scientific data management: not your everyday transaction

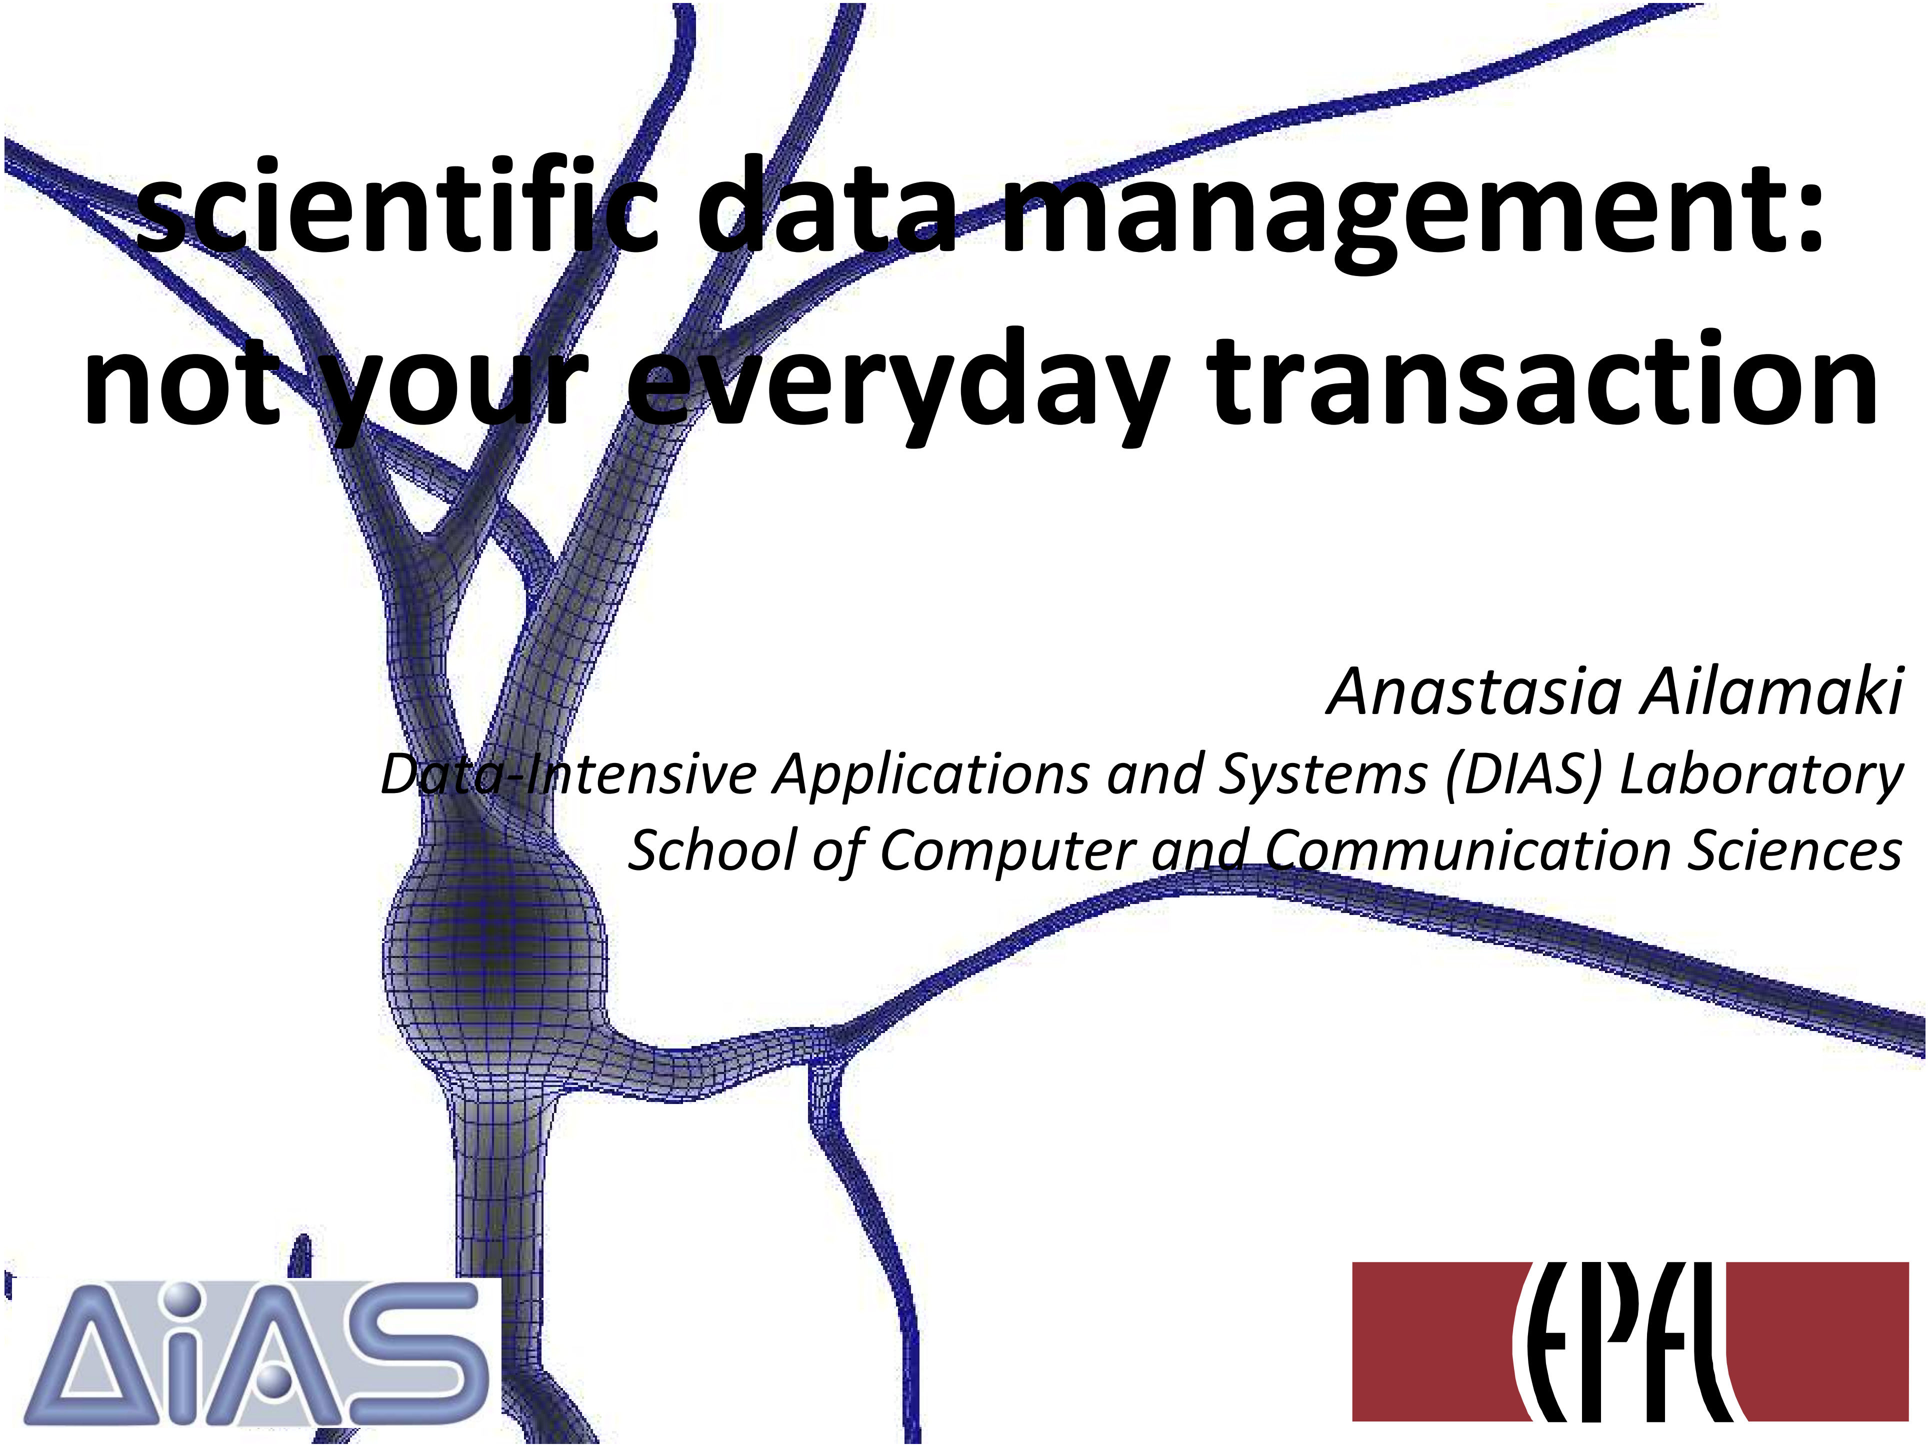

Anastasia Ailamaki
Data-Intensive Applications and Systems (DIAS) Laboratory
School of Computer and Communication Sciences

the human genome at ten (nature, 4/2010)

