Data Intensive Computing

Paul Muzio

paul.muzio@csi.cuny.edu

1-718-982-3413

Agenda

- The City University of New York
- Data intensive computing
- Data curation

CUNY-Background

- Established as the Free Academy – 1847
  - Economically disadvantaged and those precluded from attending the leading universities of the day because of ethnicity or gender.
- Today
  - 24 institutions within the 300 sq. miles of NYC

- 280,000 students in degree programs
- 230,000 in non-degree programs
- 189 different languages spoken
- 42% first time college
- Alumni include:
  - 12 Nobel Laureates
  - 2 Fields Medal winners
  - Jonas Salk (polio vaccine)
  - Andy Grove (Intel Corp)
  - Robert Kahn (TCP/IP)
  - Many Pulitizer Prize Winners

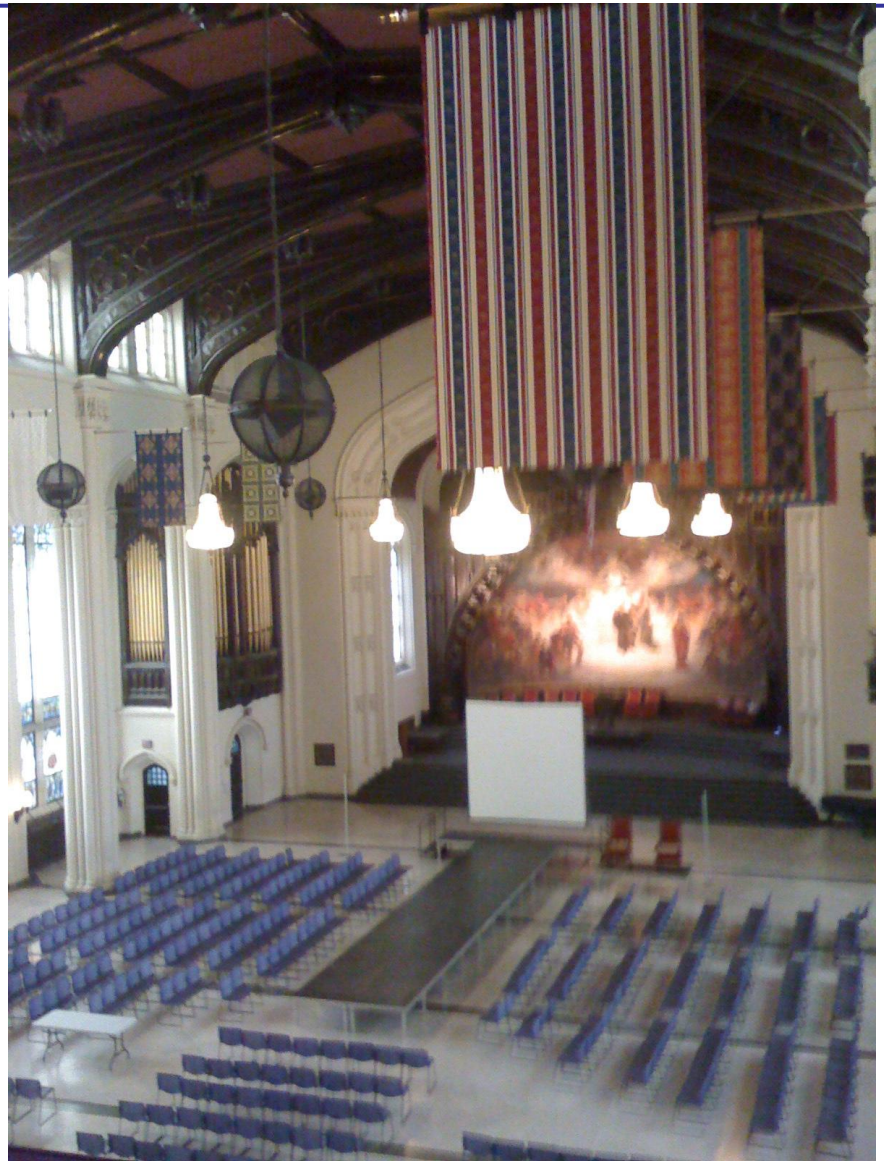

Data Intensive Computing

- Technology enables the collection of unprecedented quantities of data
  - Microprocessors
  - Sensor systems/laboratory equipment
  - Inexpensive memory
  - Inexpensive storage devices
- Examples
  - Large Hadron Collider
  - Environmental sciences
  - Genomics
  - Astronomy
  - etc.
- Increasingly, not just the physical sciences
  - Often non-numeric/semantic data

