advanced Very High Resolution Radiometer* (AVHRR) sensor scans perpendicular to the satellite's track.
- At regular intervals along a scan line measurements are gathered to form an *instantaneous field of view* (IFOV).
- Scan lines are aggregated into Level 1 data sets.

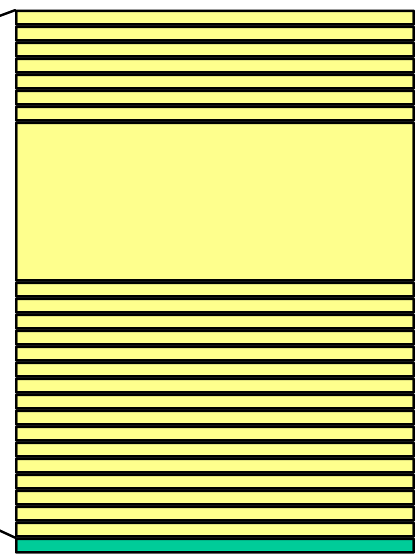

A single file of *Global Area Coverage* (GAC) data represents:

- ~one full earth orbit.
- ~110 minutes.
- ~40 megabytes.
- ~15,000 scan lines.

One scan line is 409 IFOV's

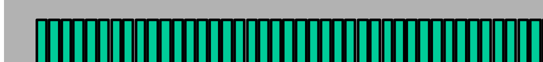

Spatial Irregularity

AVHRR Level 1B NOAA-7 Satellite 16x16 IFOV blocks.