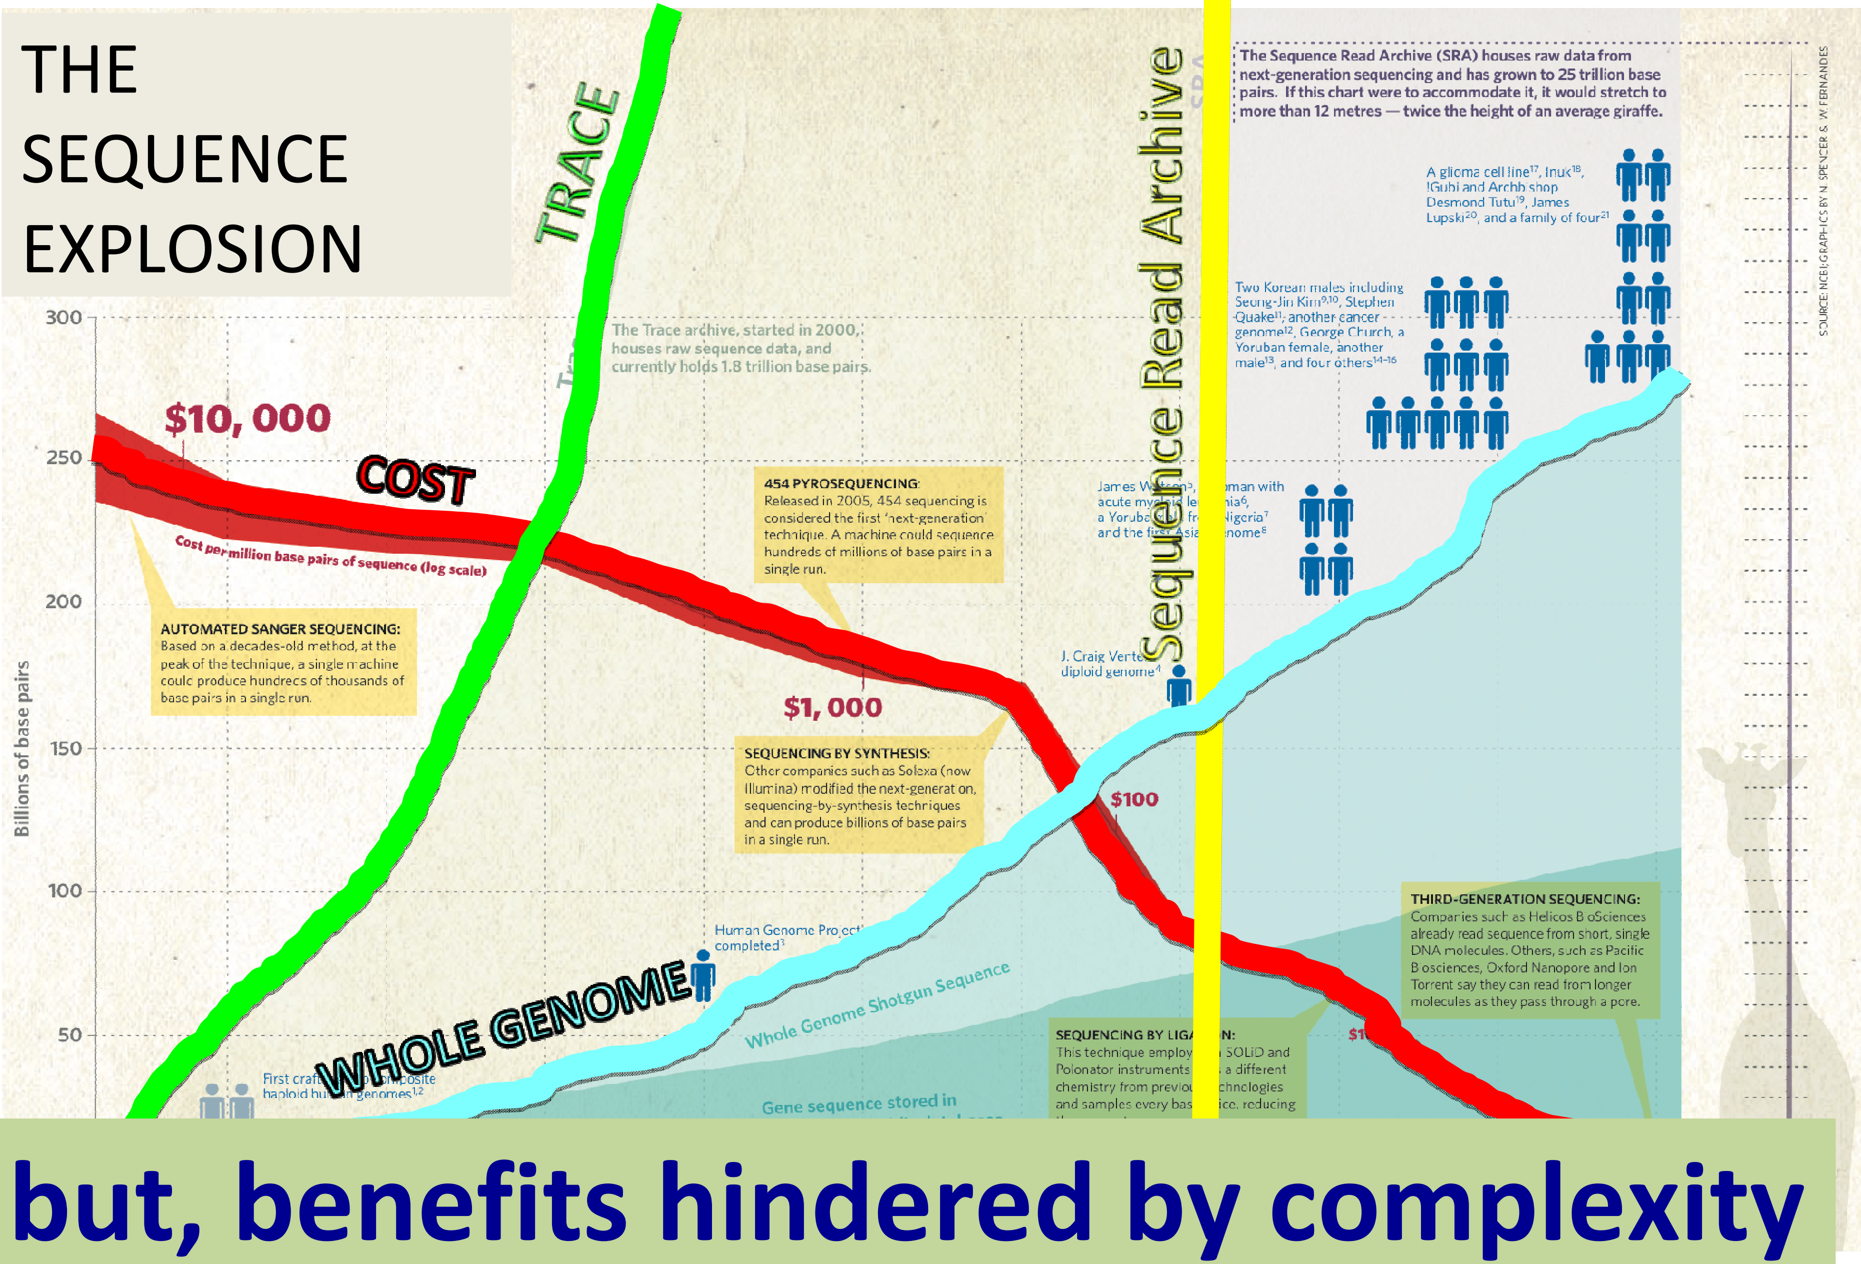

but, benefits hindered by complexity

scientific data grows much faster than technology

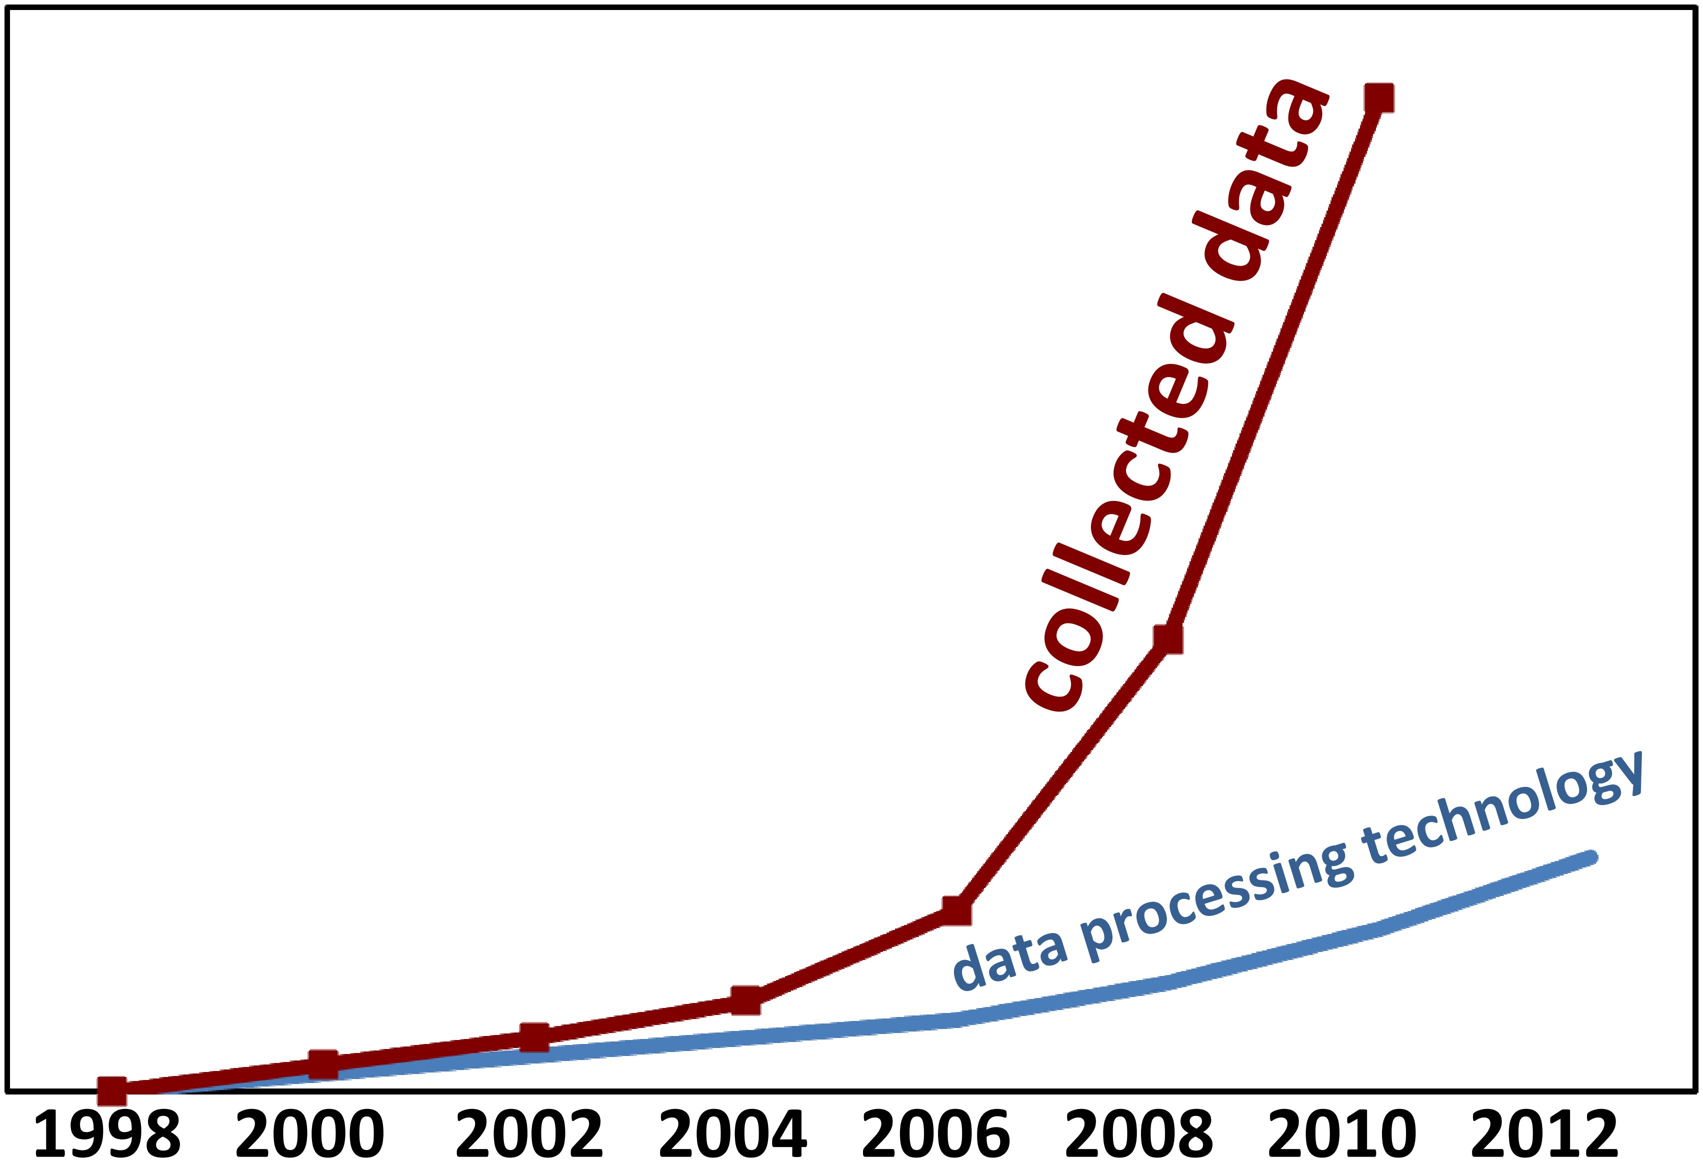

WinterCorp Survey, www.wintercorp.com

scientific data management now

- legacy software
- in main memory of supercomputers
- databases too rigid to use

As data grows, problem changes

- difficult and slow
- some data discarded

bridge CS and domain sciences

meeting domain scientists

1. "Hello, we're SO HAPPY to meet you. *We have SO MUCH data!* PLEASE come visit!"
2. visit lab, pretty pictures (we have SO MUCH data)
3. "Let's have lunch!" (we have SO MUCH data)
4. revisit lab, receive promise to get some data
5. ask for data, no reply
6. play DBA (design/normalize schema, design data, write queries, rewrite queries, talk to tech staff)
7. ask for data, receive 2GB
8. repeat (6), then ask again for data, receive 4GB

why collaboration is hard

- Data is their achievement
  - They do not understand what we will do with it
  - They are afraid of what we may do with it
  - They think that we will put it on the internet
- They are not sure how we will help them
  - Do not recognize their problems in our demos
- They have been “burned” before

support work builds foundation!