Data Intensive Computing

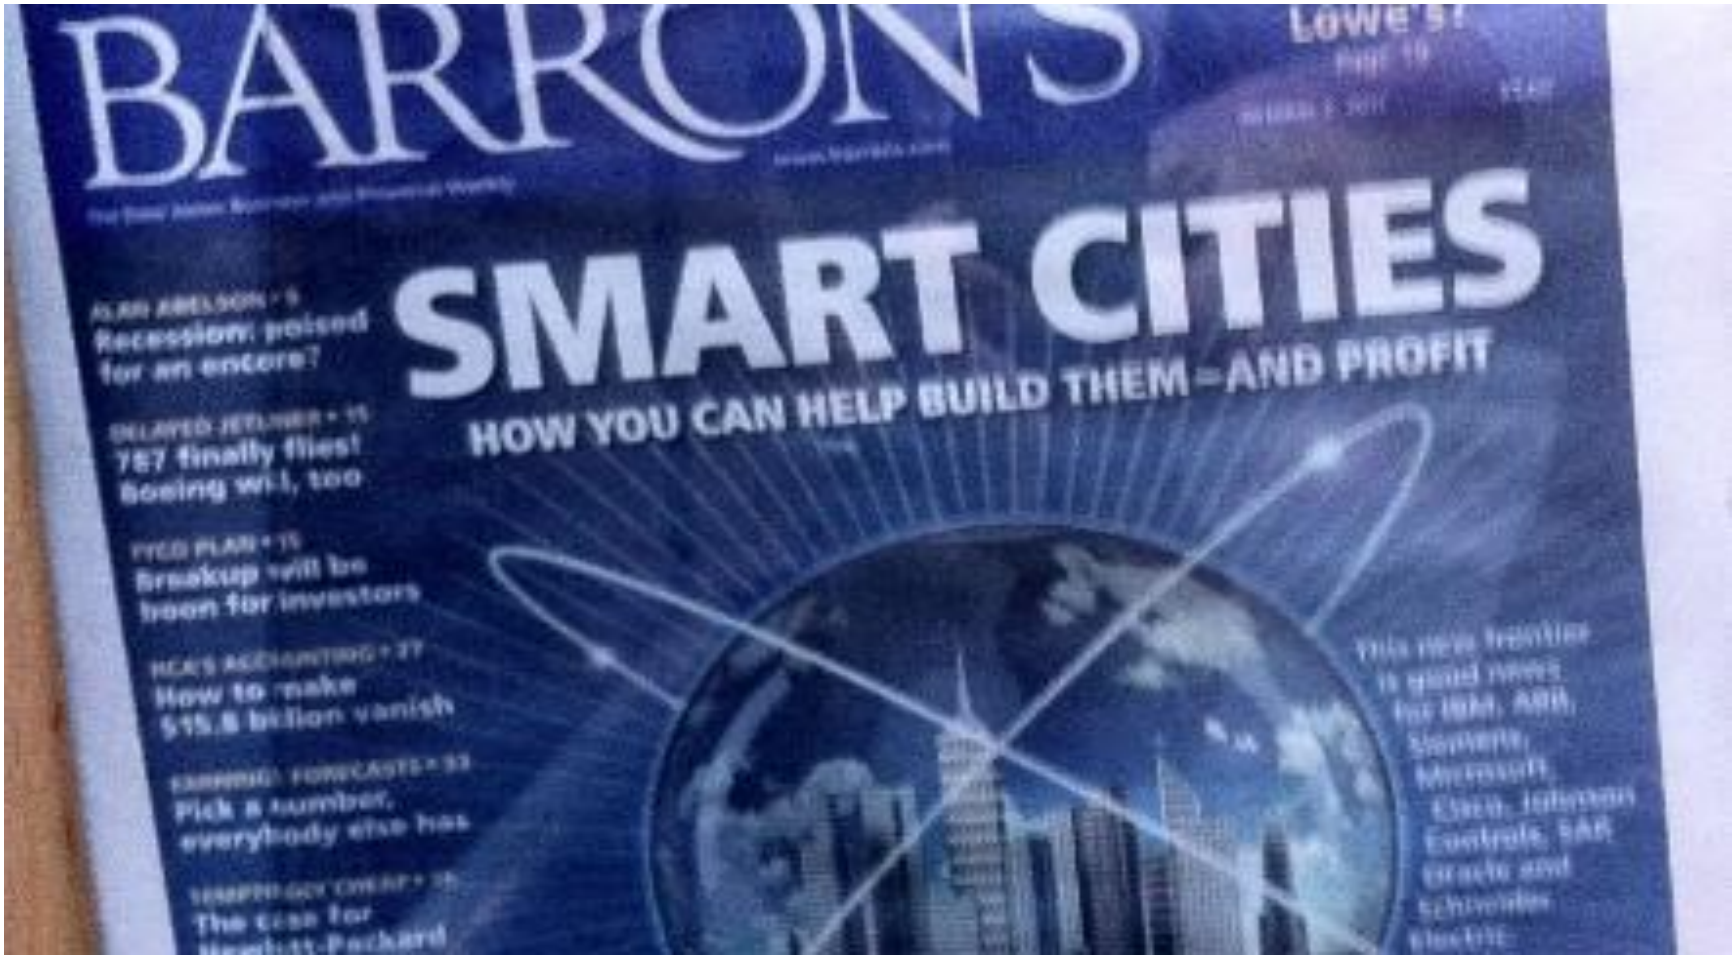

Data Intensive Computing

- Example: City of New York
  - GPS and fare tracking for 160 million taxi cab rides/year
  - Up to 2.3 billion subway/bus rides per year
  - Easypass toll tracking on tunnels/bridges, soon license plate tracking
  - Solar mapping of the entire city at 1 cm resolution, energy consumption for 1 million buildings
    - http://www.cuny.edu/about/resources/sustainability/solar-america/map.html
  - Air quality
  - Social services, health services
  - Criminal justice system
  - Real estate
  - Education
  - Etc.

“micro-city”

- Studies in autism and epilepsy
  - Psychology Department, College of Staten Island-CUNY
  - New York State Institute for Basic Research in Developmental Disabilities
- Corollary study
  - Understand behavior from the individual to societal level using “Naked African Mole Rats”
  - Controlled experiments are possible
  - Fill in the gap

Naked Mole Rats

- East Africa
- Mammal, closest relative is porcupine
- One of only 2 eusocial mammalian species (both African Mole rats)
- Hierarchical society
- Lives > 30 years
- Cancer resistant
- Poor vision, but may navigate using magnetic fields
- Poikilothermic
- Lives under extreme hypoxia in rainy season
- Each colony has a queen and breeding male
- Offspring will likely never reproduce (less than 5% chance)
- Males and females are barely distinguishable
- Communicate through chirps
- Have individual recognition