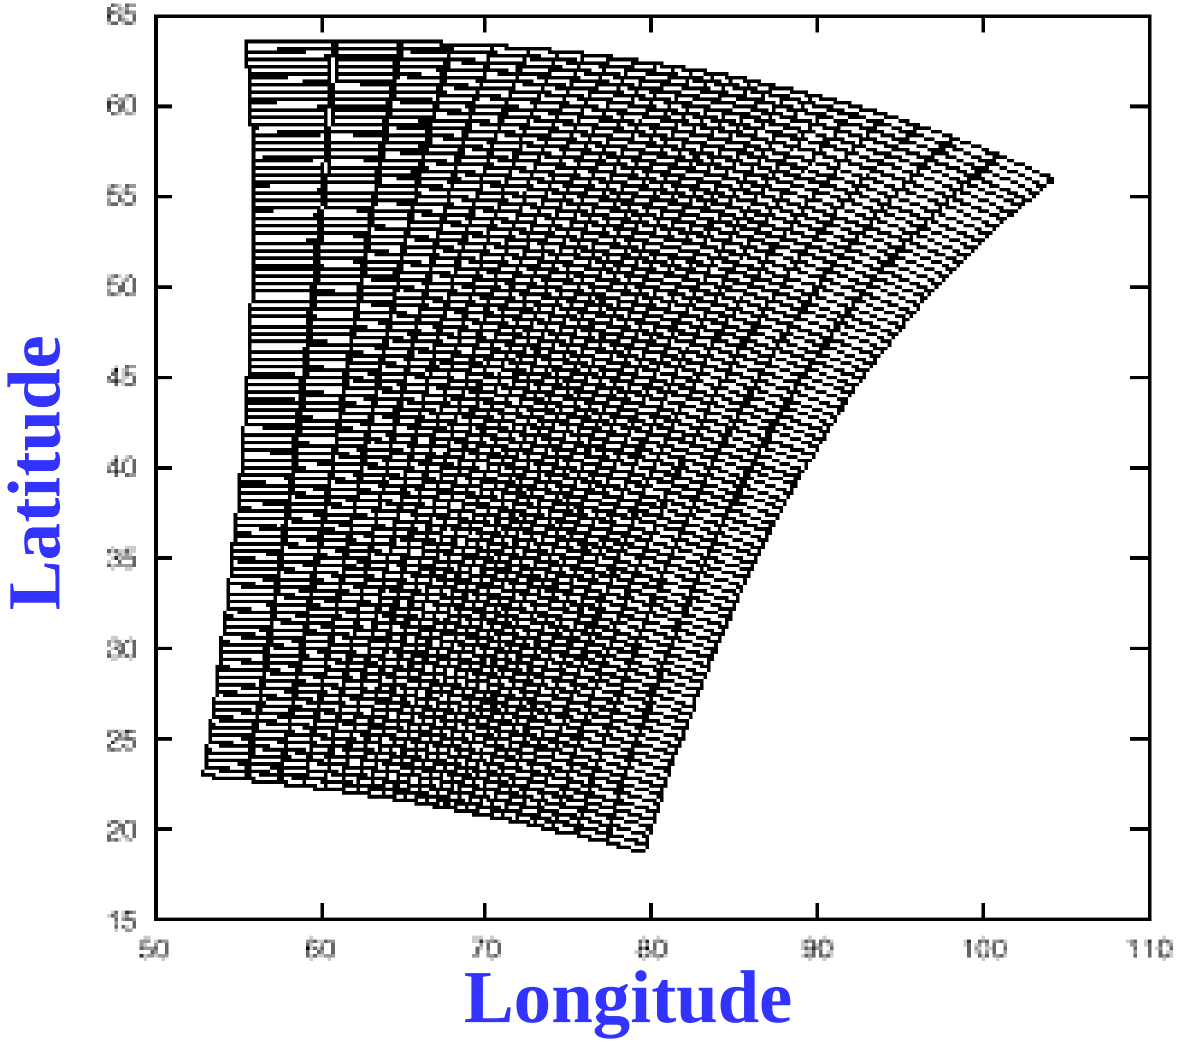

Multiscale Physics Based Simulation of Fluid Flow for Energy and Environment

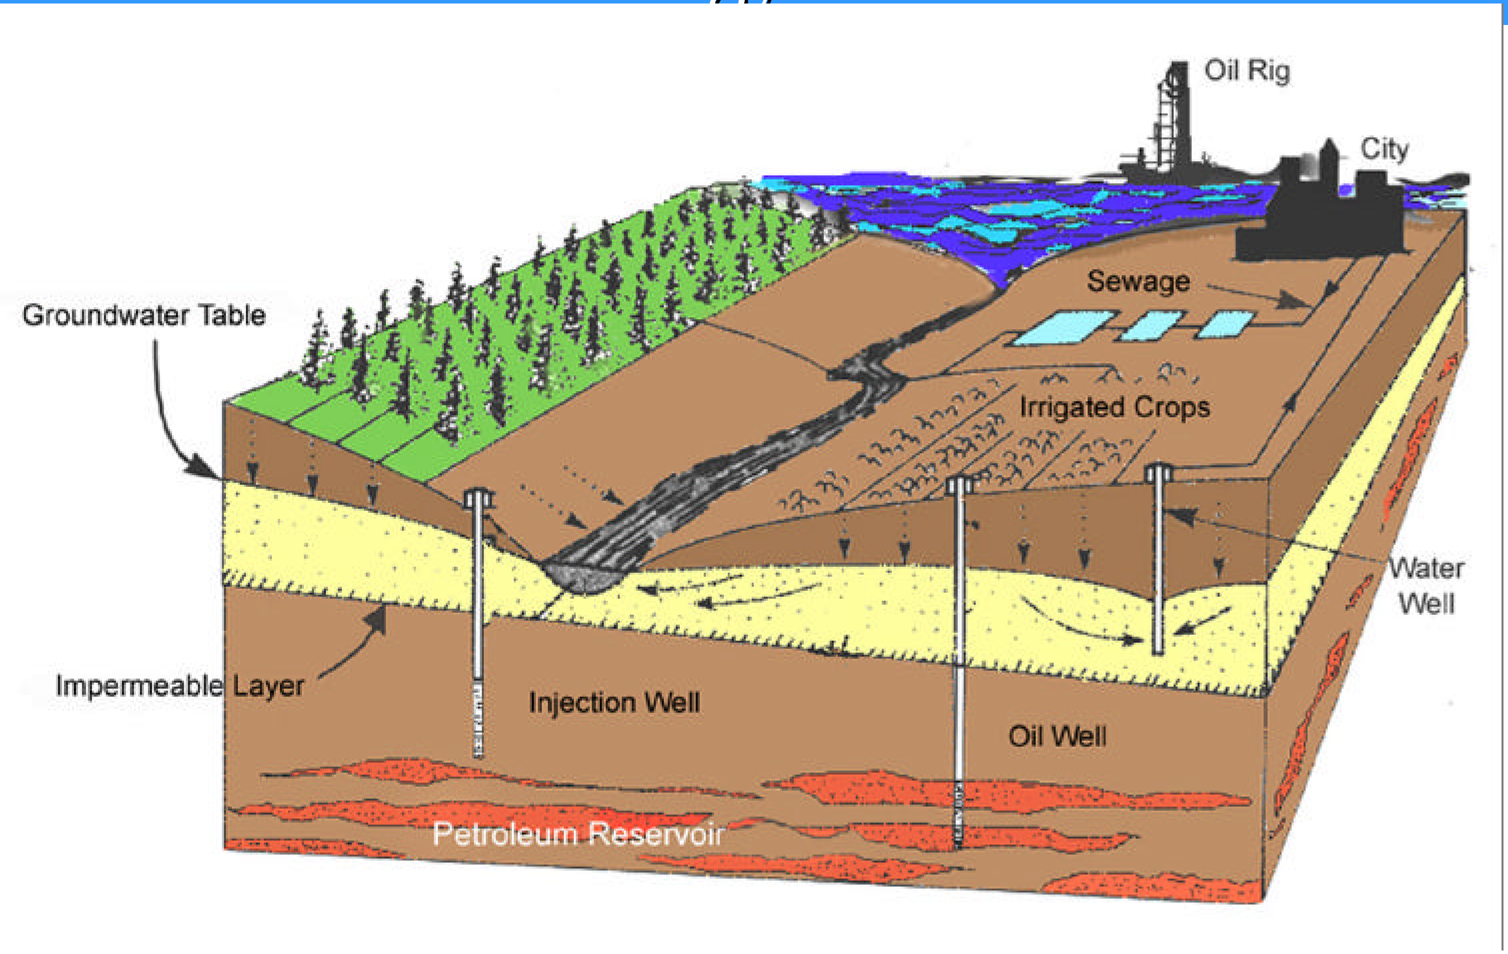

The Tyranny of Scale

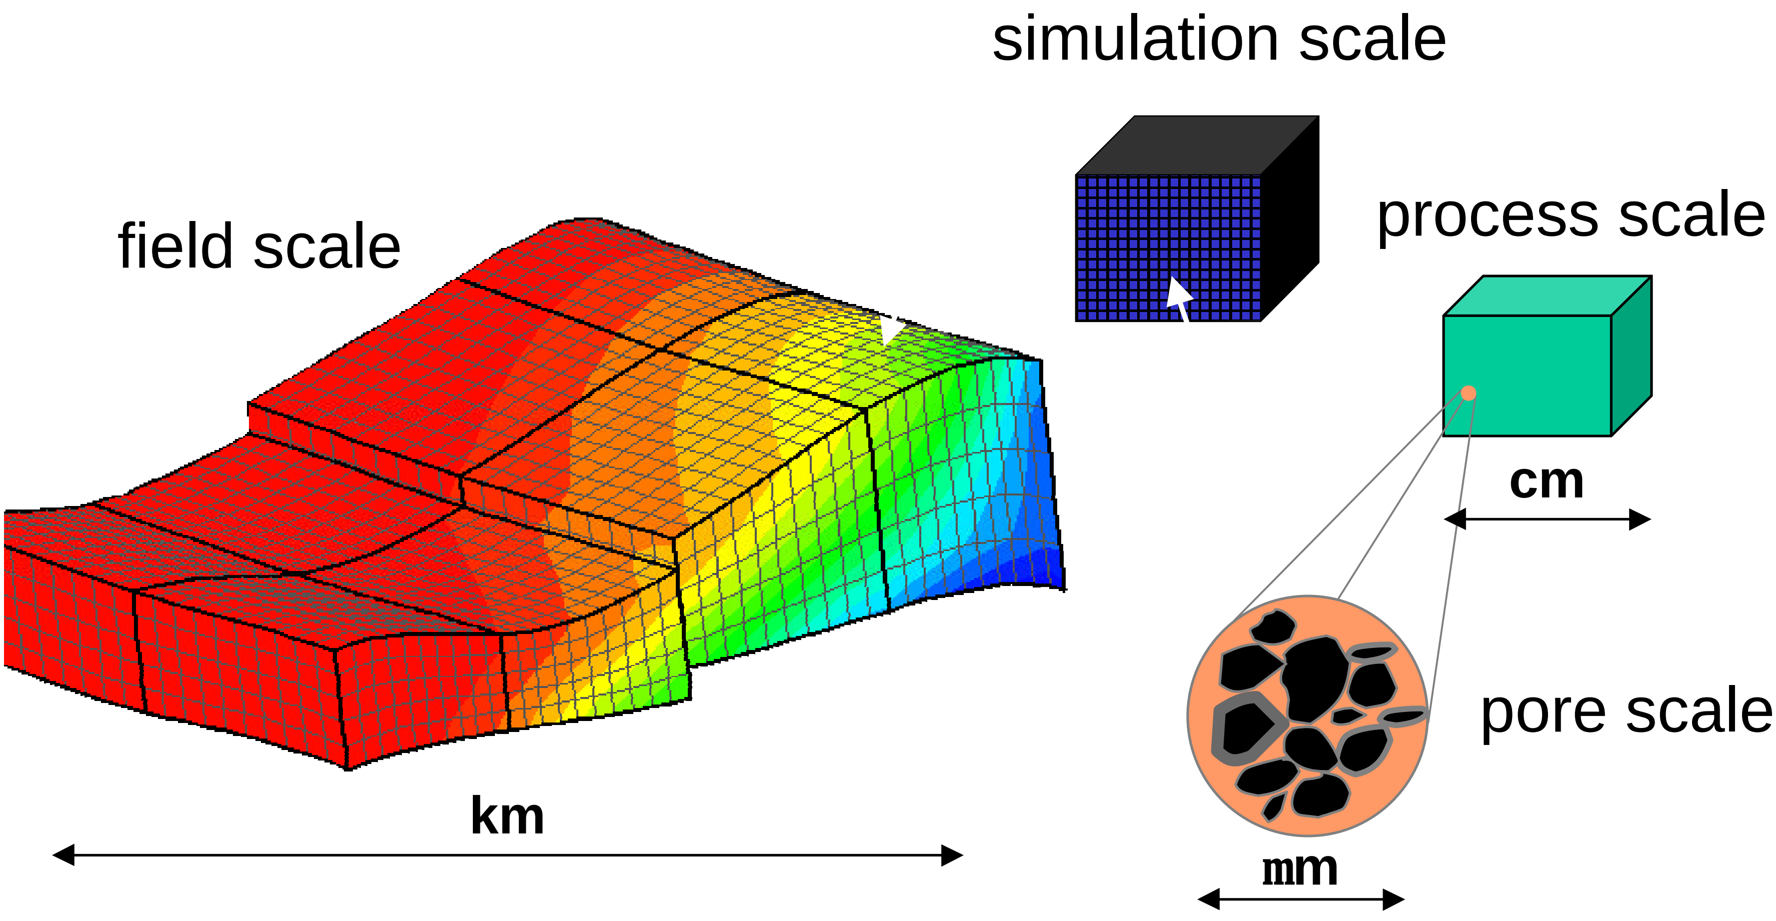

Bioremediation Simulation

abiotic reactions compete with microbes, reduce extent of biodegradation

Microbe colonies (magenta)

Dissolved NAPL (blue)

Mineral oxidation products (green)