It takes around **18 months** of **discussing** and **learning** about a scientific application and dataset *(while providing DBA services)* until a problem which calls for **true innovation** reveals itself.

ORGANIZING SCIENTIFIC DATA

hydrological cycling, 1997

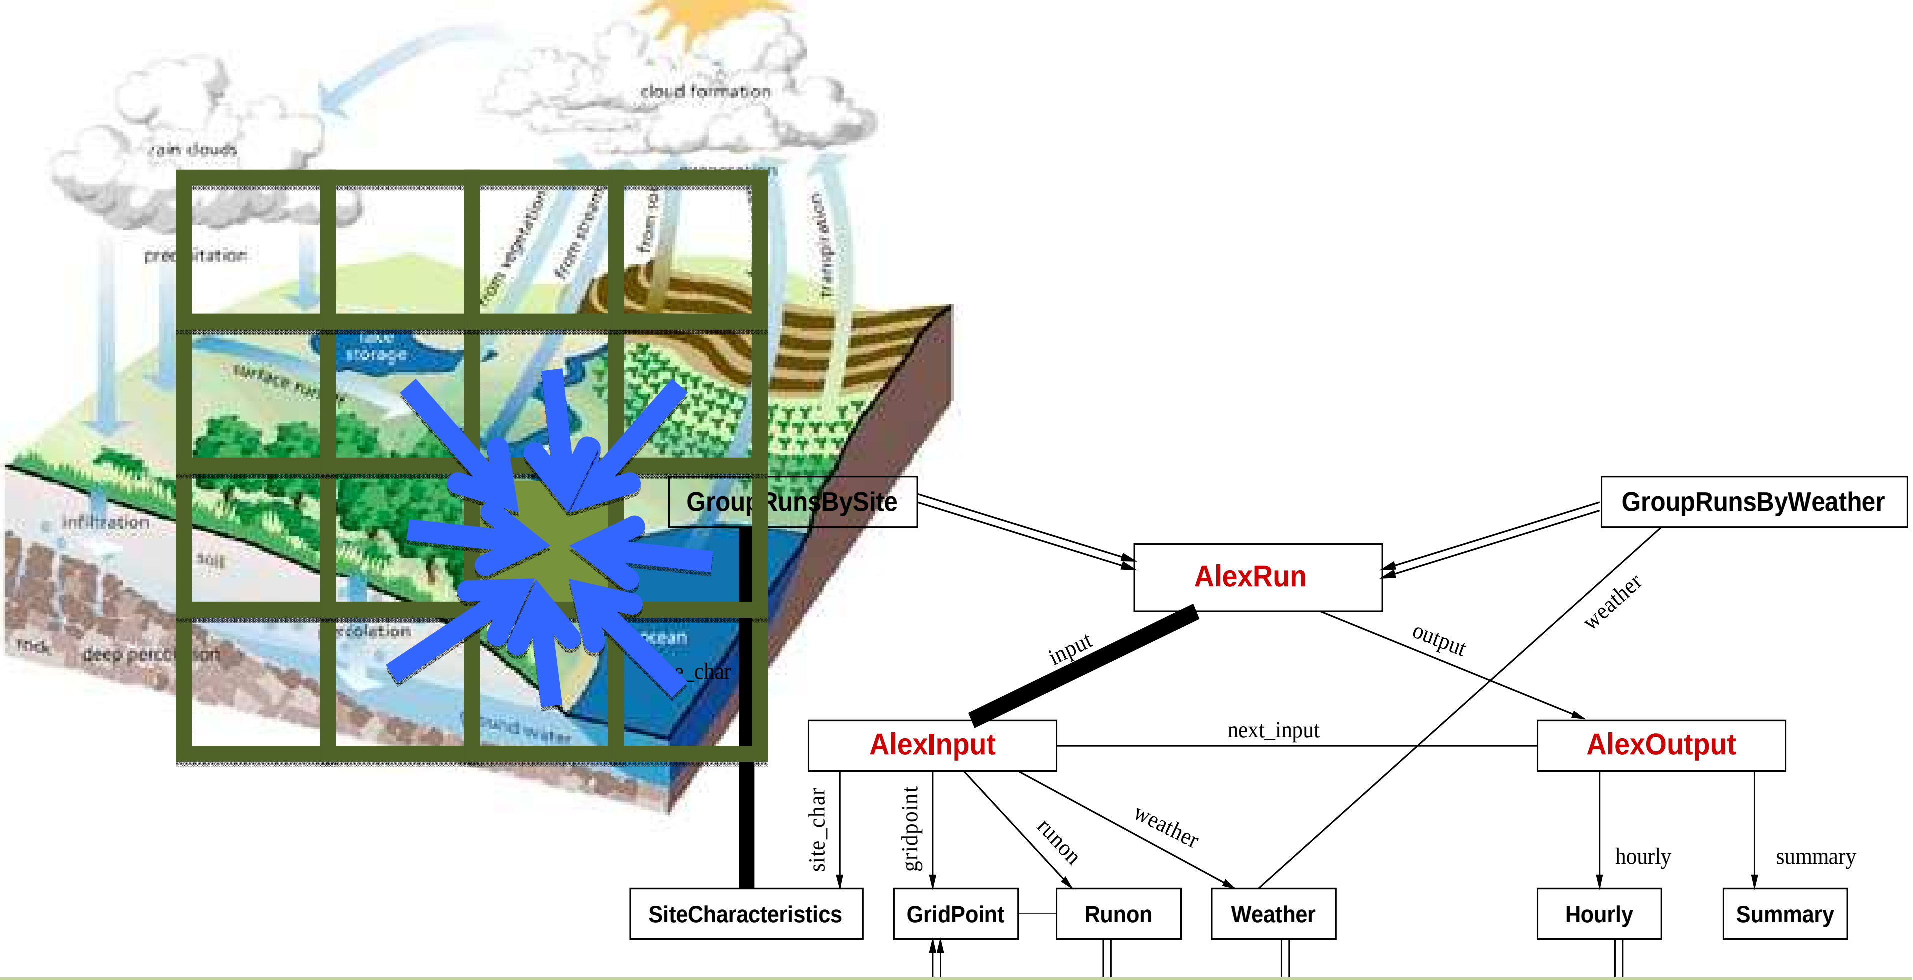

data-centric scientific workflow enactment from within DBMS

the large synoptic array telescope, 2017

date	time	type	category	Cx	Cy	Cz	…
mon	22:00	star	*proto*	…	…	…	…
mon	22:15	star	*red giant*	…	…	…	…
mon	22:20	galaxy	*spiral*	…	…	…	…
mon	22:27	star	*dwarf*	…	…	…	…
mon	23:00	galaxy	*elliptical*	…	…	…	…
tue	22:10	galaxy	*spiral*	…	…	…	…
tue	22:25	star	*proto*	…	…	…	…
…	…	…	…	…	…	…	…
…	…	…	…	…	…	…	…

100 billion objects, 20 PB/night

PARINDA DB designer

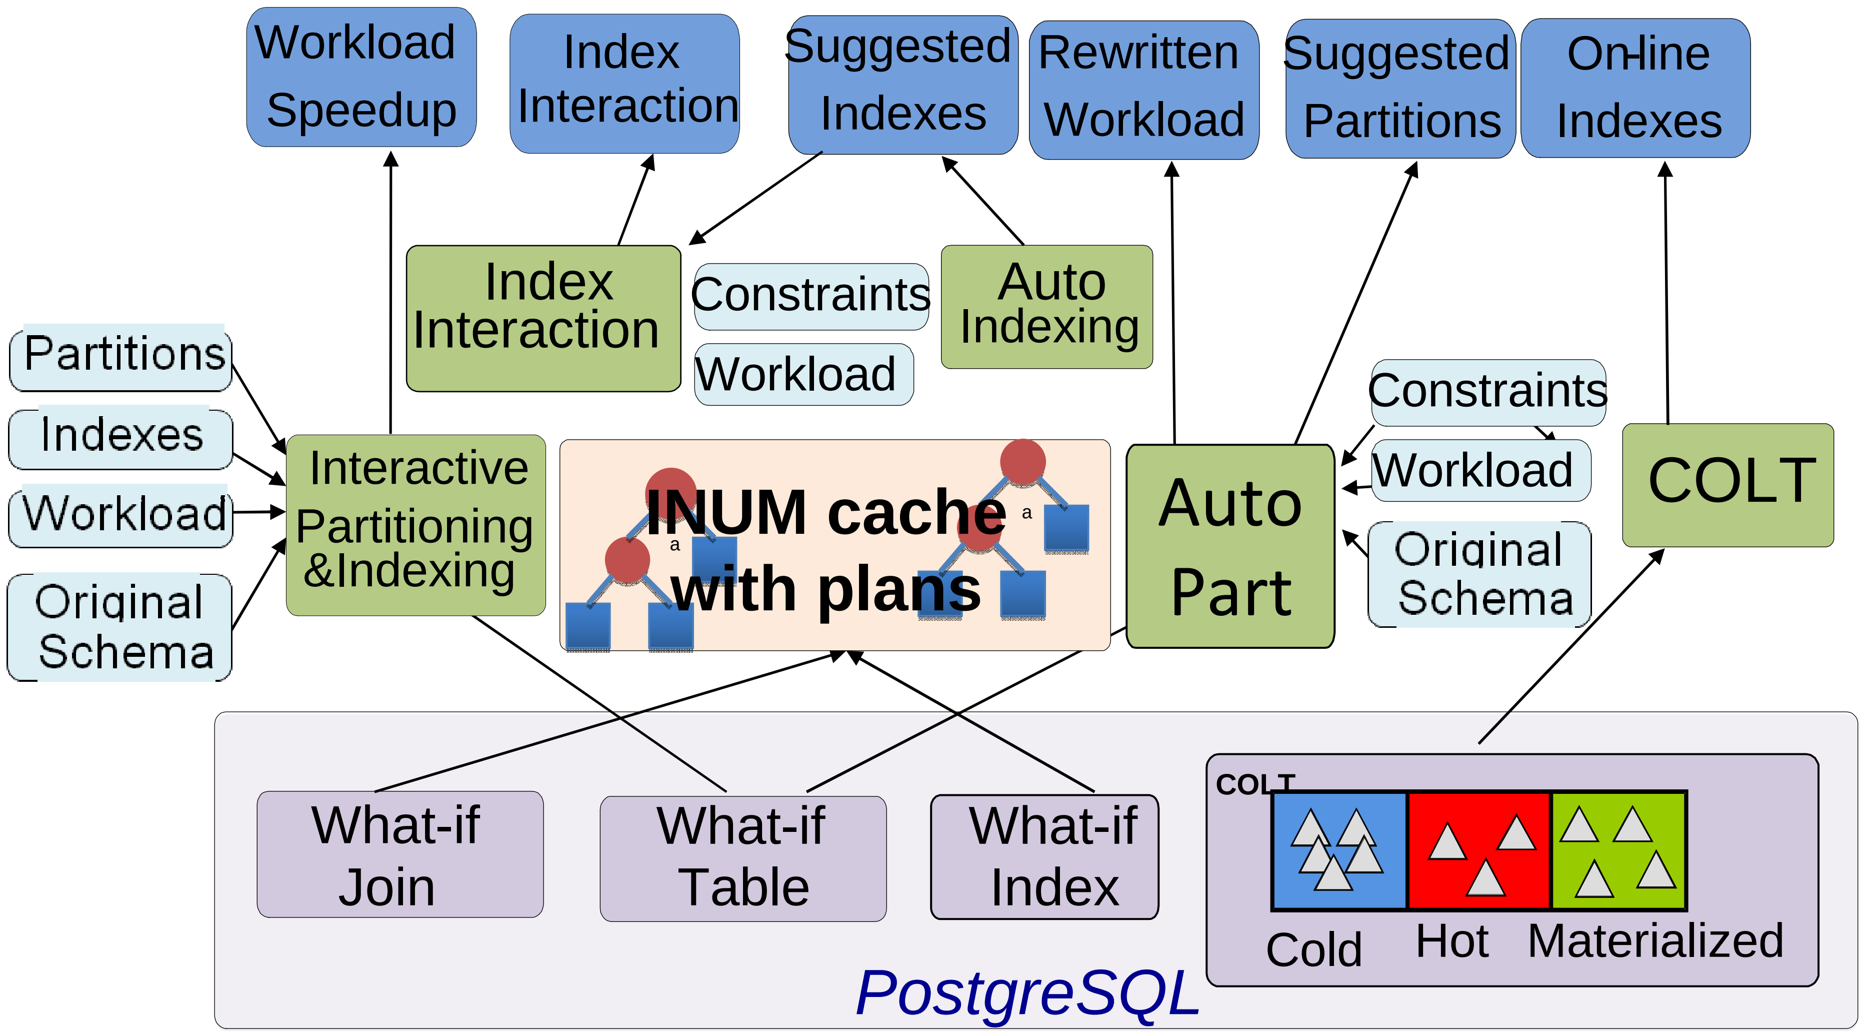

IMPLEMENTING PHYSICAL MODELS

earthquake simulation/analysis

Quake Group [www.cs.cmu.edu/~quake]