Courtesy: D. McCloskey, Psychology, College of Staten Island, CUNY

Society Structure

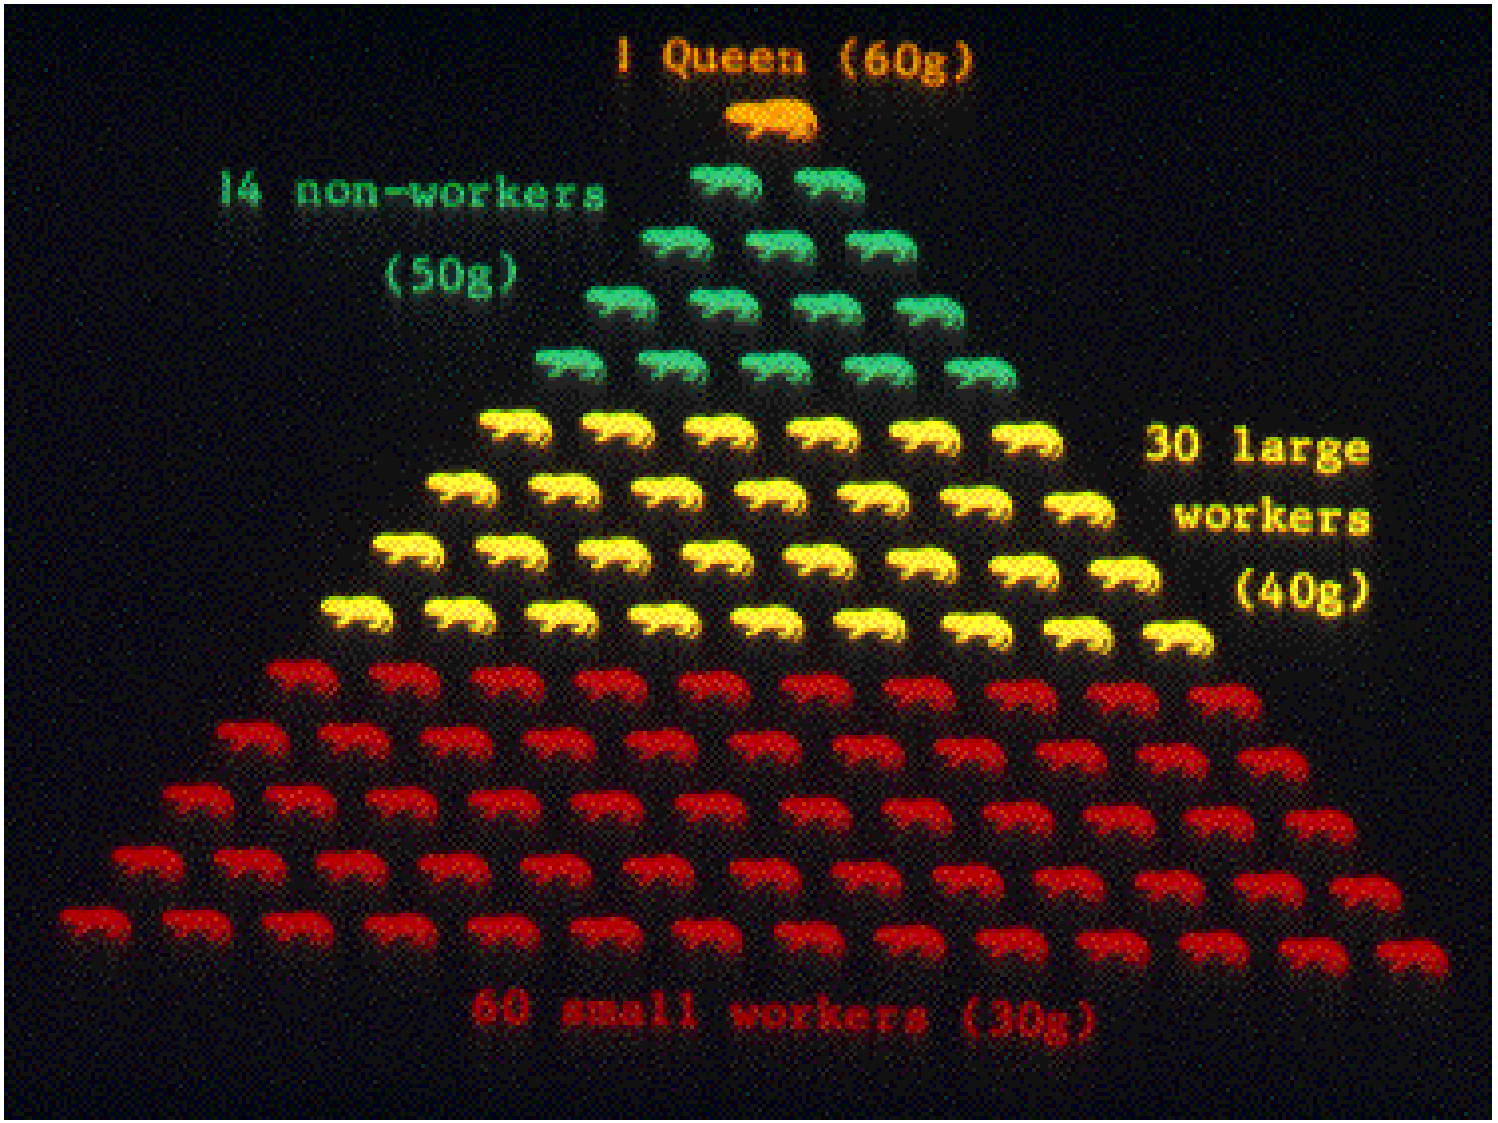

Mole Rat Laboratory

- RFIDs imbedded in each mole
  - Track location/interaction
- DNA
- Blood chemistry
- Brain structure
  - MRI
  - Biopsy
- Audio surveillance
- Video surveillance

"Social Interaction"