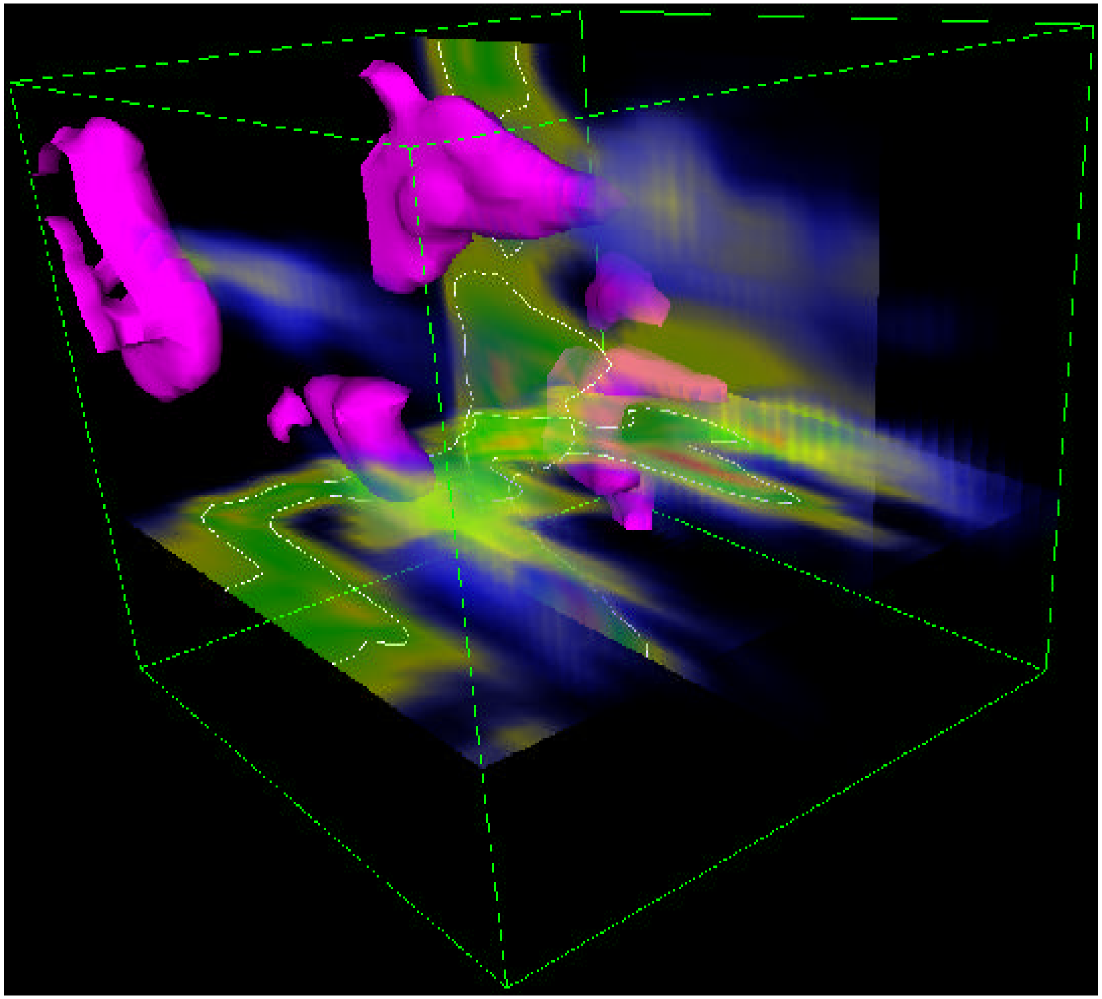

Multiblock Methods Lead to Irregular Datasets

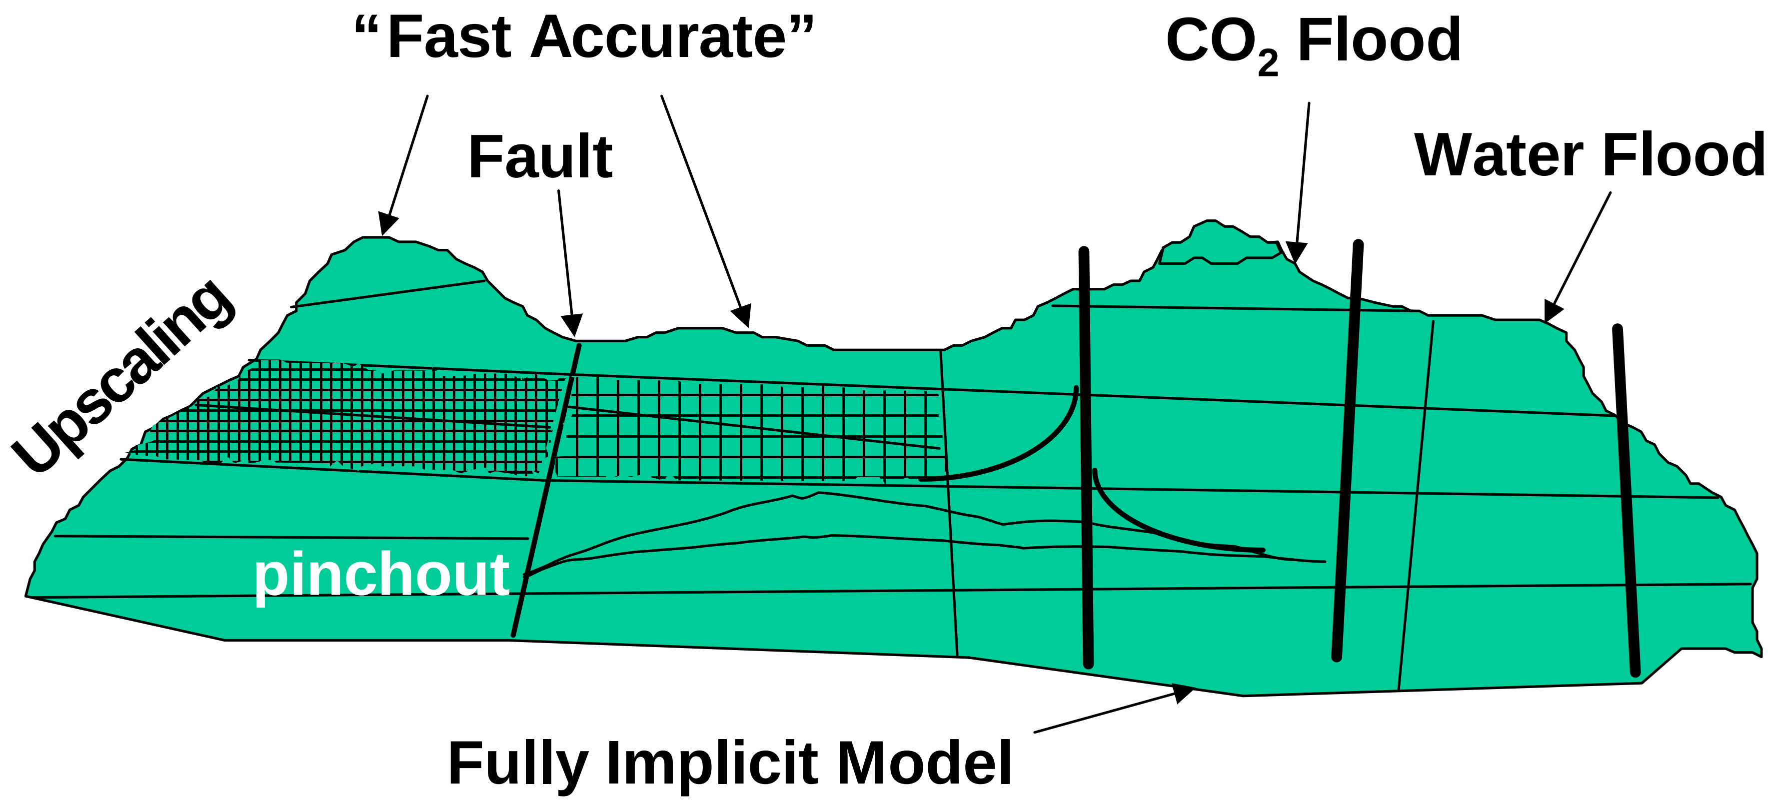

Virtual Microscope Client

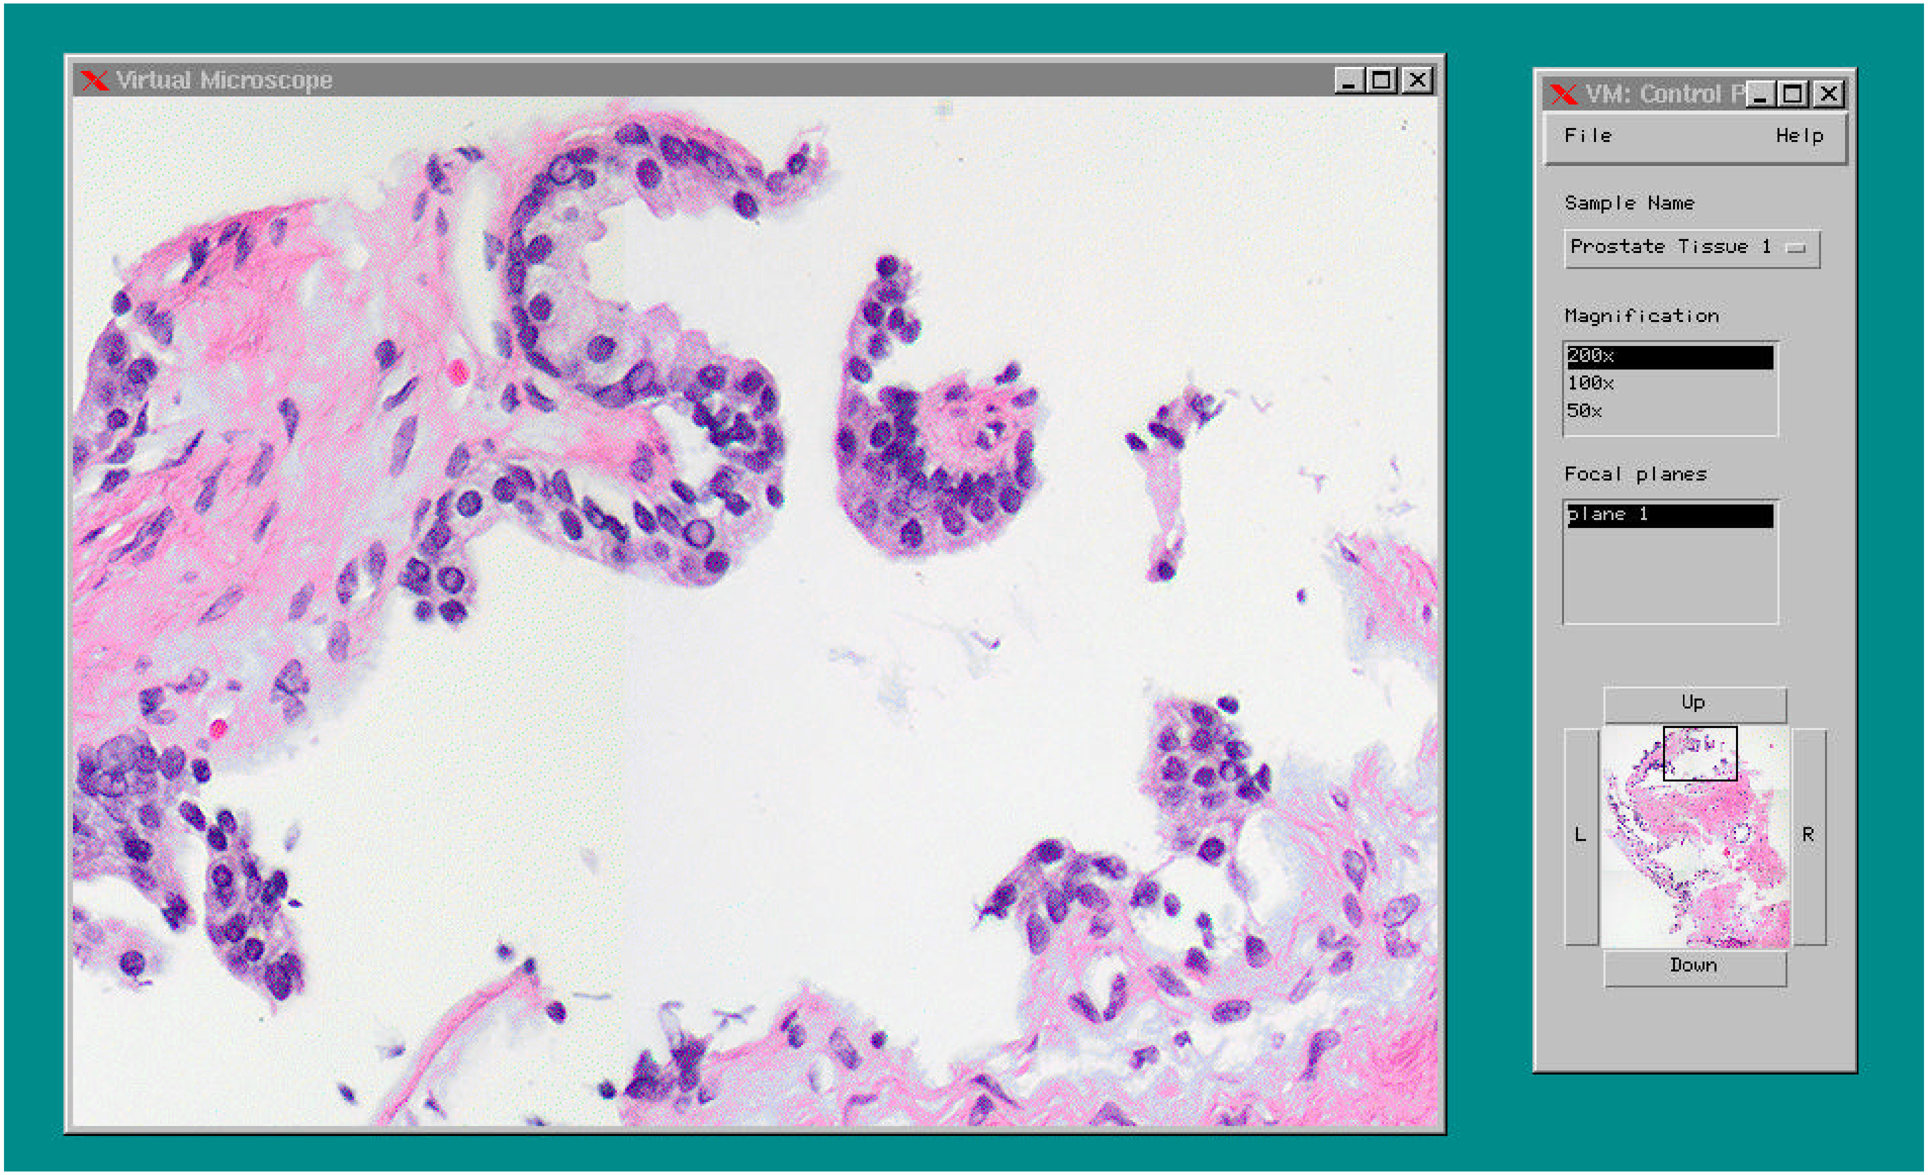

Coupling of Different Models
(Coupling of Flow Codes with Environmental Quality Codes)

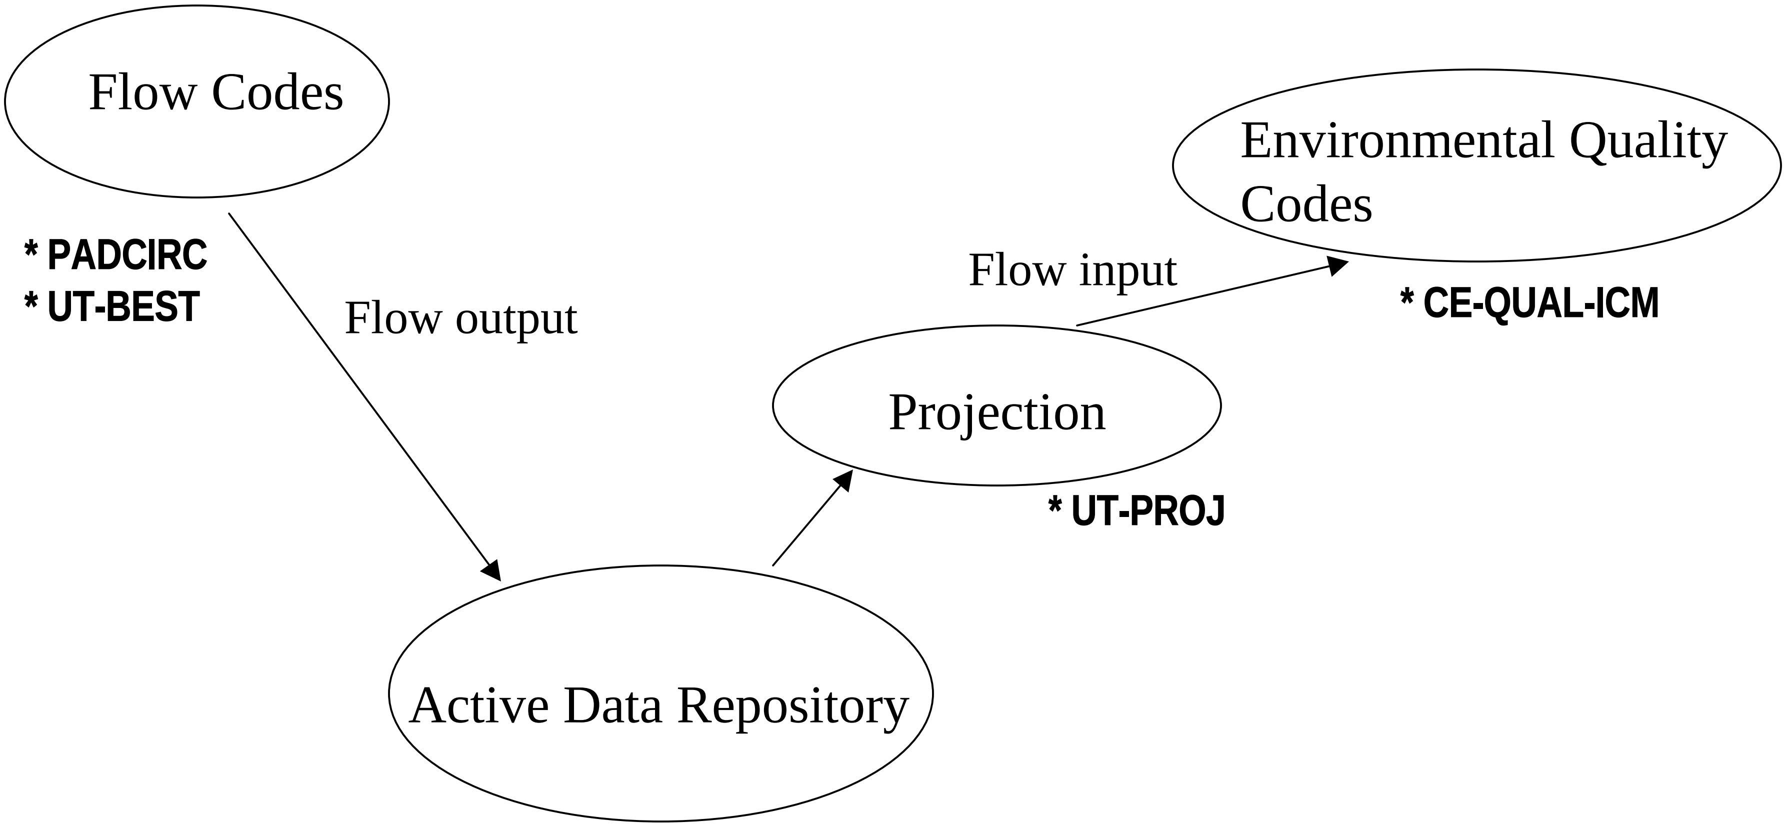

* *Storage, retrieval, processing of multiple datasets from different flow codes*