Scientist

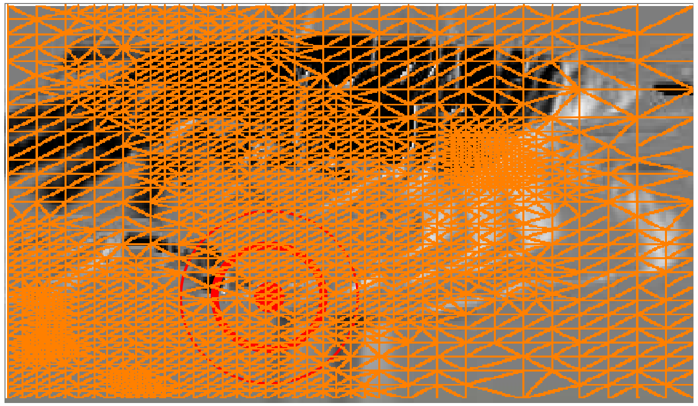

Simulation Model

esh)

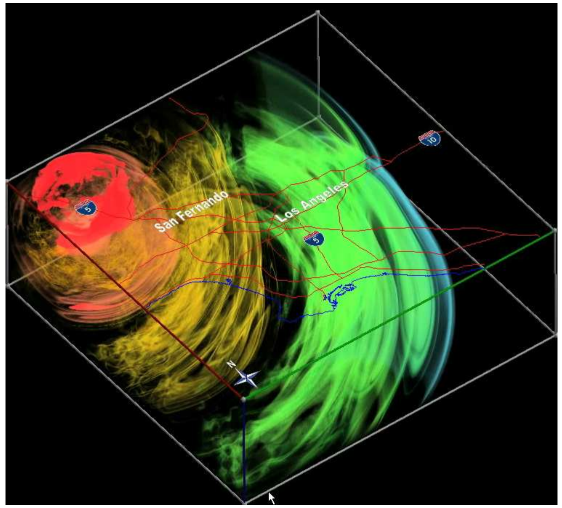

Simulation Output

Analy

fopen('time050.dat')

time030

time020

time010

100GB per time-step, 20000 time-steps

performance? access to complex structures?

tetrahedral mesh models

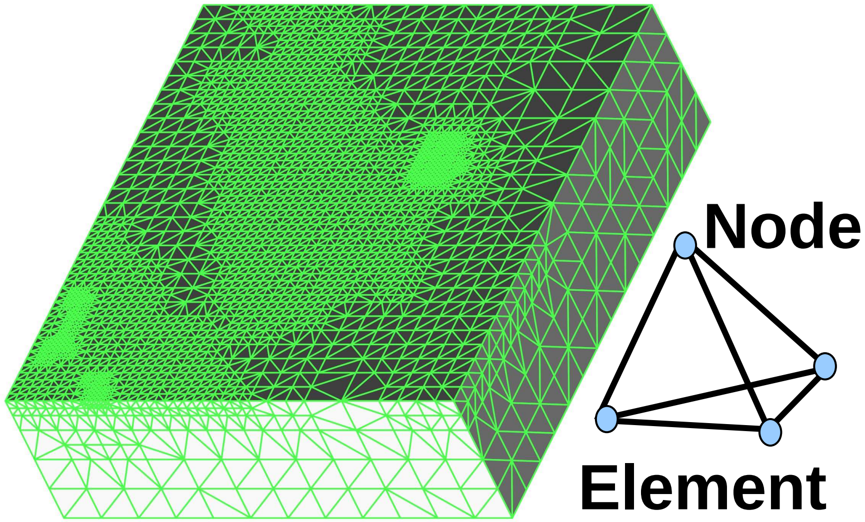

- Queries
  - Point
  - Range
  - Feature

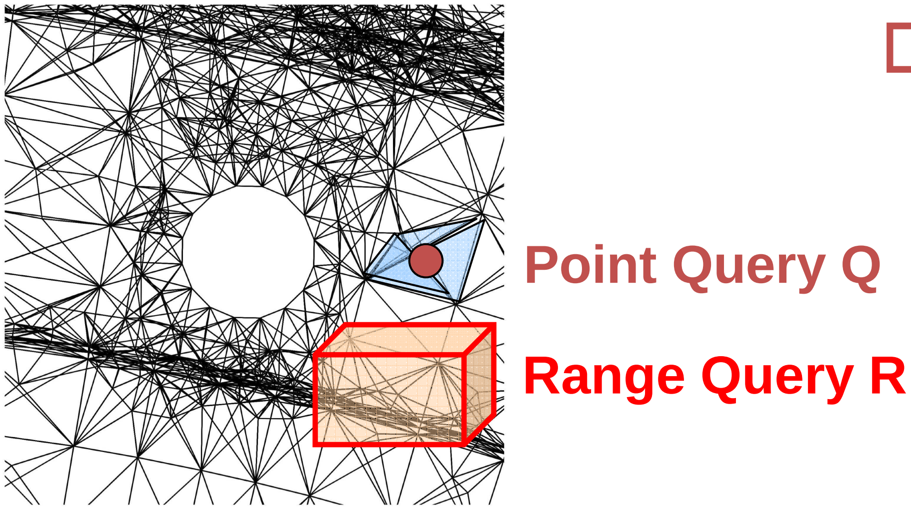

Courtesy Cornell Fracture Group

- Example: Visualization
  - Show ground velocity at **Q**
  - Draw the temperature of **R**

goal: efficiently process *geometric* queries

R-Tree indexing

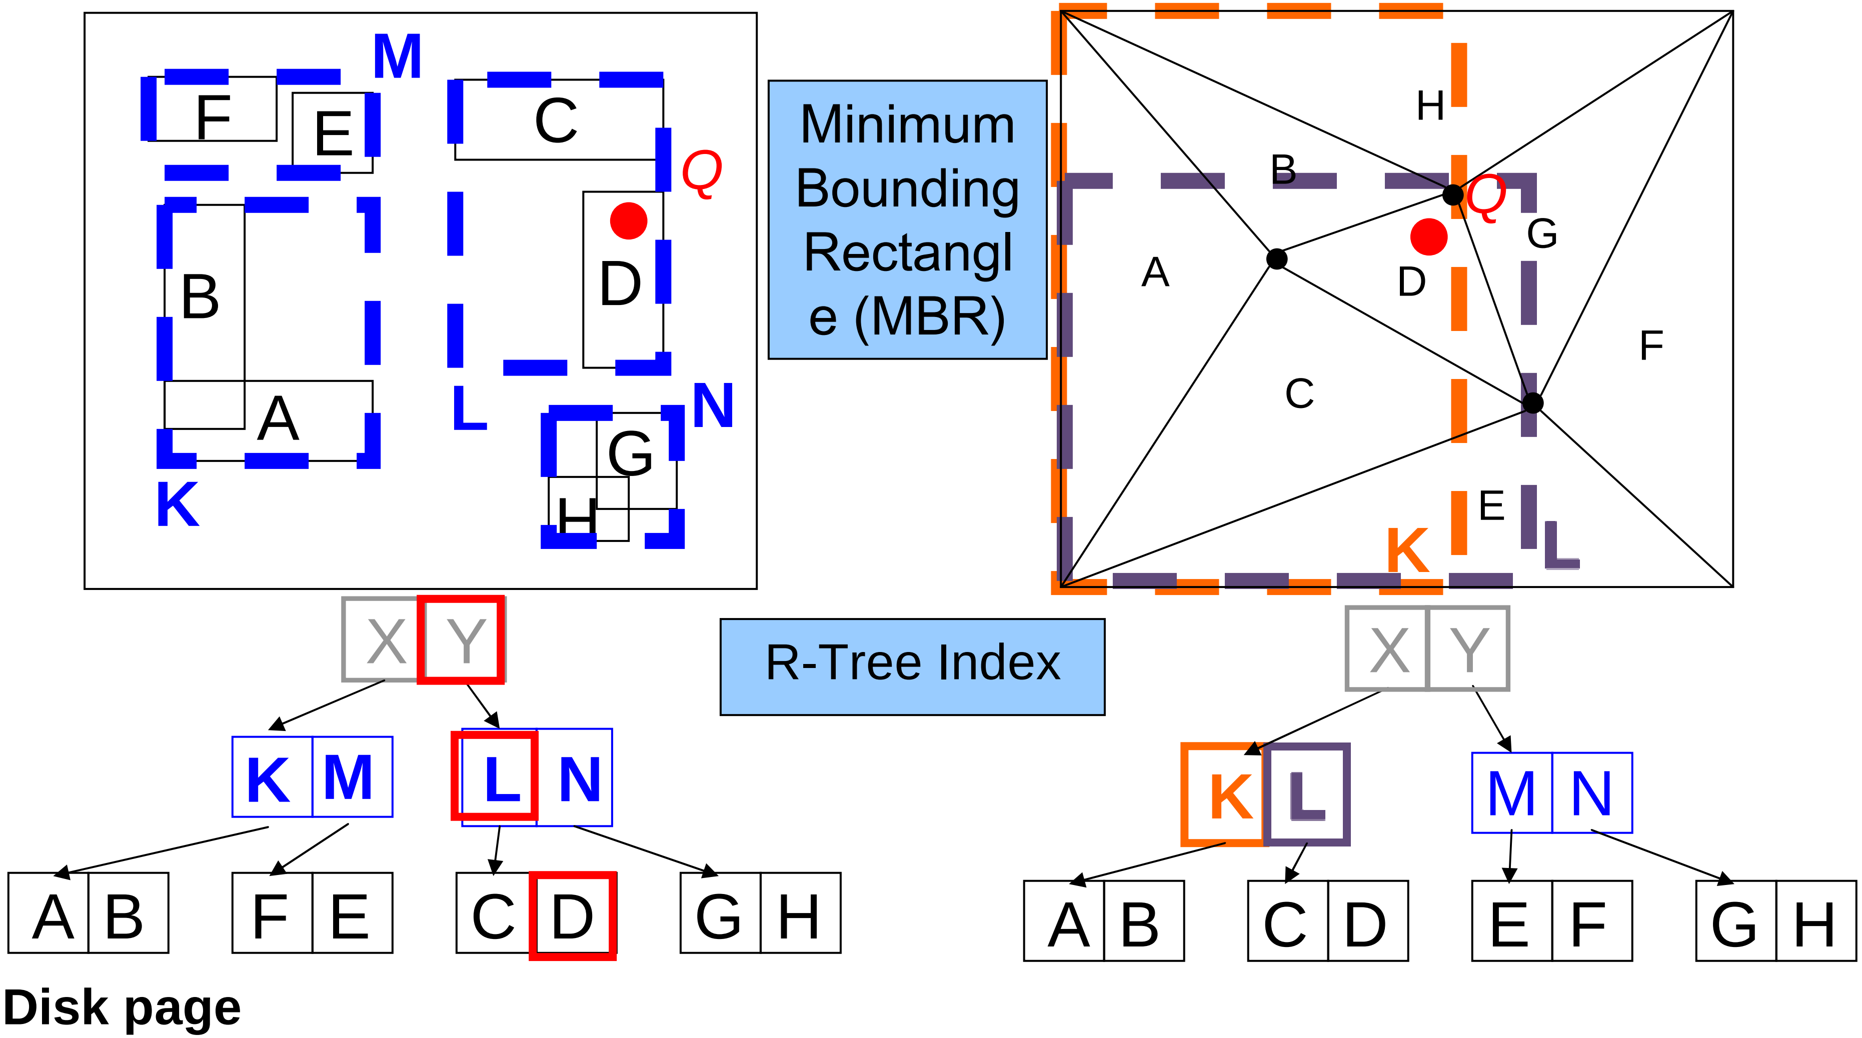

tight mesh connectivity hurts performance

directed local search

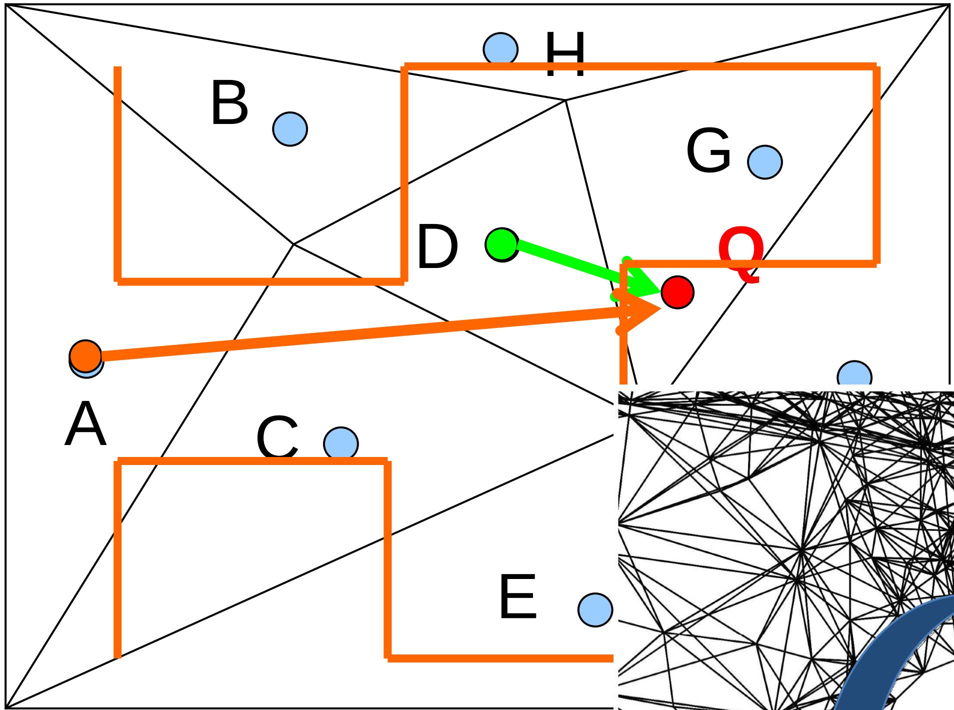

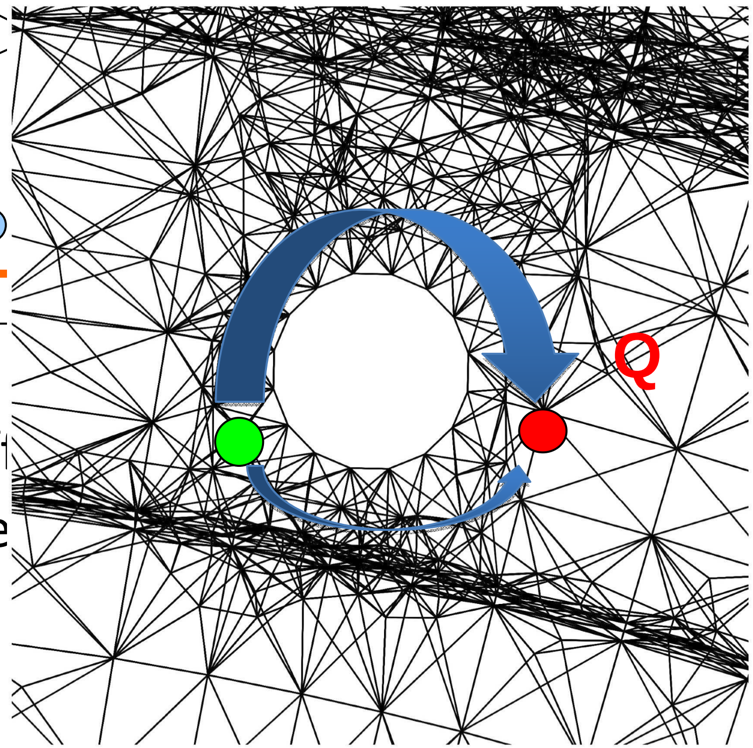

- Point queries
  - Step 1: Find a “nearby” element eff
  - Step 2: Follow path of adjacent ele
- Range queries are similar

databasing the brain

a spatial data management challenge

IBM Journal for Research and Development, Issue 52, Number 1/2, 2008

brain simulation

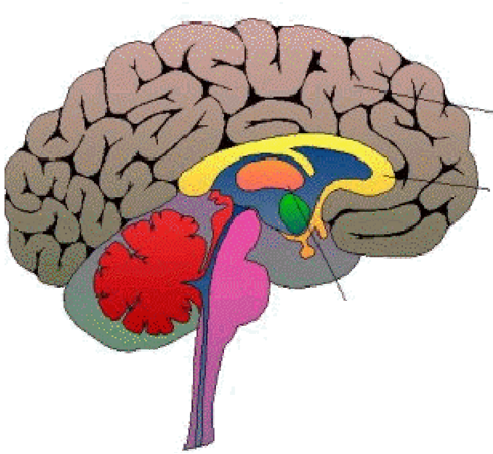

neo
cortex

images courtesy of the Blue Brain Project

a neuron

dye loading &
raster scanning

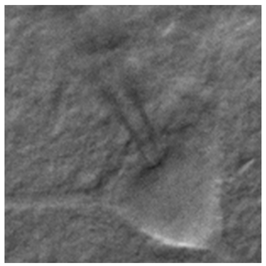

3D reconstruction

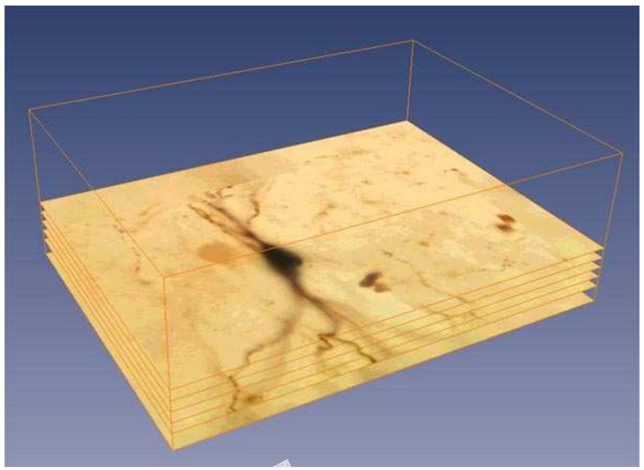

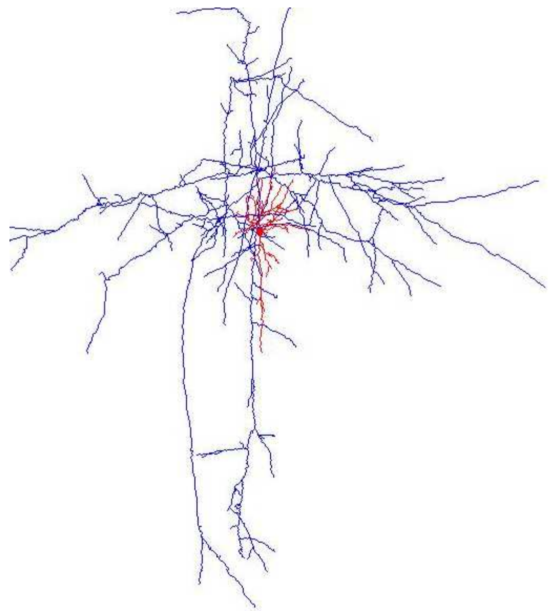

morphologies

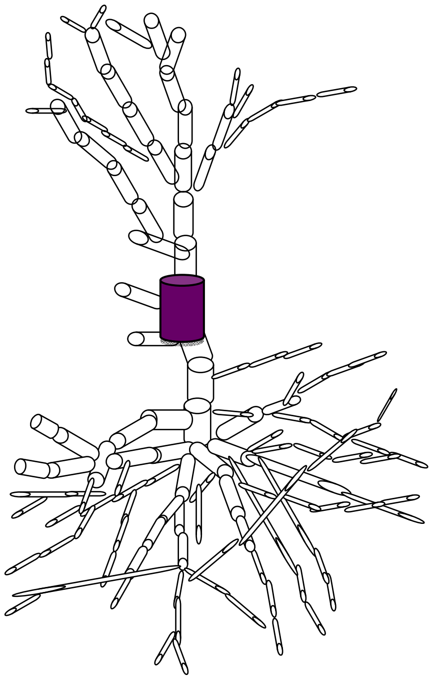

**single neuron,
modeled with 3D cylinders**

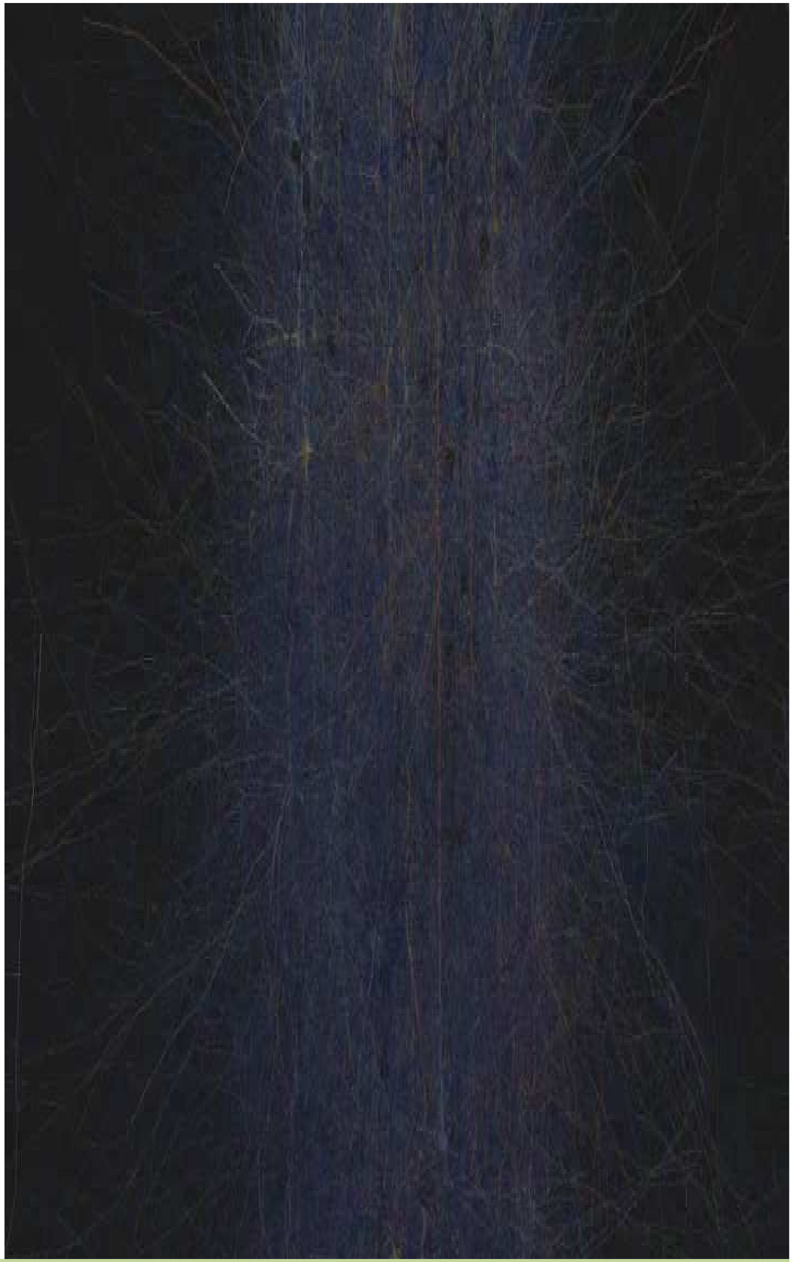

human brain: ~86 billion neurons

brain data deluge

	before	now	future
# of neurons	1692 1 Layer	100K 10 Columns	100B Full Brain
segment representation	4.5M 140MB	450M 13.4GB	4.5×10^{14} 14PB
mesh representation	173M 5.8GB	17.3B 580GB	1.7×10^{16} 0.5EB

molecular representation more fine grained
simulation trace = infinite data

spatial range

morphometric analysis

visualization

model construction

Goal: execute *efficiently* 3D spatial range queries.