Active Data Repository and MetaChaos

- Active Data Repository -- Parallel database infrastructure to query and preprocess multi-scale multiresolution datasets
  - Projections, interpolations, combining data from different timesteps, single value of new variable from multiple existing variables
  - Performance prediction techniques for database configuration, query optimization, software and hardware design
- MetaChaos -- Tools to couple parallel databases, parallel and distributed application programs

Typical Query

Output grid onto
which a projection
is carried out

Specify portion of raw
sensor data corresponding
to some search criterion

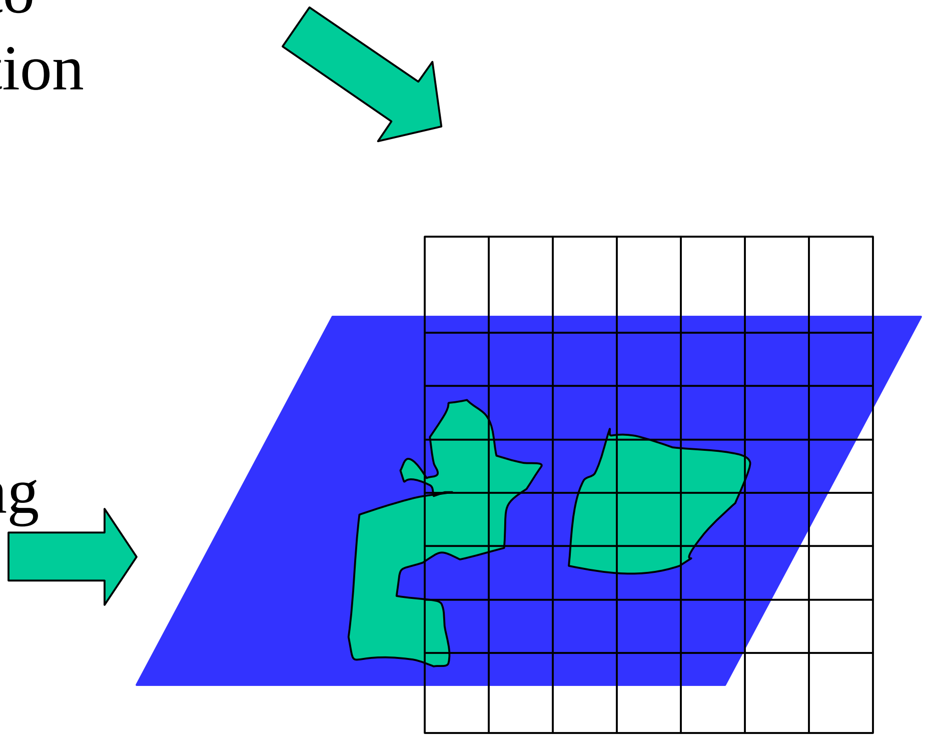

Active Data Repository Design Objectives

- *Support optimized associative access and processing of multiresolution and irregular persistent data structures*
- Integrate and overlap a wide range of user-defined operations, in particular, order-independent operations with the basic retrieval functions
- Targets *parallel and distributed architectures* that have been configured to support high I/O rates
- Applications -- Titan: Satellite sensor data; Virtual Microscope Server, Bay and Estuary Simulation