- even if data no longer fits into main memory
- even if density of dataset increases

traditional spatial indexing

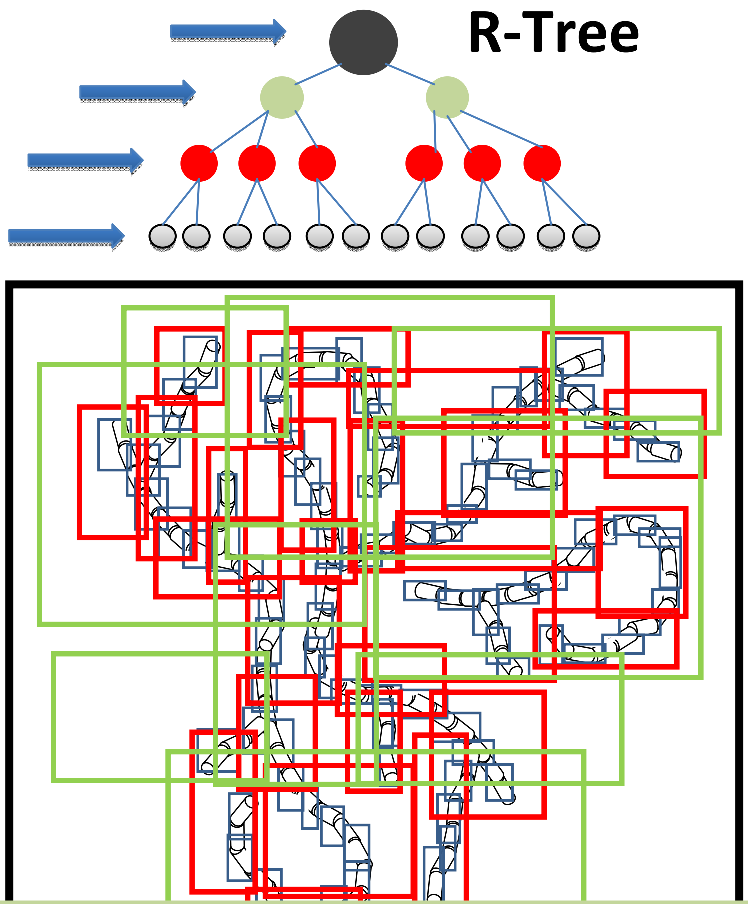

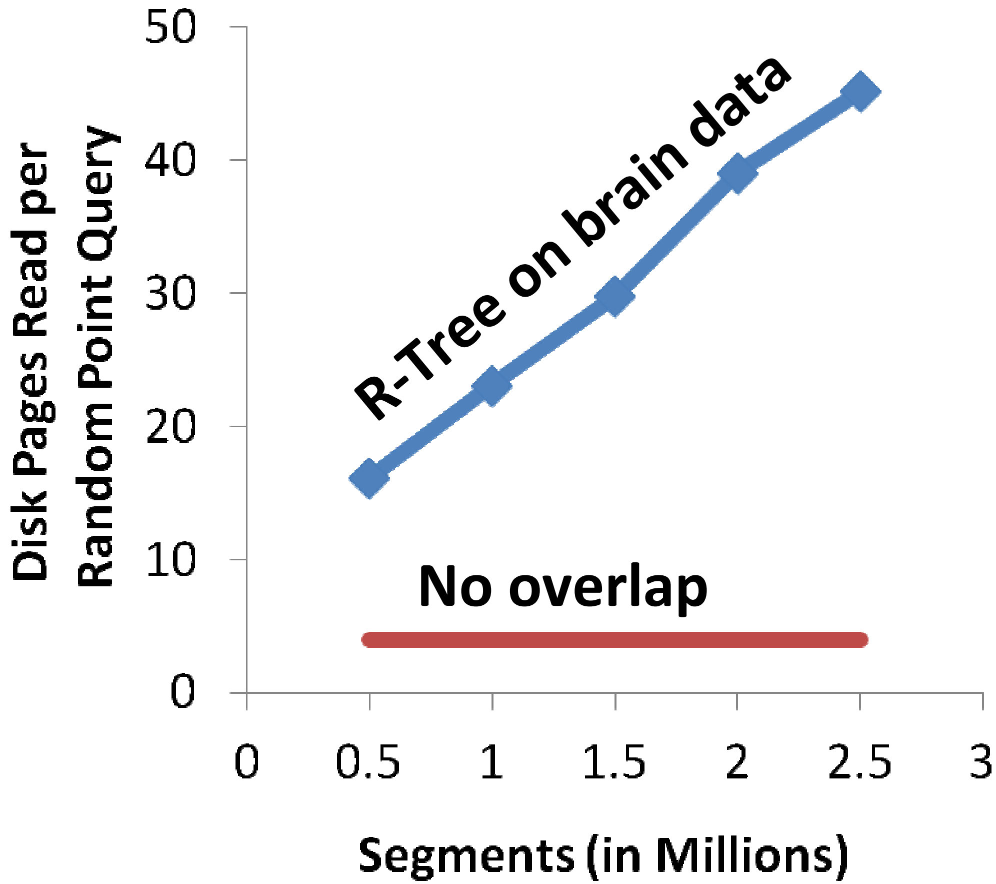

too much structural overlap

FLAT: insensitive to density

- Indexing
  - group spatially close objects on disk pages
  - add links between neighboring groups
  - use traditional R-Tree to index disk pages
- Querying
  - ***seed phase***: find random element inside range in R-Tree (not affected by overlap)
  - ***crawl phase***: use seed element and recursively traverse all neighbors

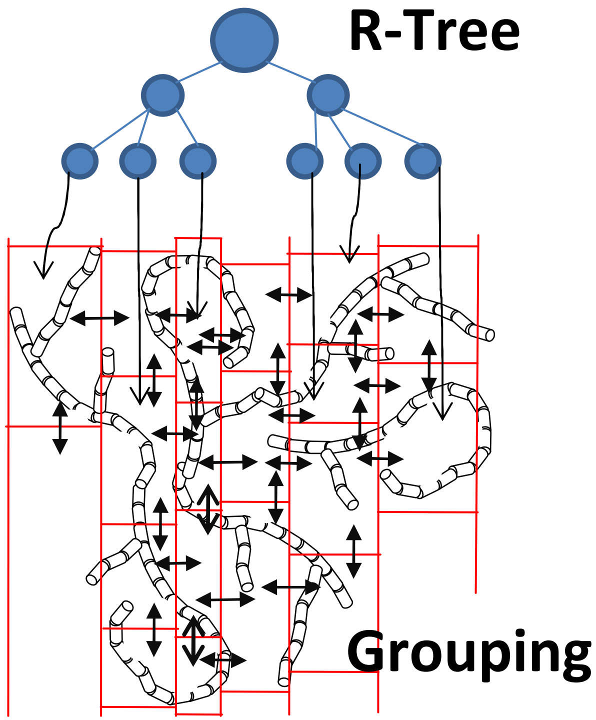

FLAT indexing

from 10K to 1 million neurons!

SCOUT: moving range queries

Neural tissue density analysis:

User issues stream of queries

I/O too costly

even with spatial index

Predict next query

smart&efficient prefetching

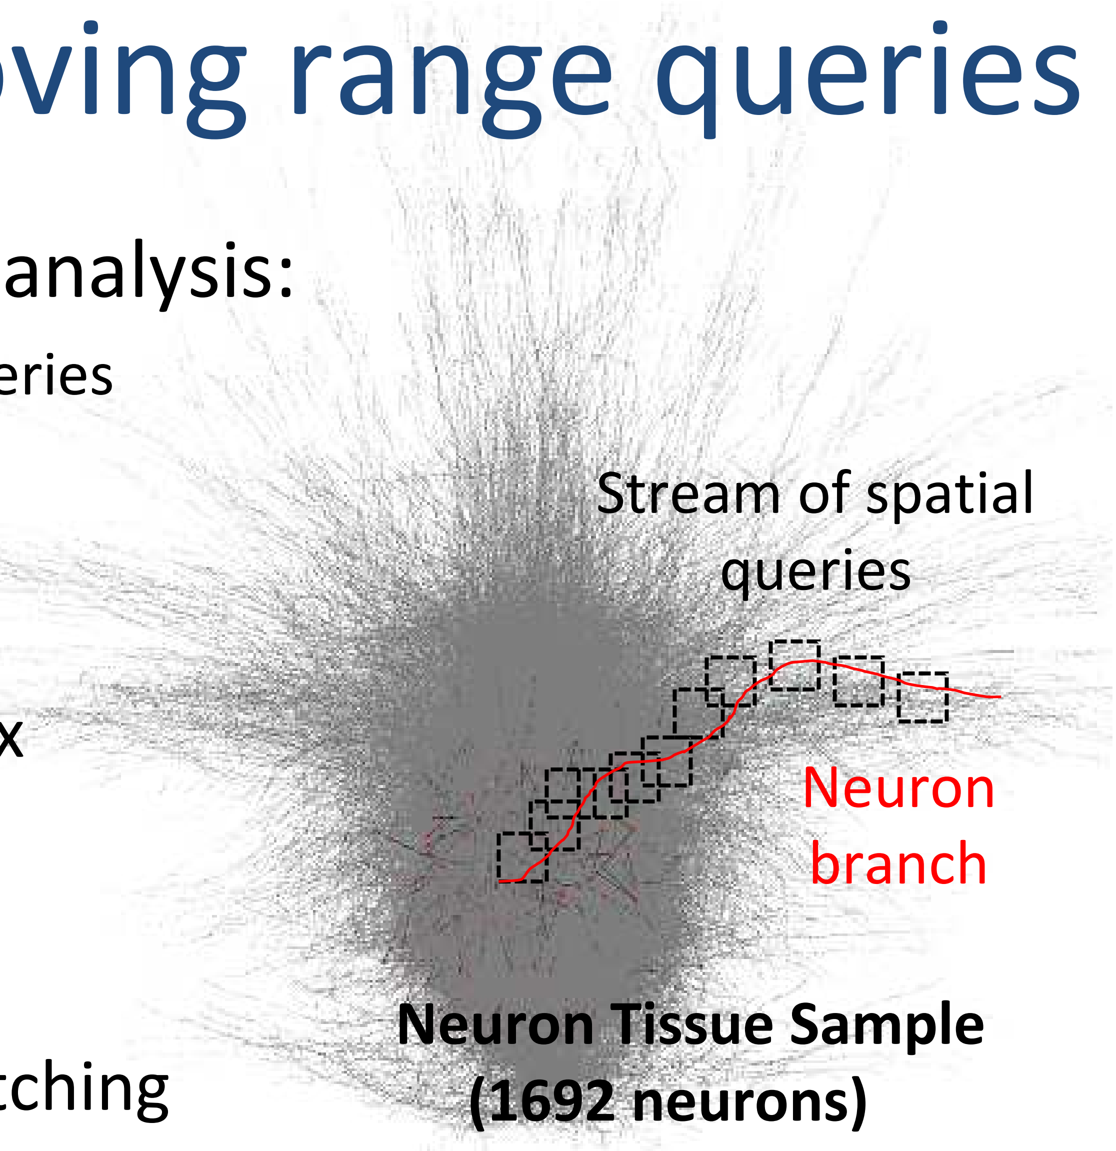

query history, density, or content

touch detection

Model Synapses

electrical connections between axons and dendrites

Data Challenge

100K neurons => 5B synapses

30GB addl space to store synapse data

Human Brain => 100B neurons ~PB space

Need efficient spatial proximity queries and precise distance calculation

a major bottleneck in brain simulation

simulation trace analysis

Need *accurate* data statistics to

- Discover and explore neuro-circuit behavior
- Compare to behavior of biological tissue
- Understand plasticity

Typical trace file ~0.8TB

for 100K neurons

for only 1 second of simulation

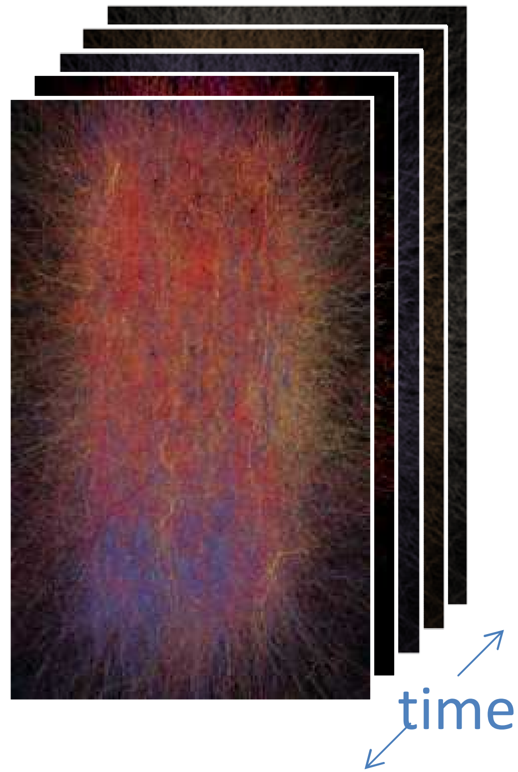

In-memory efficient access method limit use of complex query analysis

Storage capacity limits longer simulation time

need aggressive spatial compression

vision: the human brain project, 2021

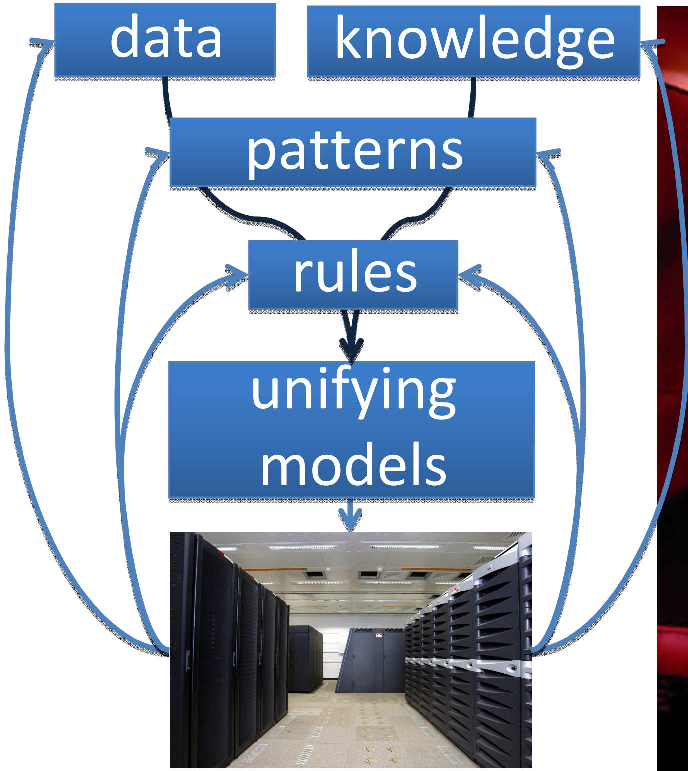

from molecules to cognition:

a (big) data integration problem

BROADER IMPACT

data-driven science

Past:

- theory
- simulation
- experiments

The “fourth paradigm”

scientific breakthrough through computing on massive data

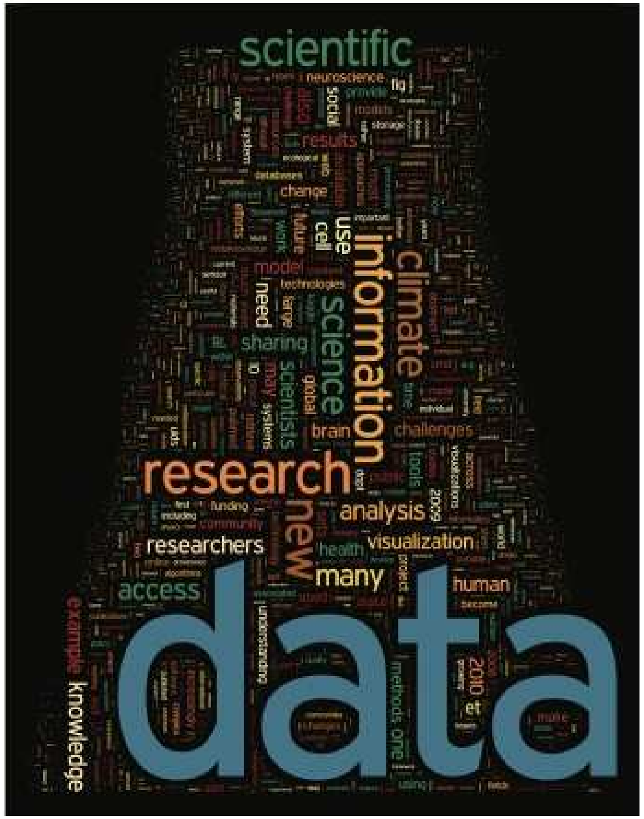


the CERN large hadron collider, now
27km
100M sensors/detection
40M detections/sec
15 PB/year
© 2006 CERN

ATLAS experiment (simplified)