Active Data Repository, MetaChaos Grid Scenarios

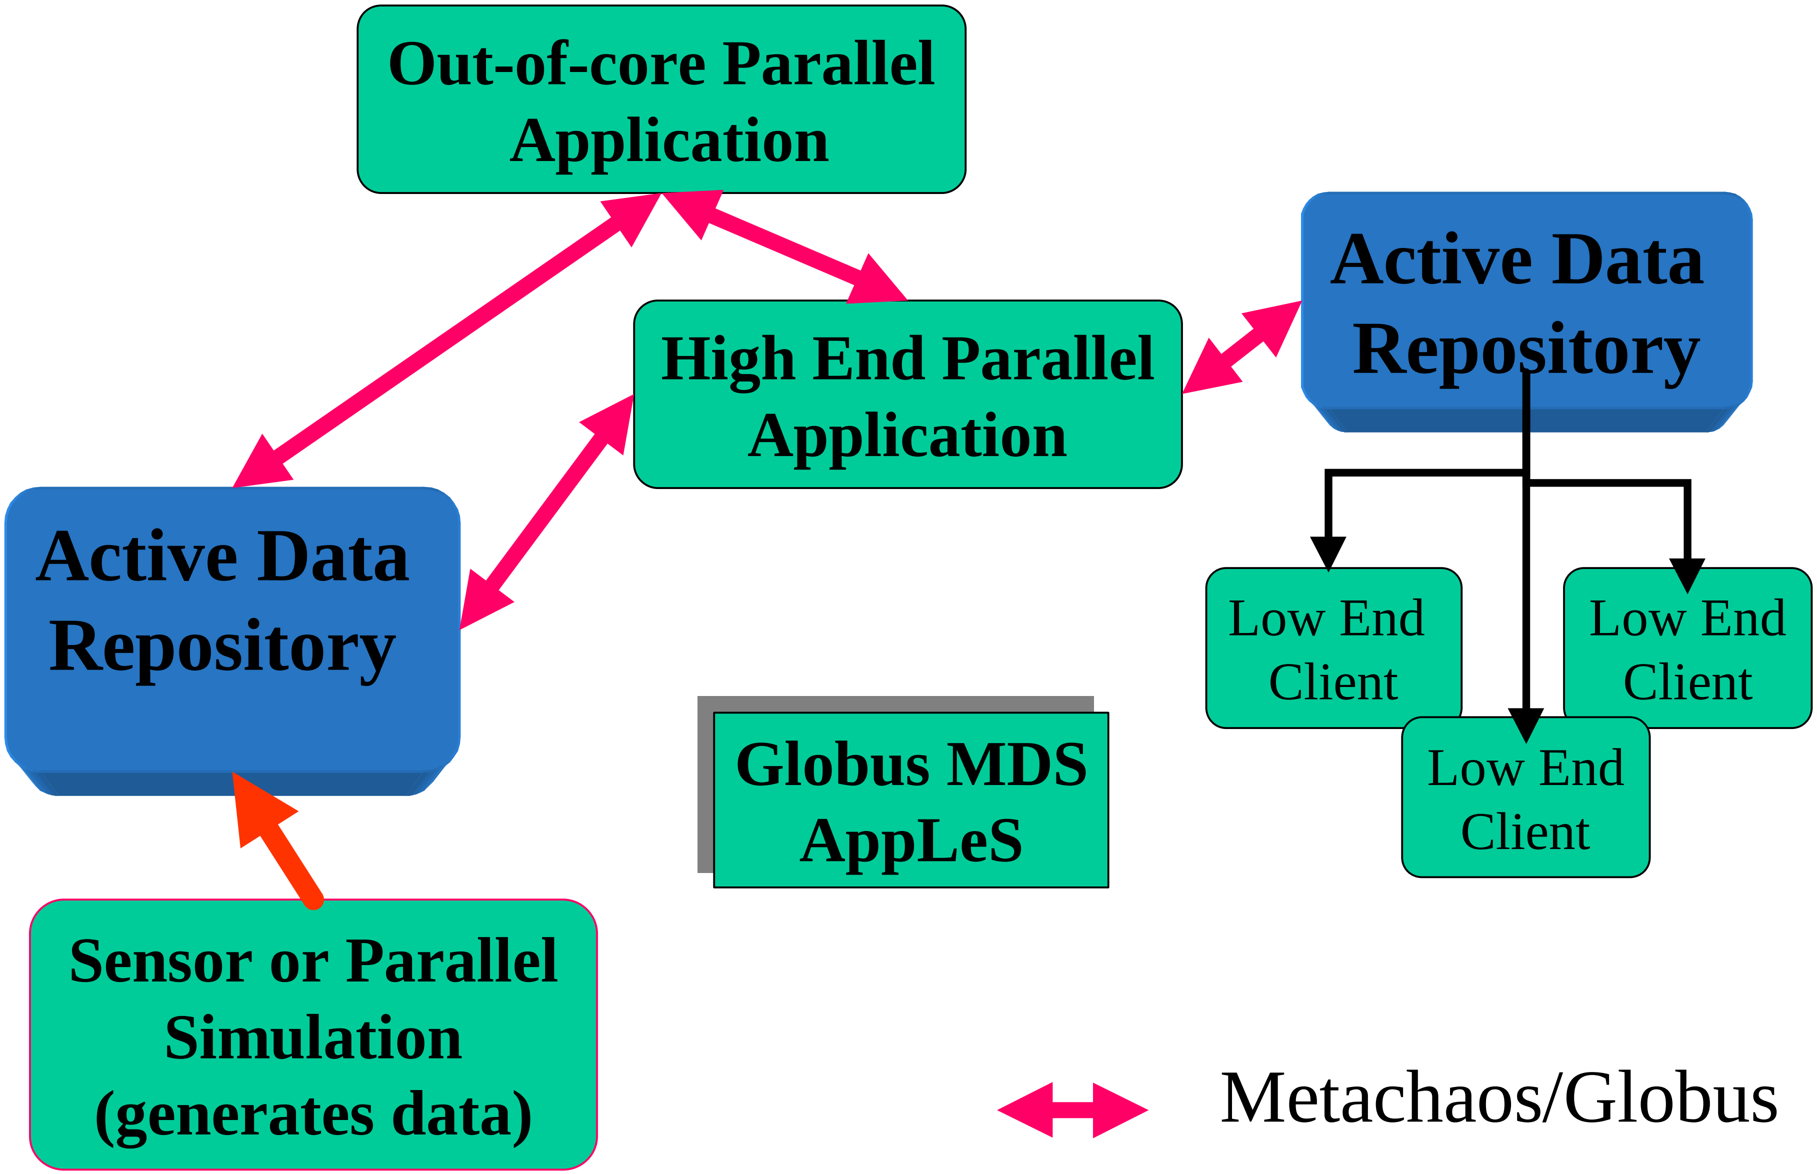

Grid Scenarios

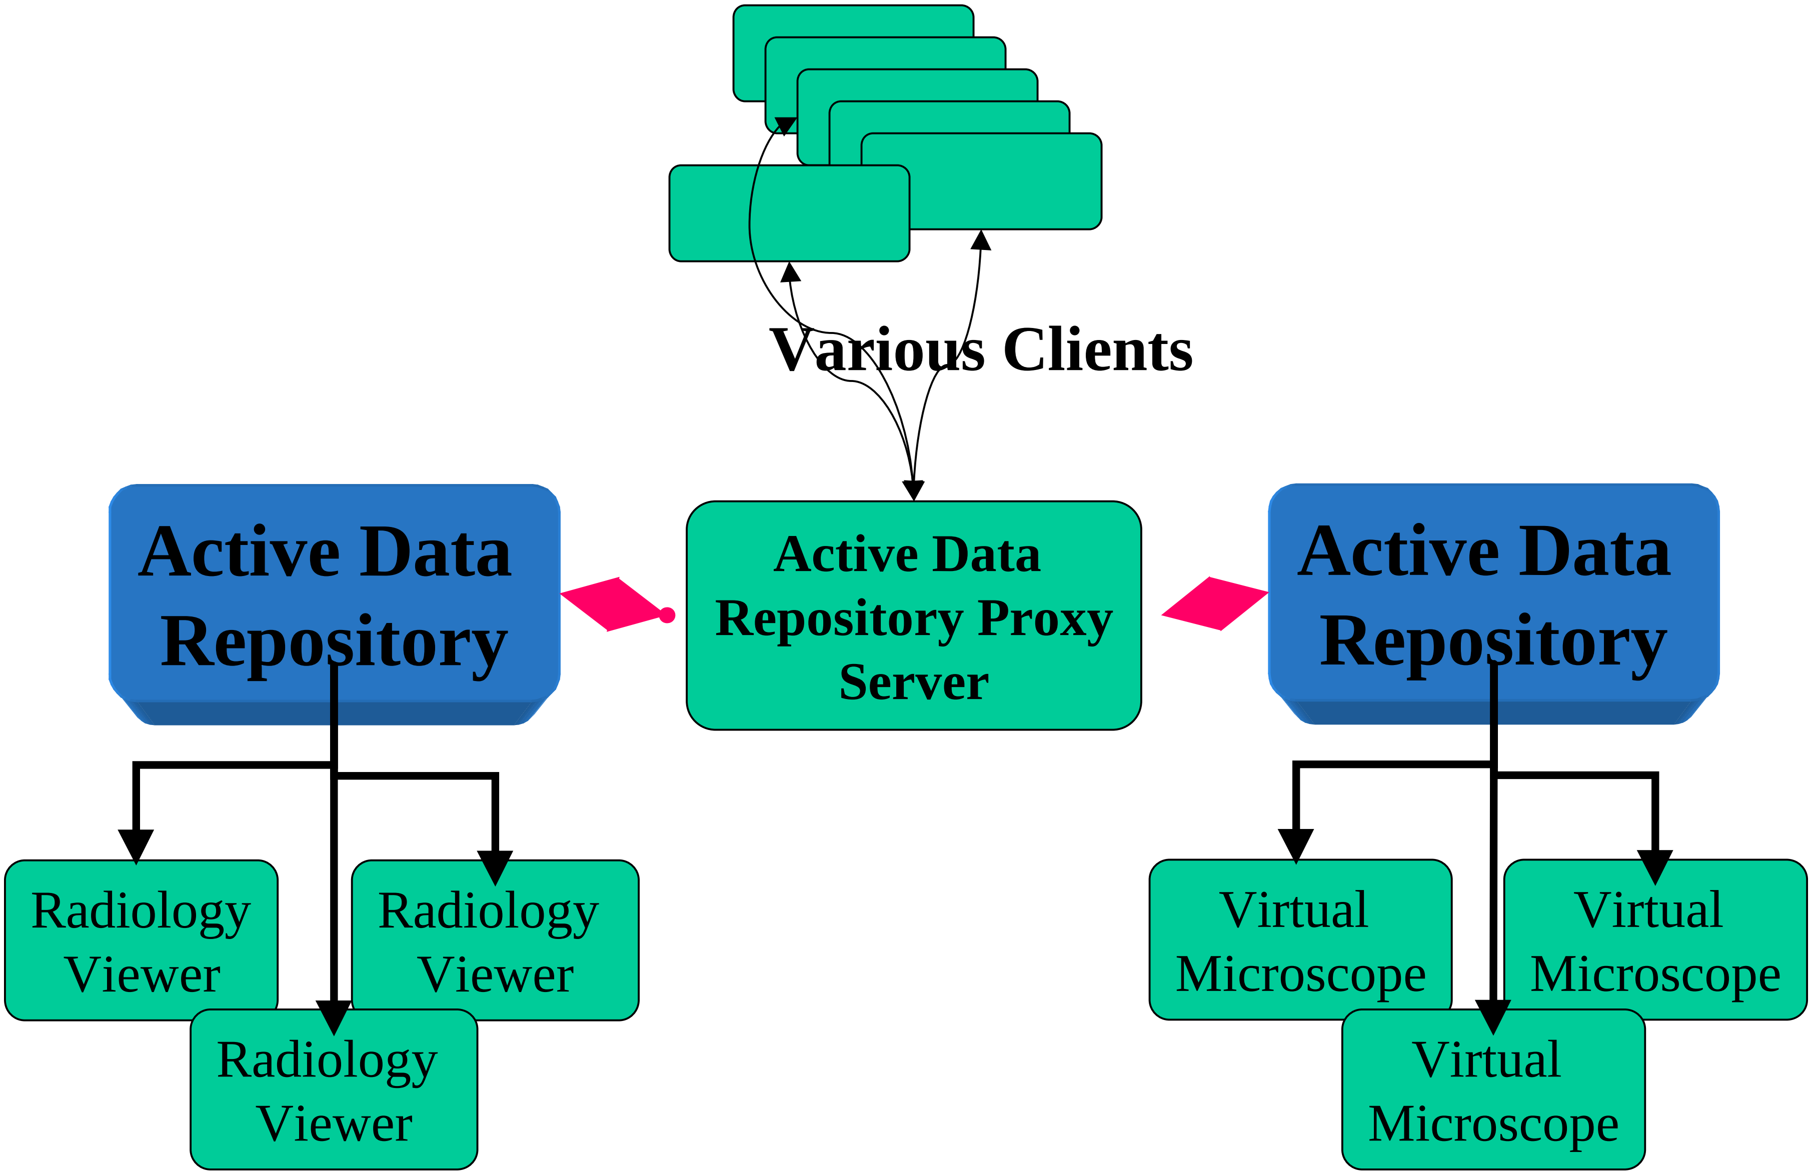

Architecture of Active Data Repository

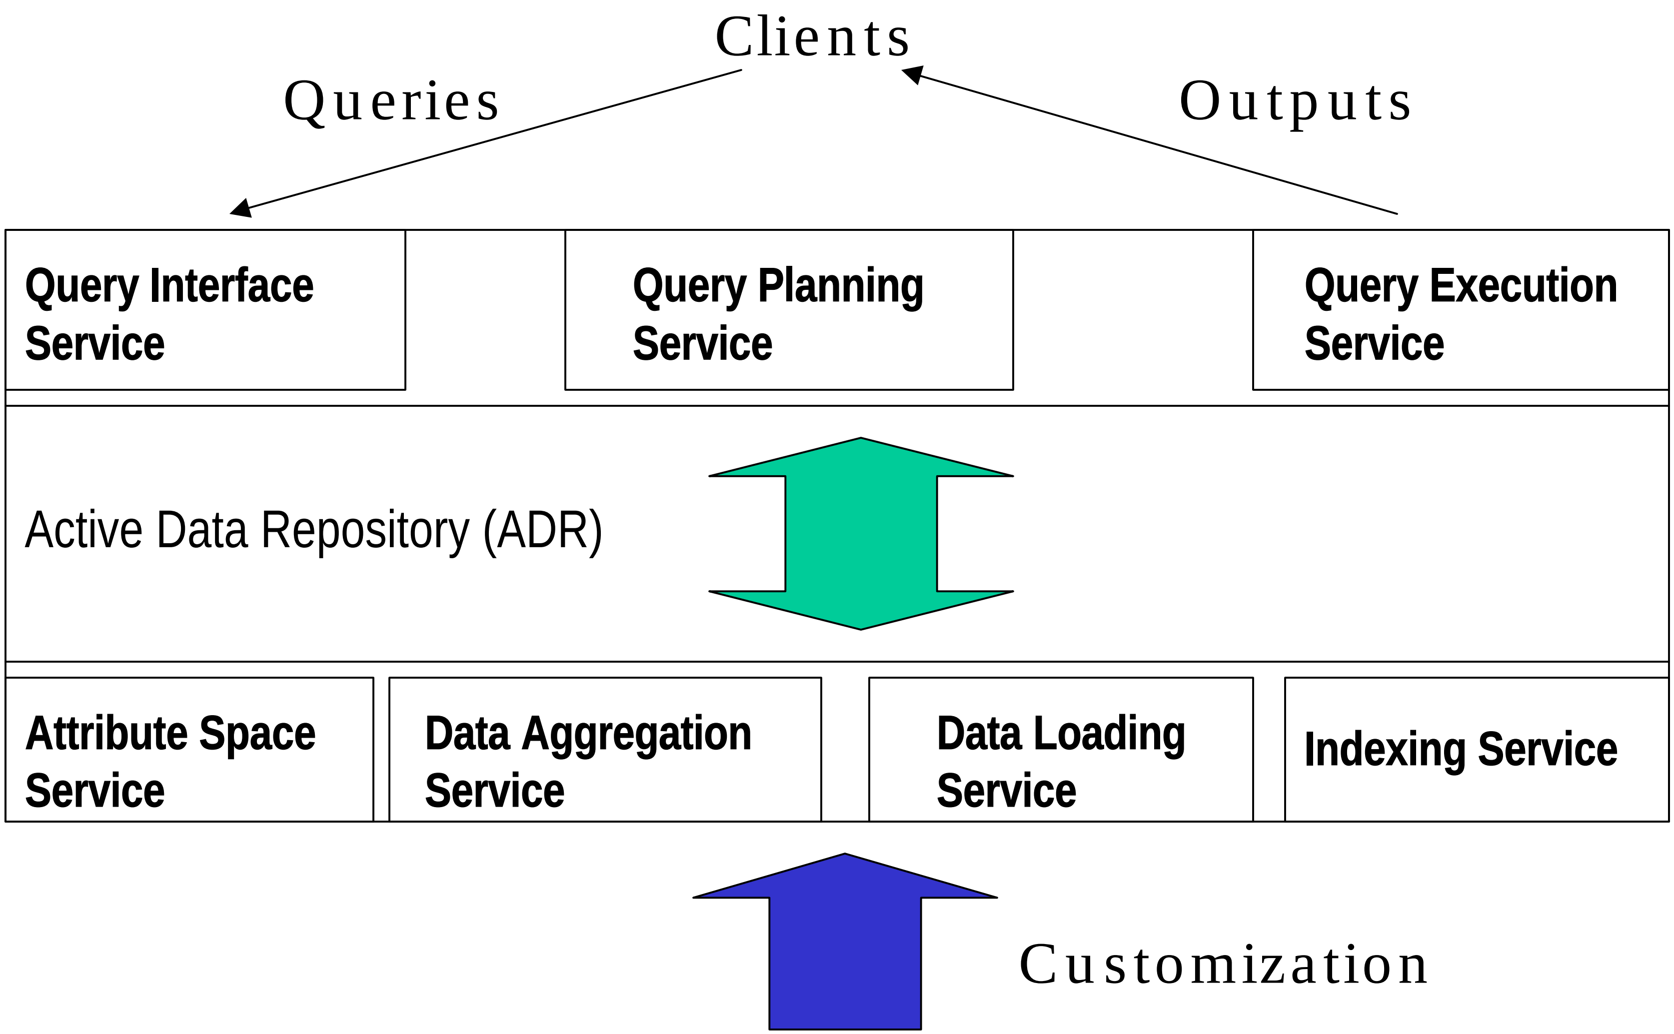

Water Contamination Studies

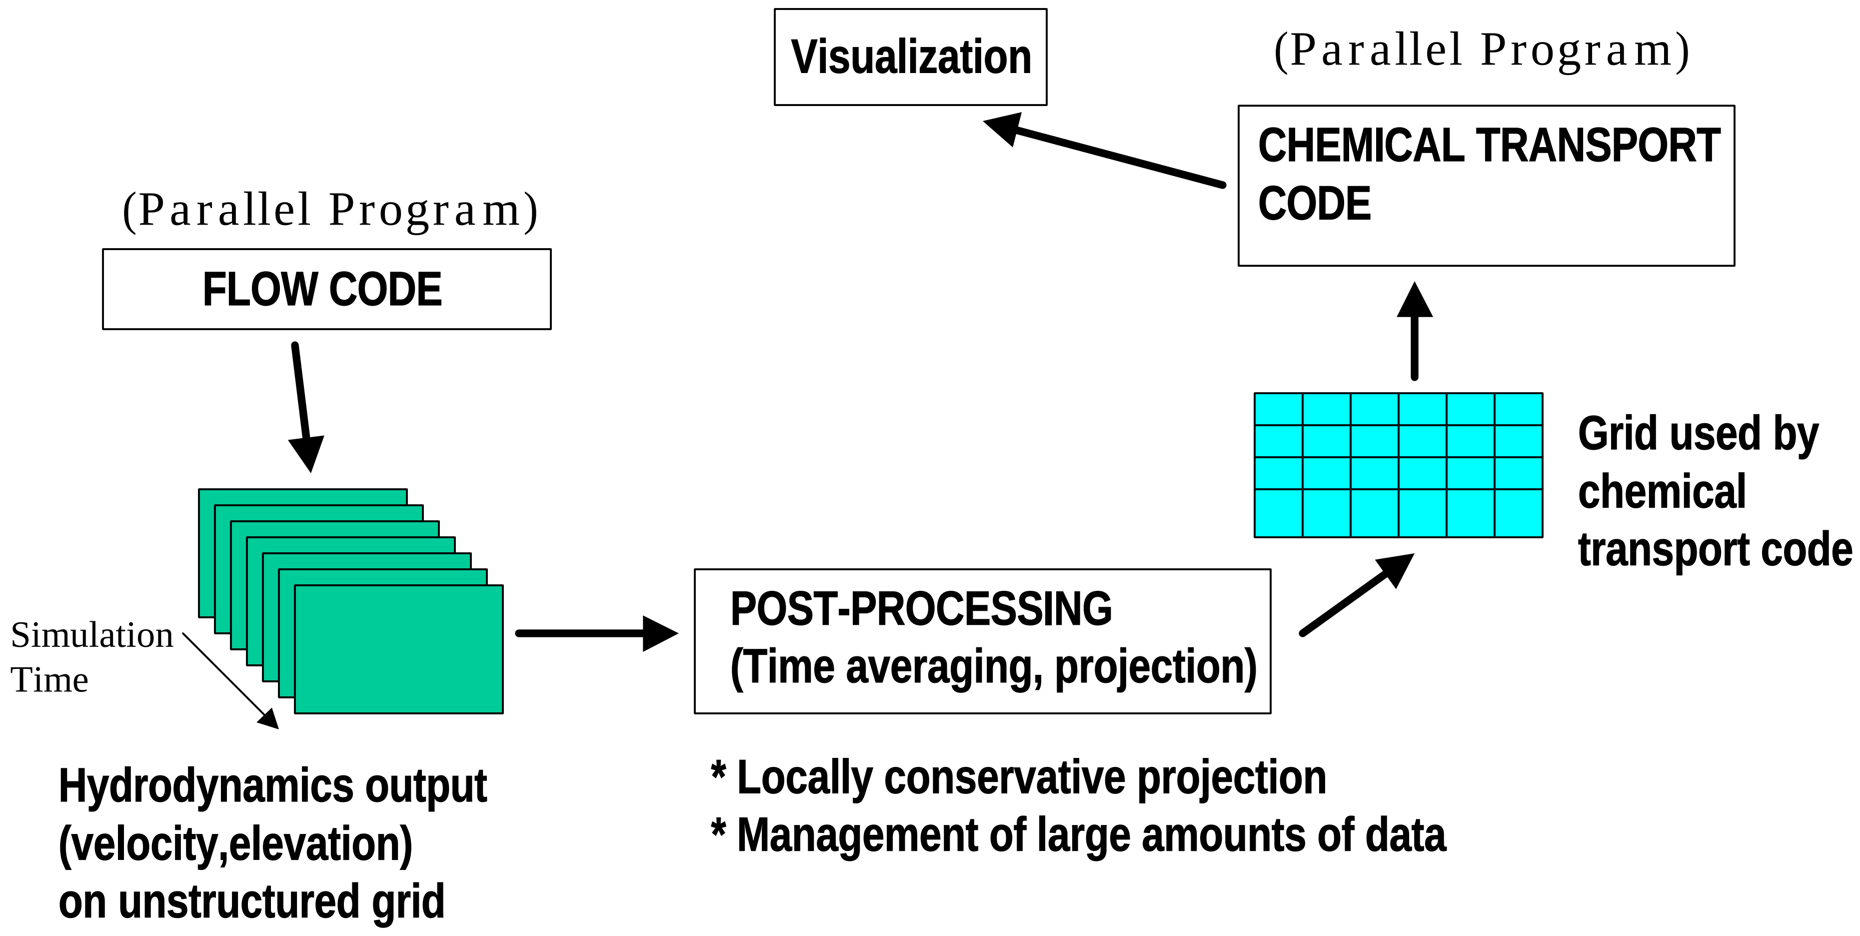

Water Contamination Studies

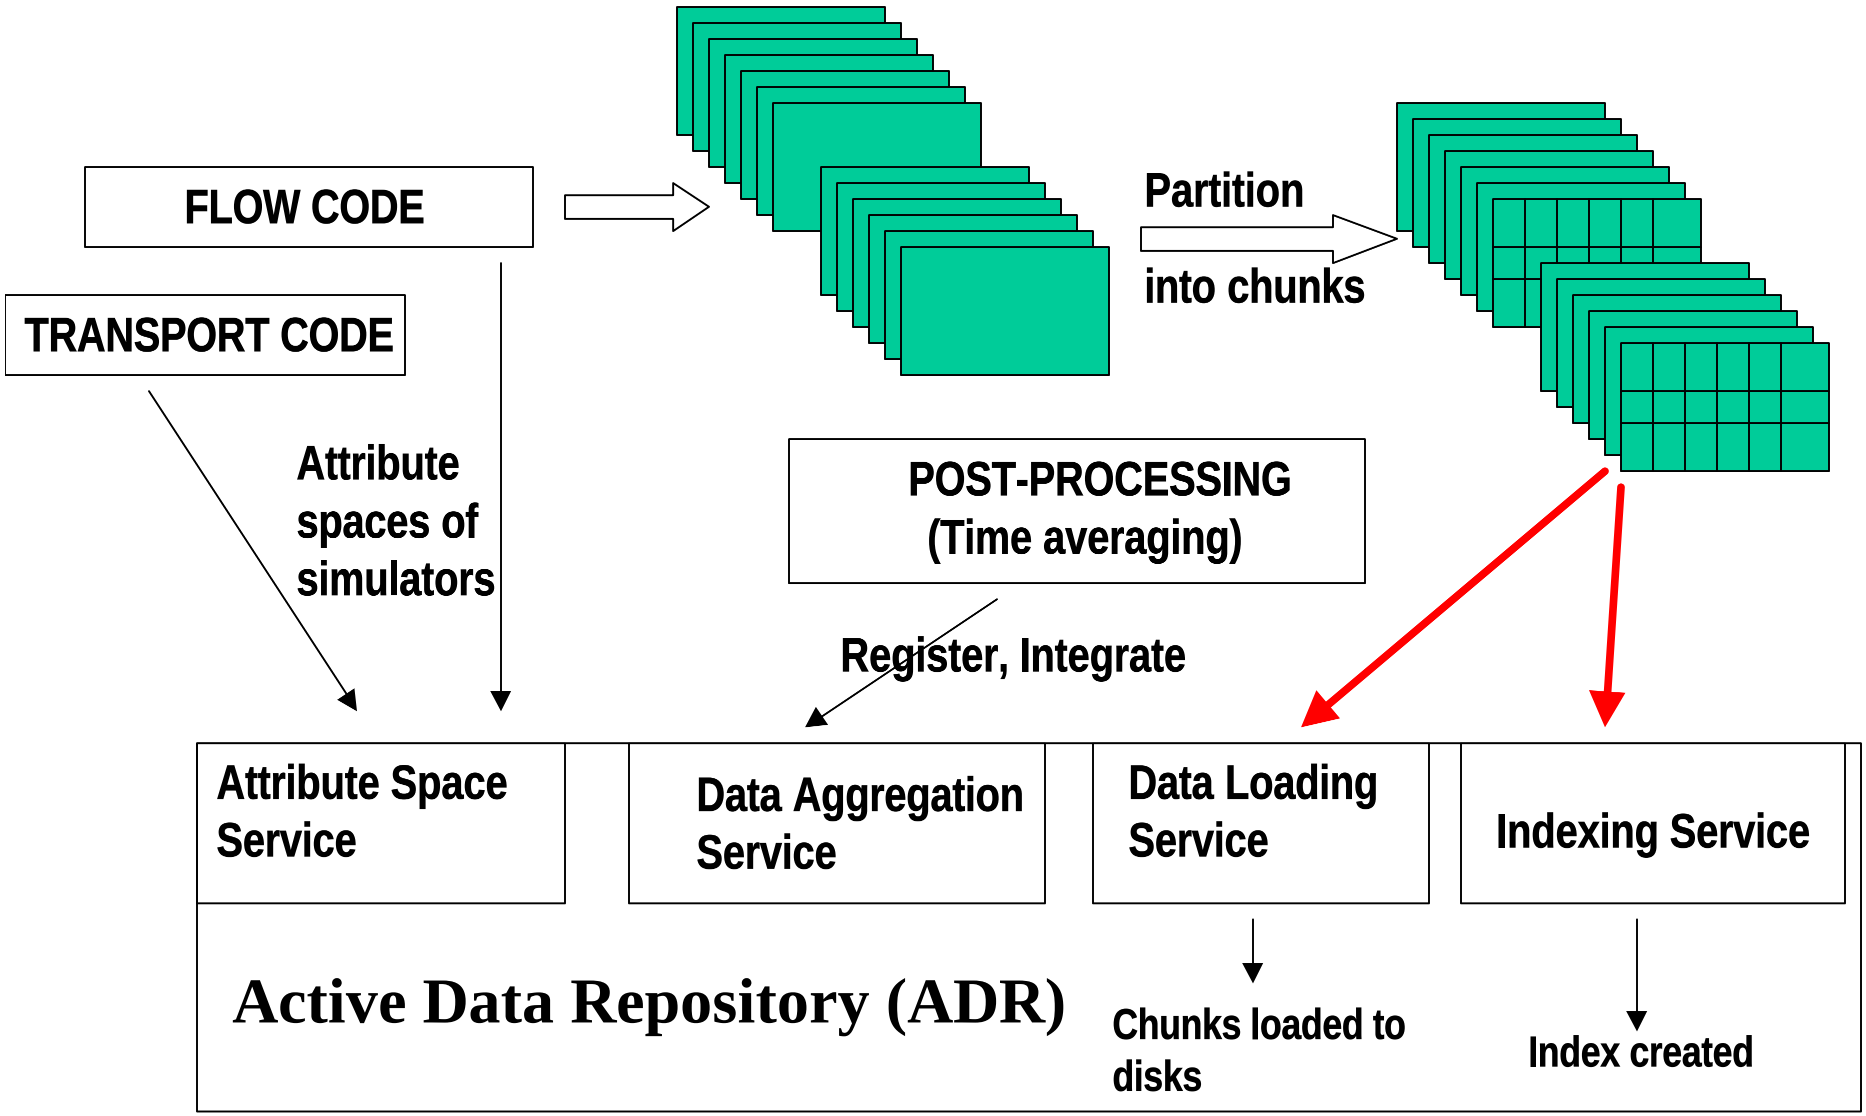

Water Contamination Studies

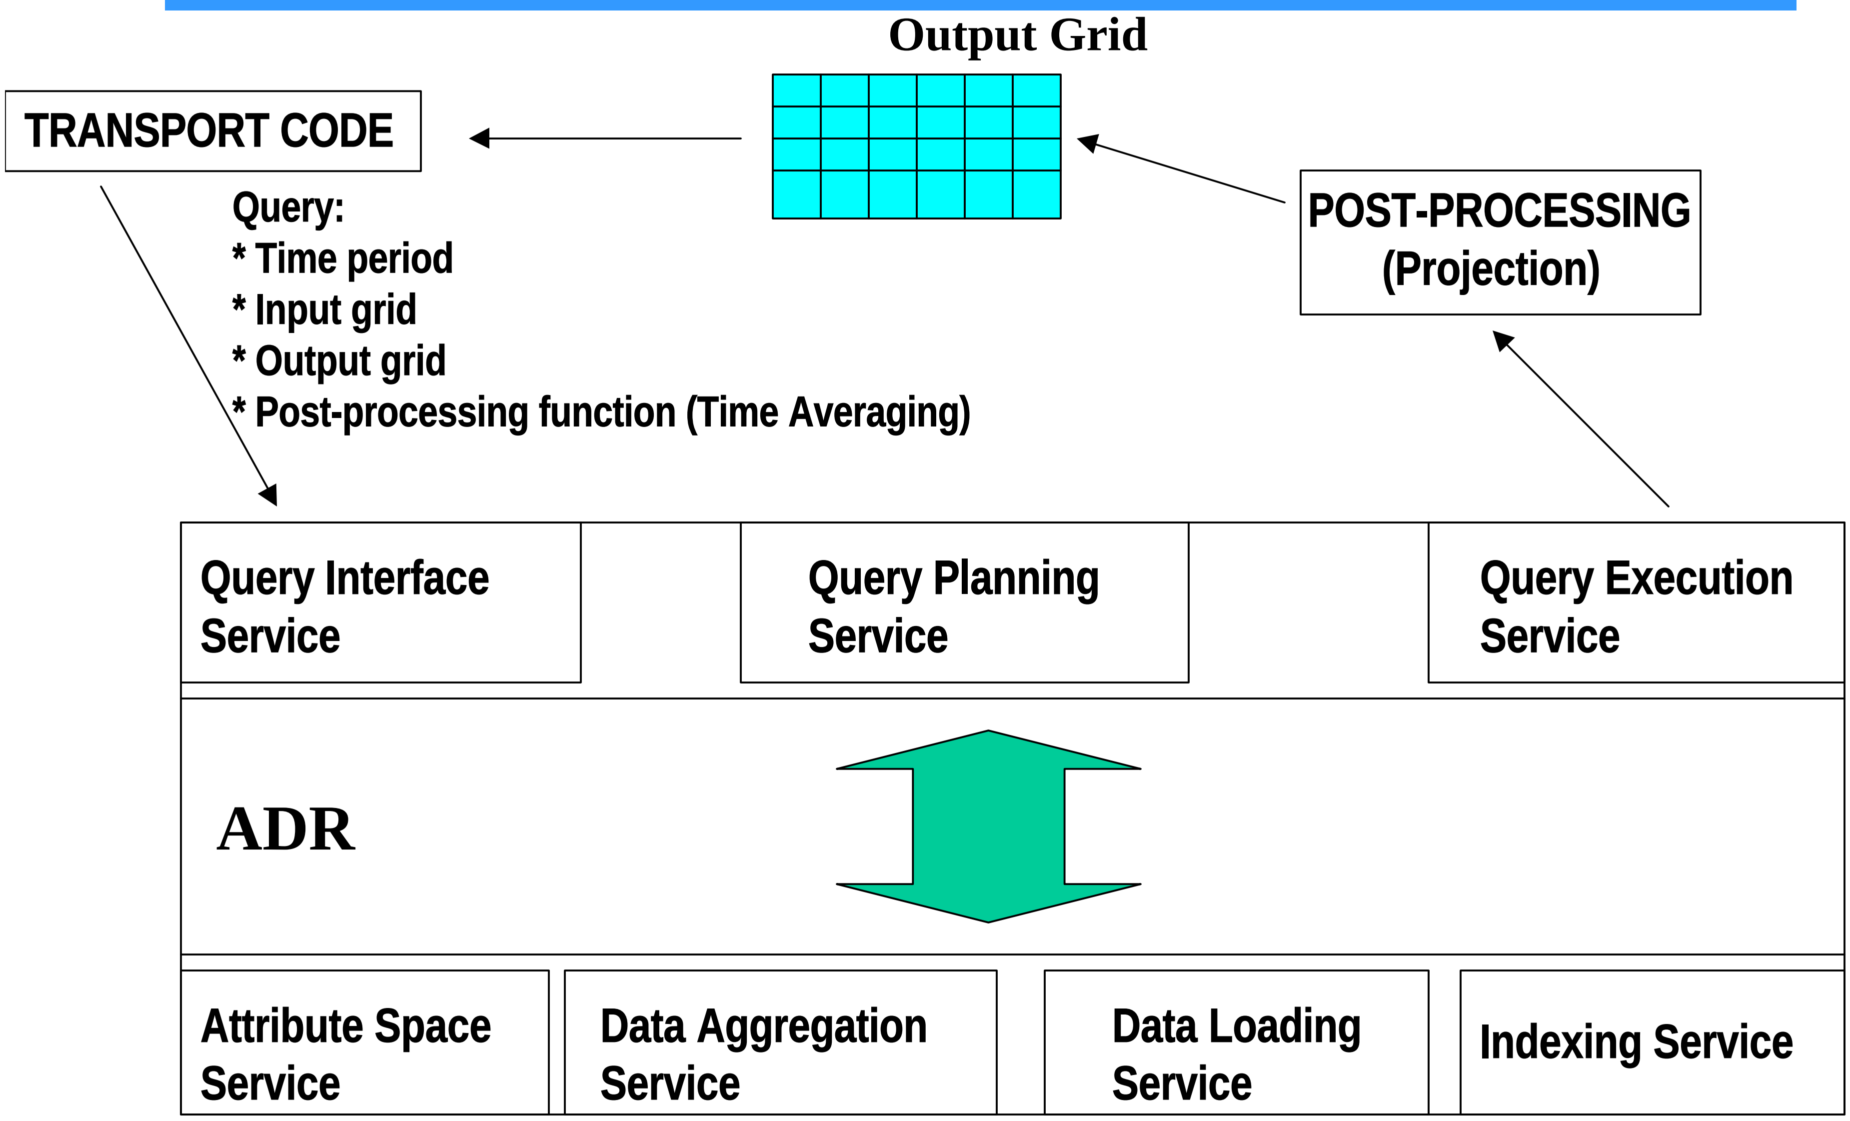

Performance Prediction -- Application Emulators