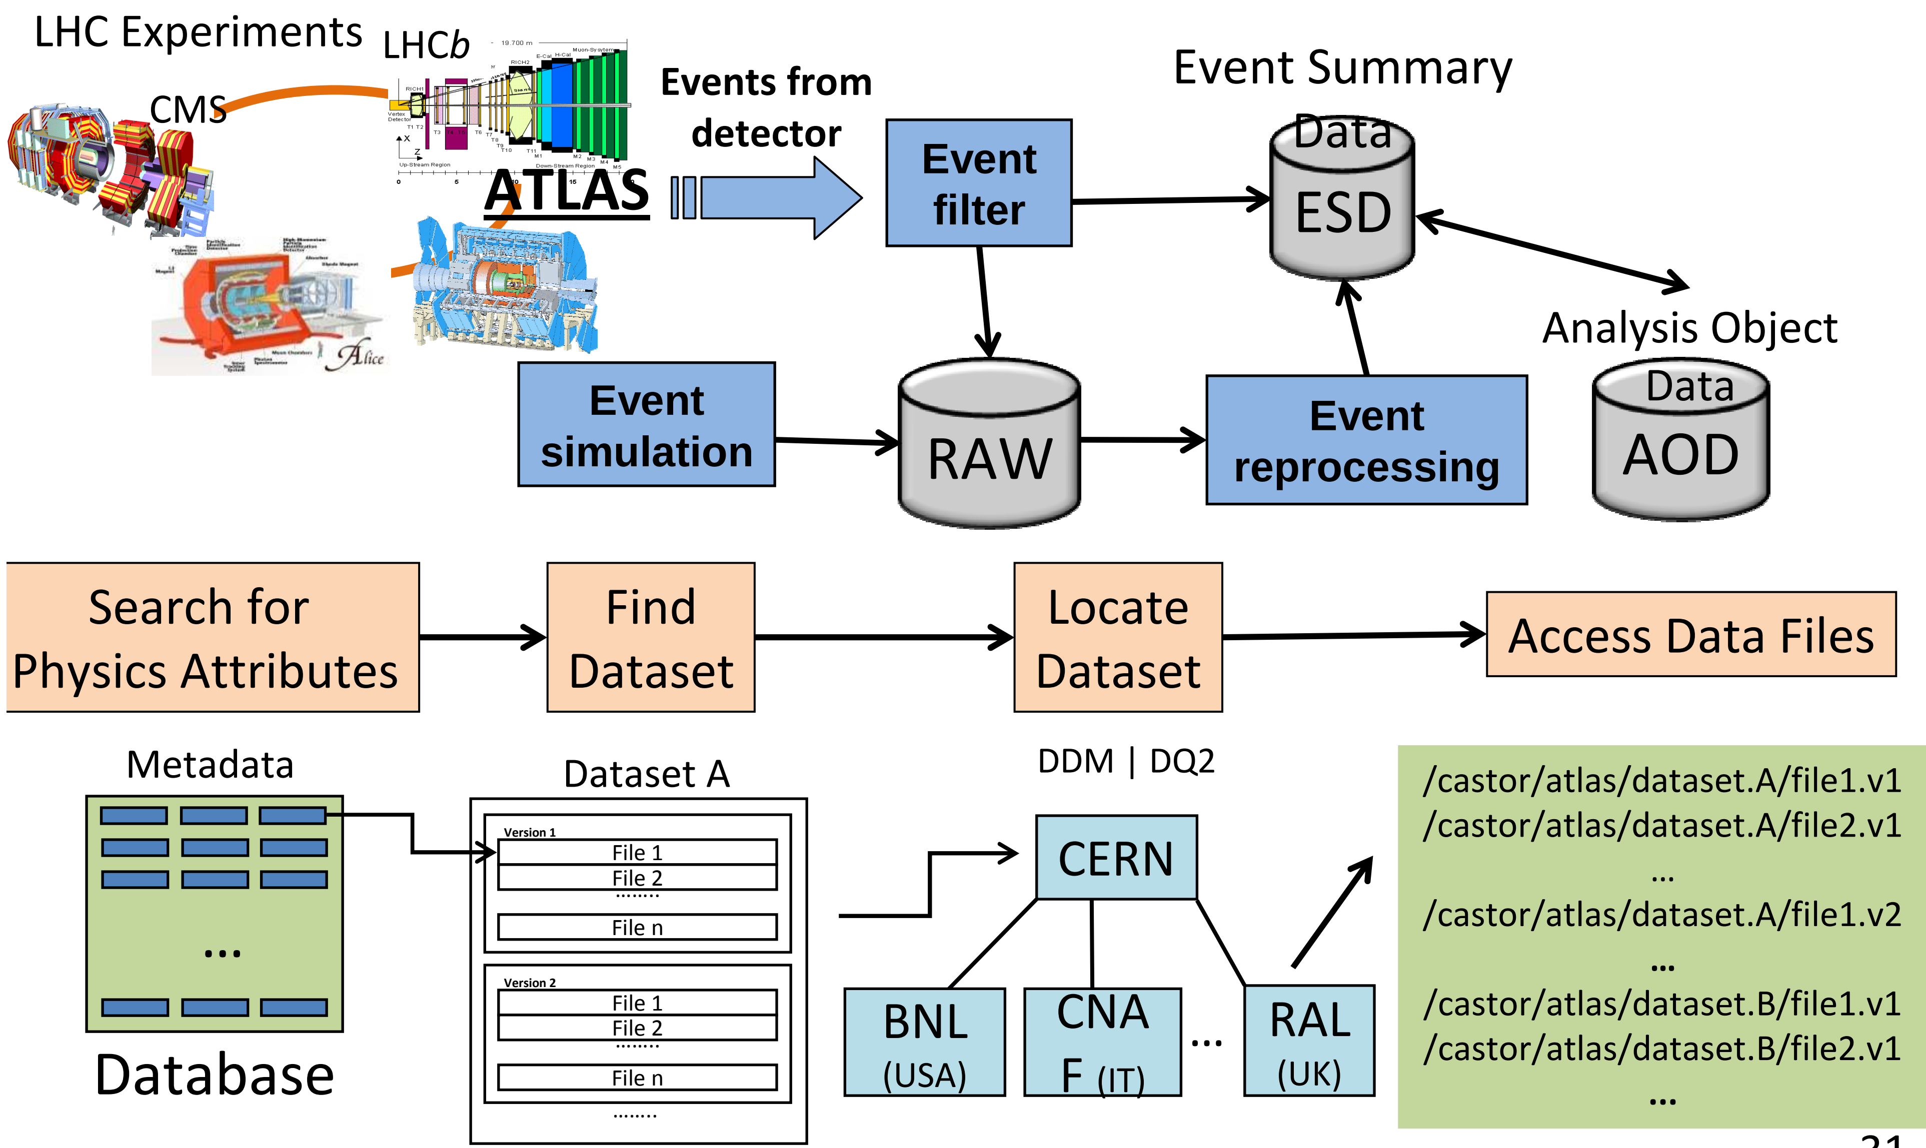

ATLAS-Amsterdam tier-1 re-processing data flow

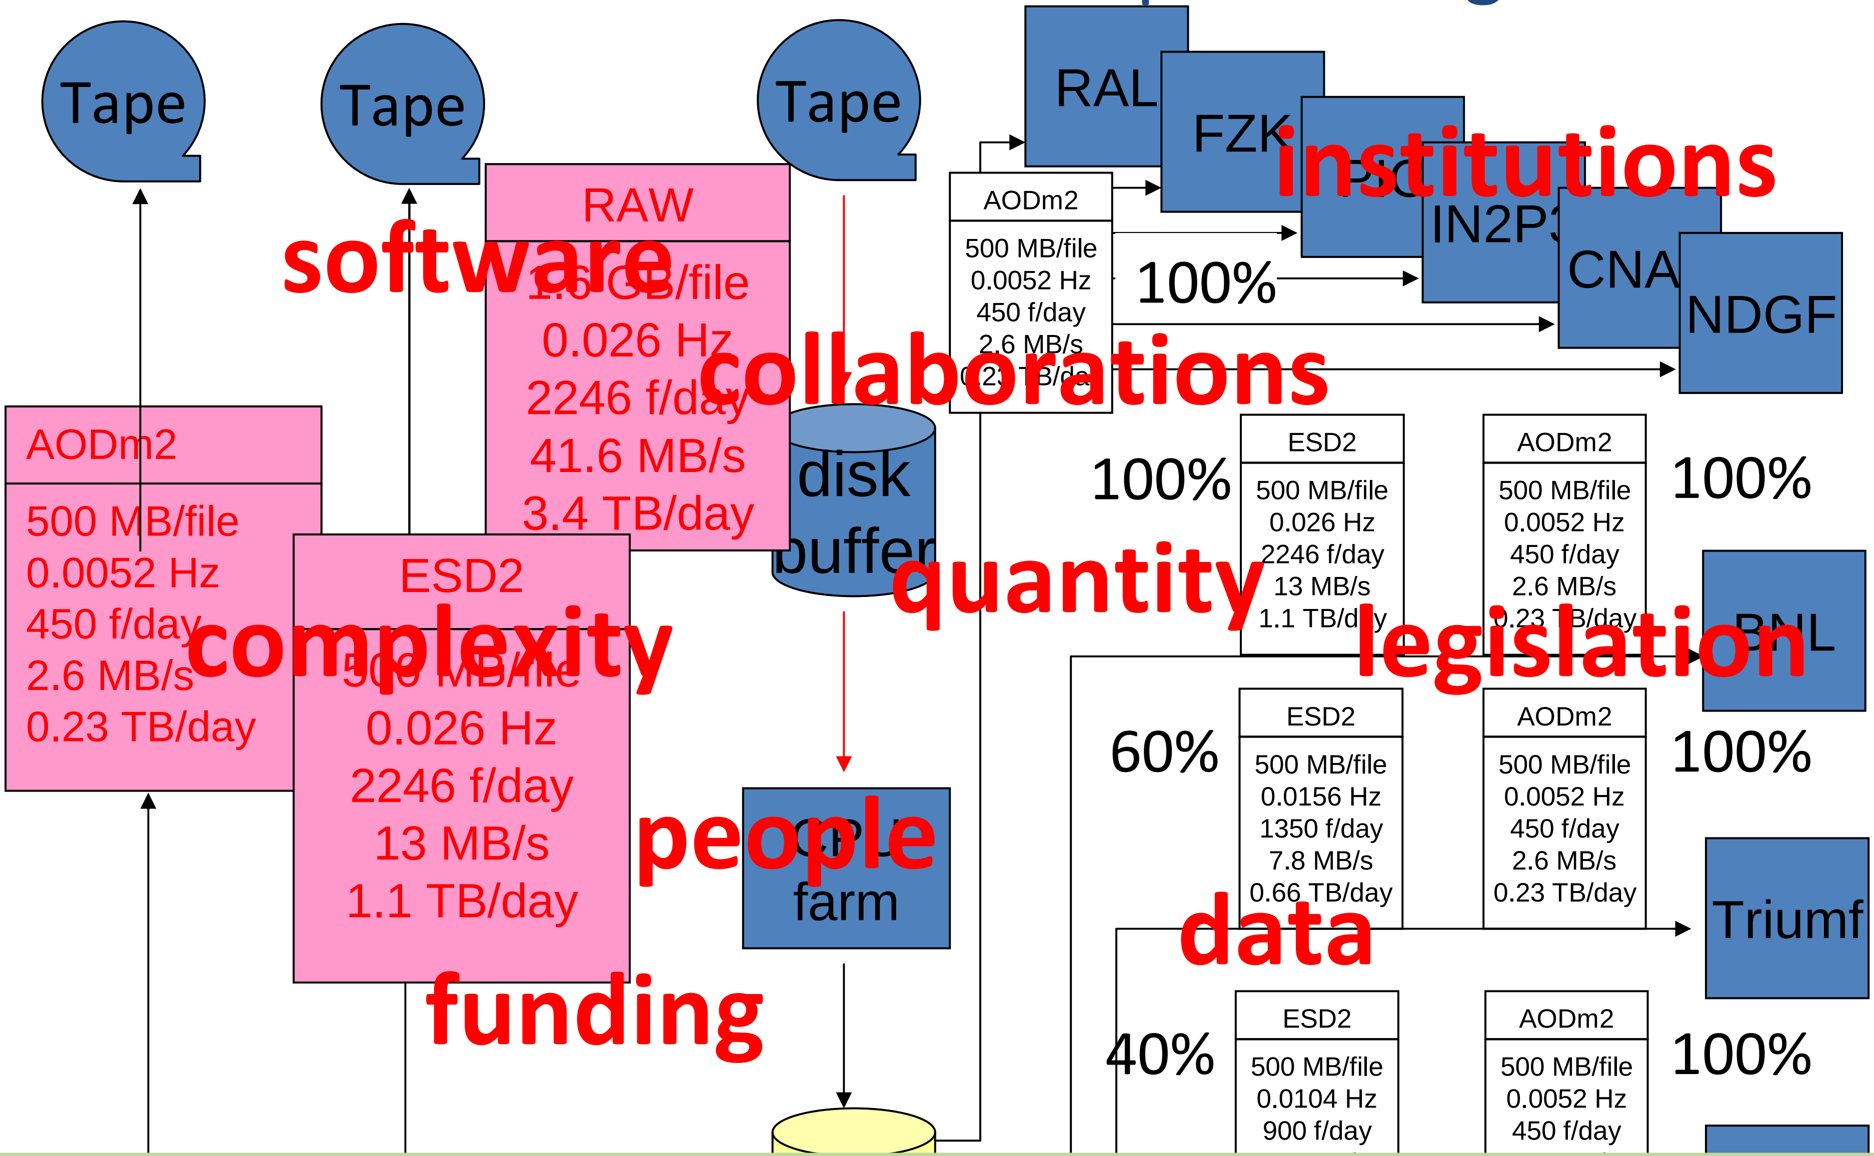

not your everyday transaction

DB: extract, transform, load, query

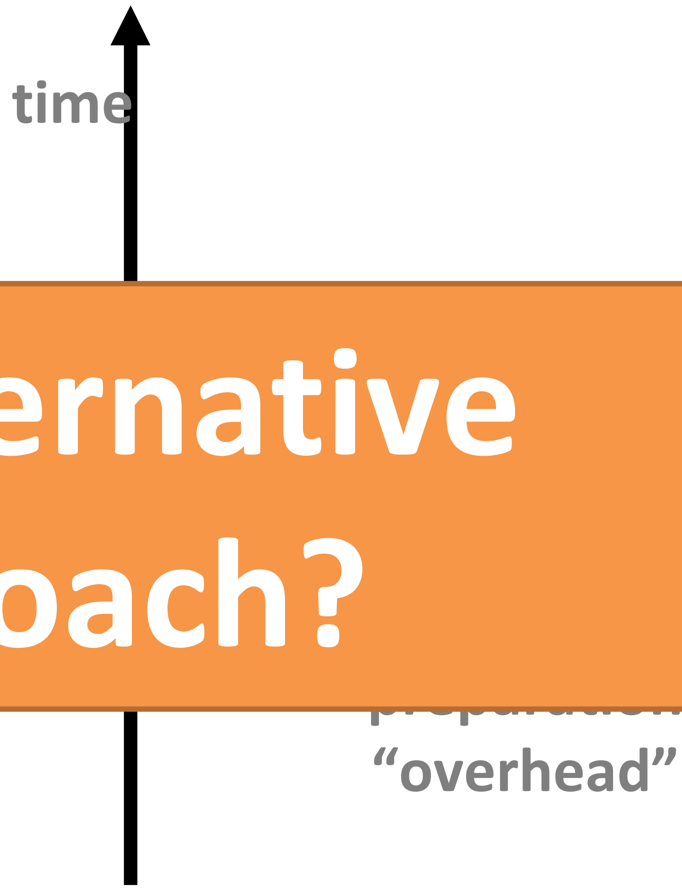

any alternative approach?

… lack of trust in database vendors
… databases “forever owning” the data

NoDB: query, query, query, query

Idea: Queries run **in-situ**, i.e., over raw data files

✓ Large collections of files

✓ Multiple data formats

✓ Integration with existing tools

✓ Changing areas of interest

New technology: fast queries on raw data

lessons learnt

- Be patient
  - Aim at solving the scientist’s problems, not ours
  - Go back to the lab and apply findings
- Many solutions necessary, one* does not fit all
  - *query language, data model, data type, index
- Open our minds to build *many* bridges to sciences

key to turn data into information

…an endless source of inspiration

Thank you!

Martha Anderson, Yannis Ioannidis, Miron Livny
Jim Gray, Alex Szalay, Randal Burns, Tanu Malik
Andy Connolly, Bob Nichol, Jeff Gardner
Gerd Heber, Dave O'Hallaron, Julio Lopez
Henry Markram, Felix Schuermann

…and to everyone who collects data