- Performance prediction:
  - How to configure database
    - data partitioning
    - how many compute processes and where should they run
  - Performance of database various possible architectures
    - standard architectures, petaflop, active disks
  - Insight into how to improve software design to optimize performance

Performance Prediction -- Application Emulators

- Application Emulators:
  - Parameterized program designed to mimic application computation and data movement patterns
  - Focus is on lower layers of memory hierarchy, computational details, cache behavior are abstracted
  - Coarse grained, executable description of patterns of data movement and computation
- Simulator suite to project performance to varying levels of accuracy

Comparison of Real Application and Emulator (Maryland 16 Processor SP-2)

Execution Times, 10-day data

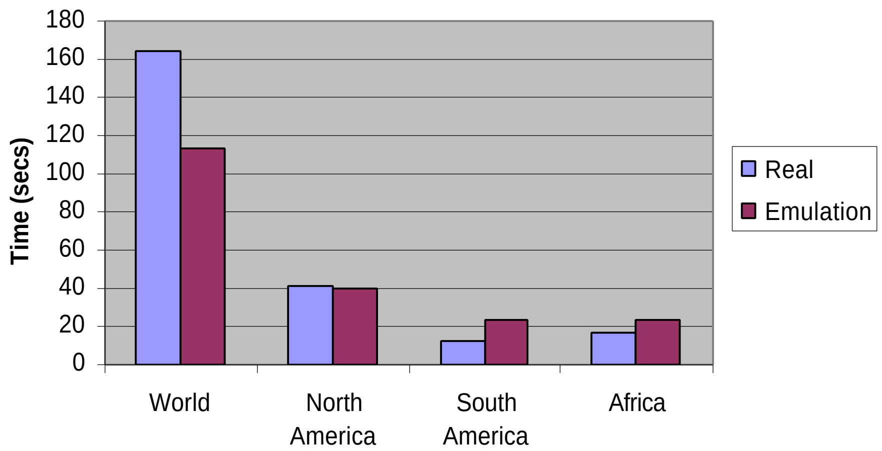

Execution Times, 60-day data

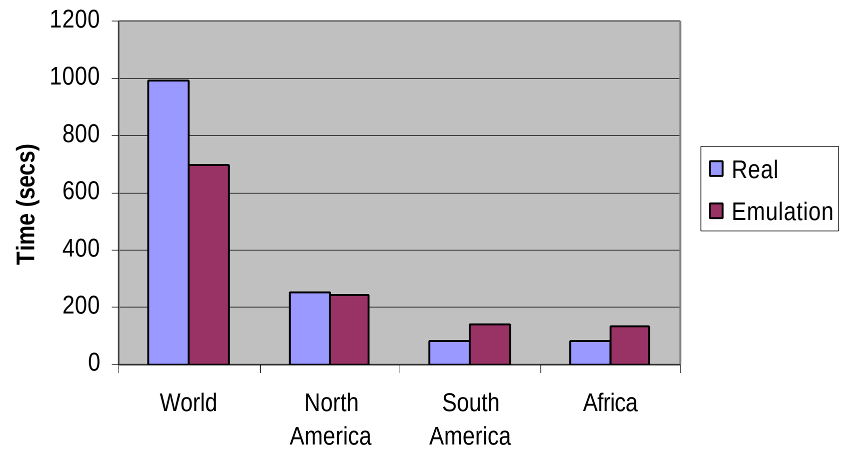

Satellite Data Processing (SP-2)

Scaled Input

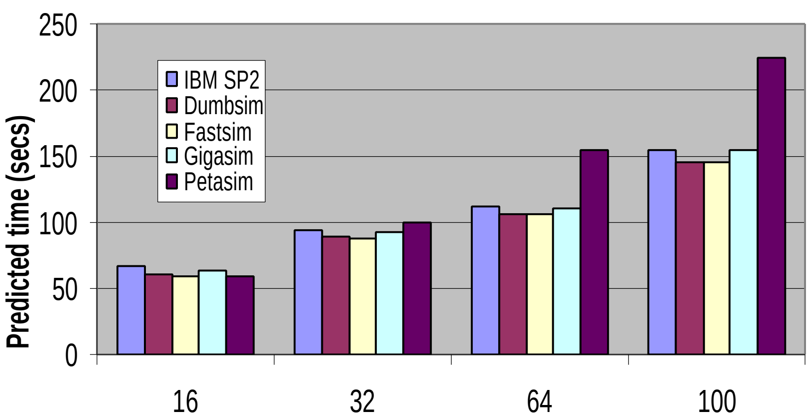

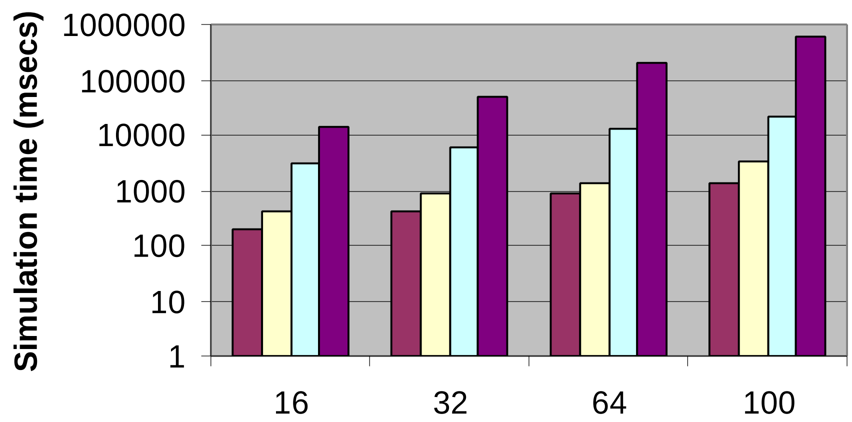

Non-scaled Input

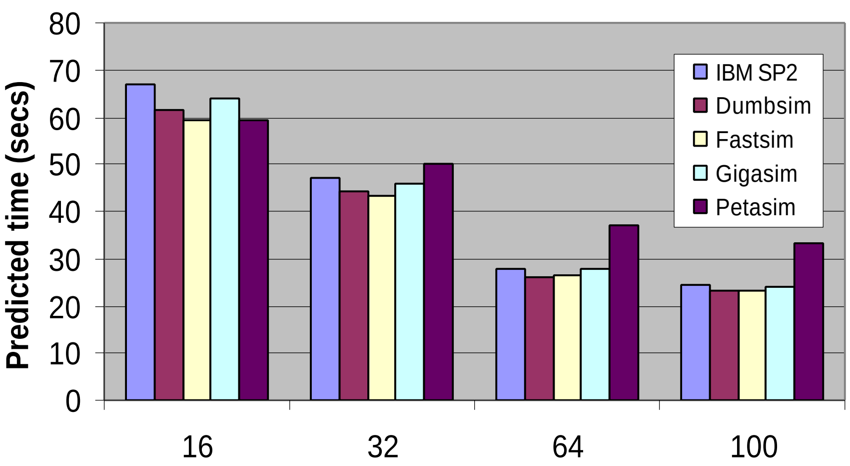

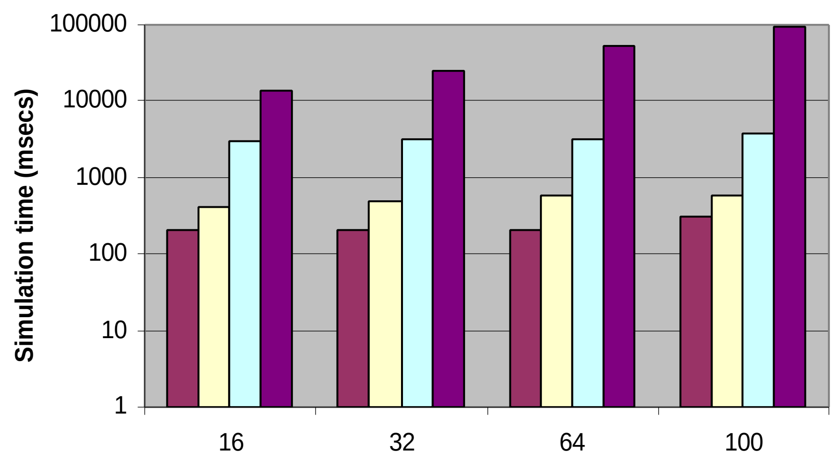

Virtual Microscope (SP-2)

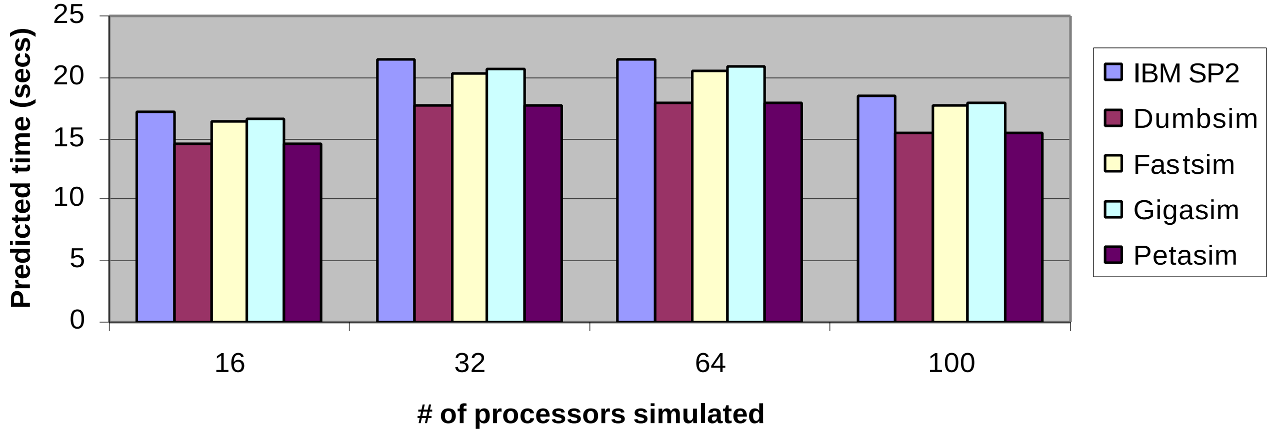

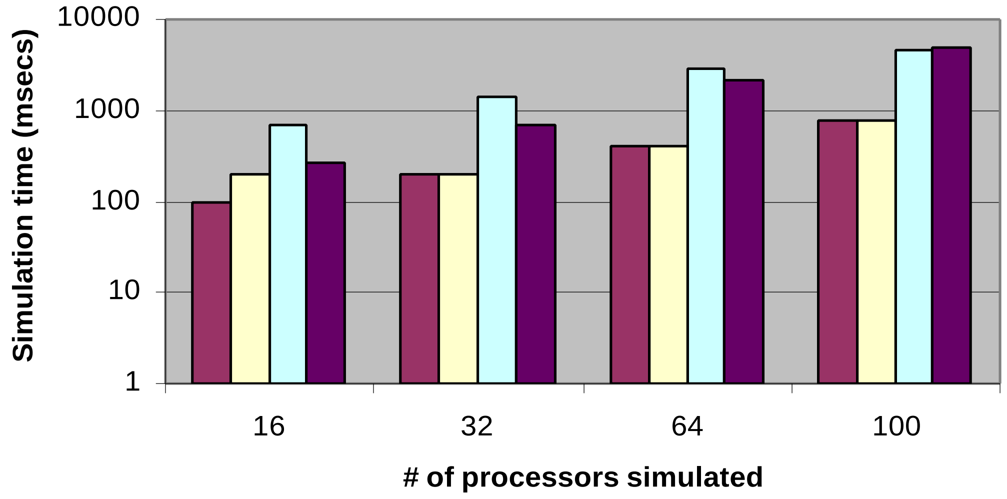

Active Disk Architecture

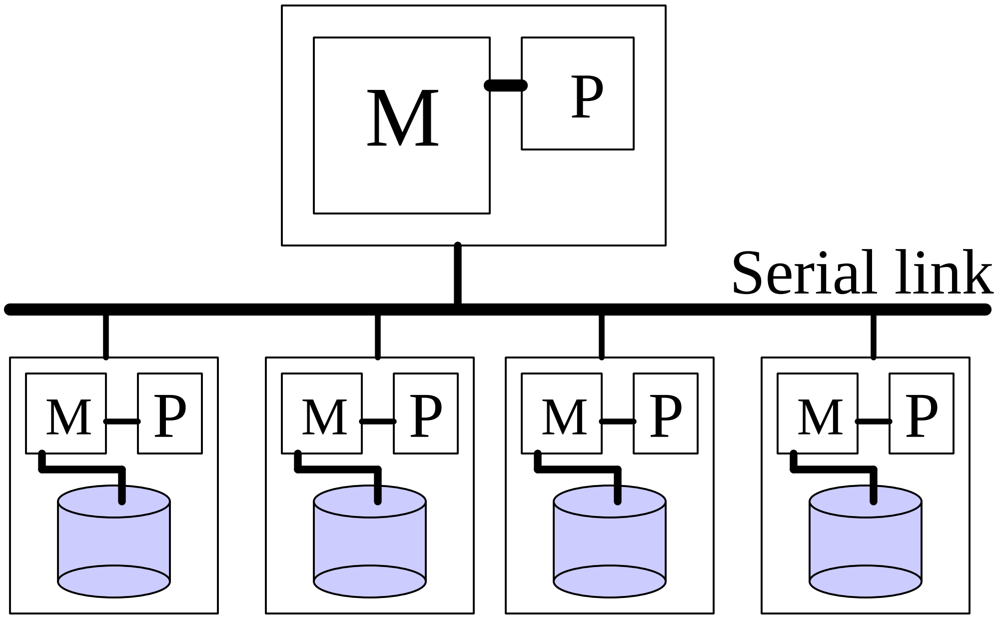

- Restructure apps
- Disk-resident code
  - bulk processing
  - *disklet*
- Host-resident code
  - coordination
  - communication
  - combination

- Processing power scales naturally with storage capacity
- Processing power evolves with storage

Experiments

- Compared algorithm-architecture combinations
  - current and future configurations
- Evaluated scalability
  - configuration: 4-32 disks
  - dataset size
- Evaluated impact of host upgrades

Conclusion

- Large irregular datasets are coming to your computational grid
- Database software can hide complexity of exploiting large irregular datasets
- Performance methodology for configuring database, for choosing target architectures and in predicting cost of queries

Ongoing and Open Issues

- Integration into metacomputing environment
  - ADR metadata to be stored in NPACI SRBs and Globus MDS
  - Process placement and communication coordinated by Globus, scheduled by AppLeS (Generalization of SARA!)
  - Coupling ADR with parallel programs coordinated by MetaChaos
- Integration into object-relational framework
  - Appropriate query language
  - Compile-time/runtime optimizations
    - intermodule tiling