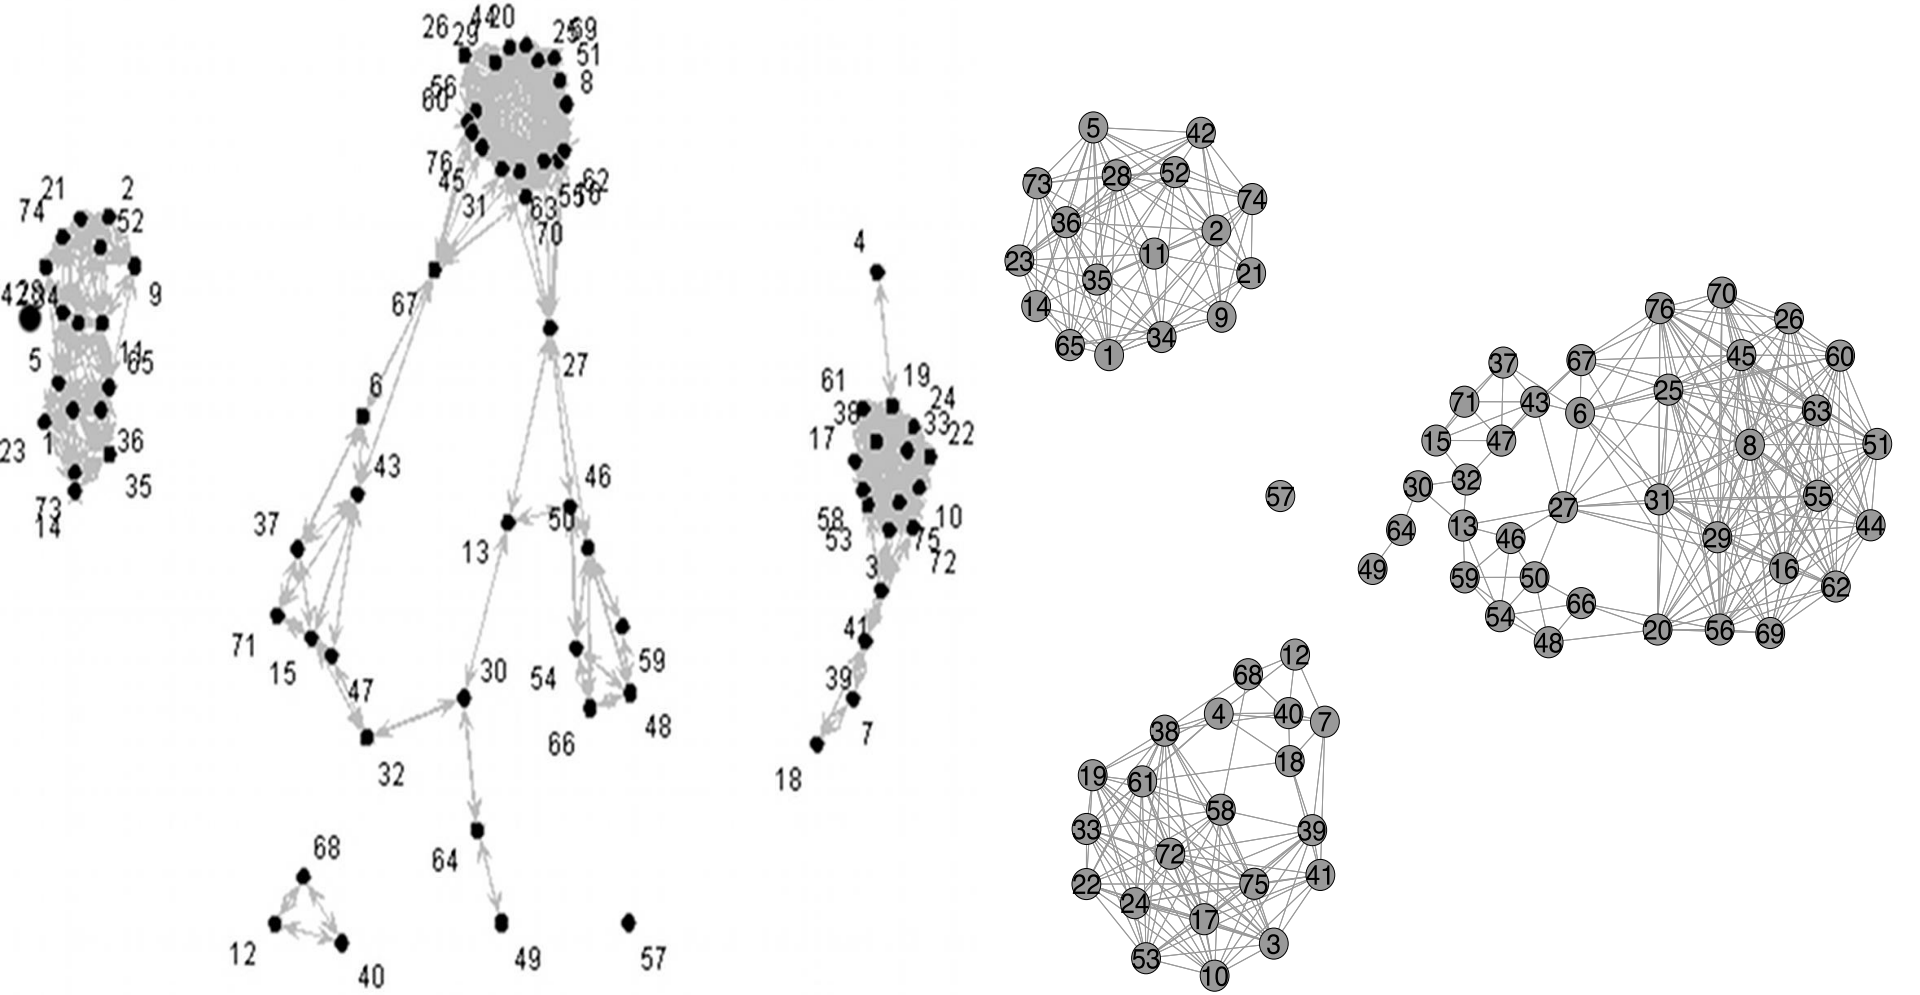

McCloskey, Kress, Imberman, Kushnir, Briffa-Mirabella (2011). *Proceedings of the 16th ACM SIGKDD* (In Press).

Social Networking

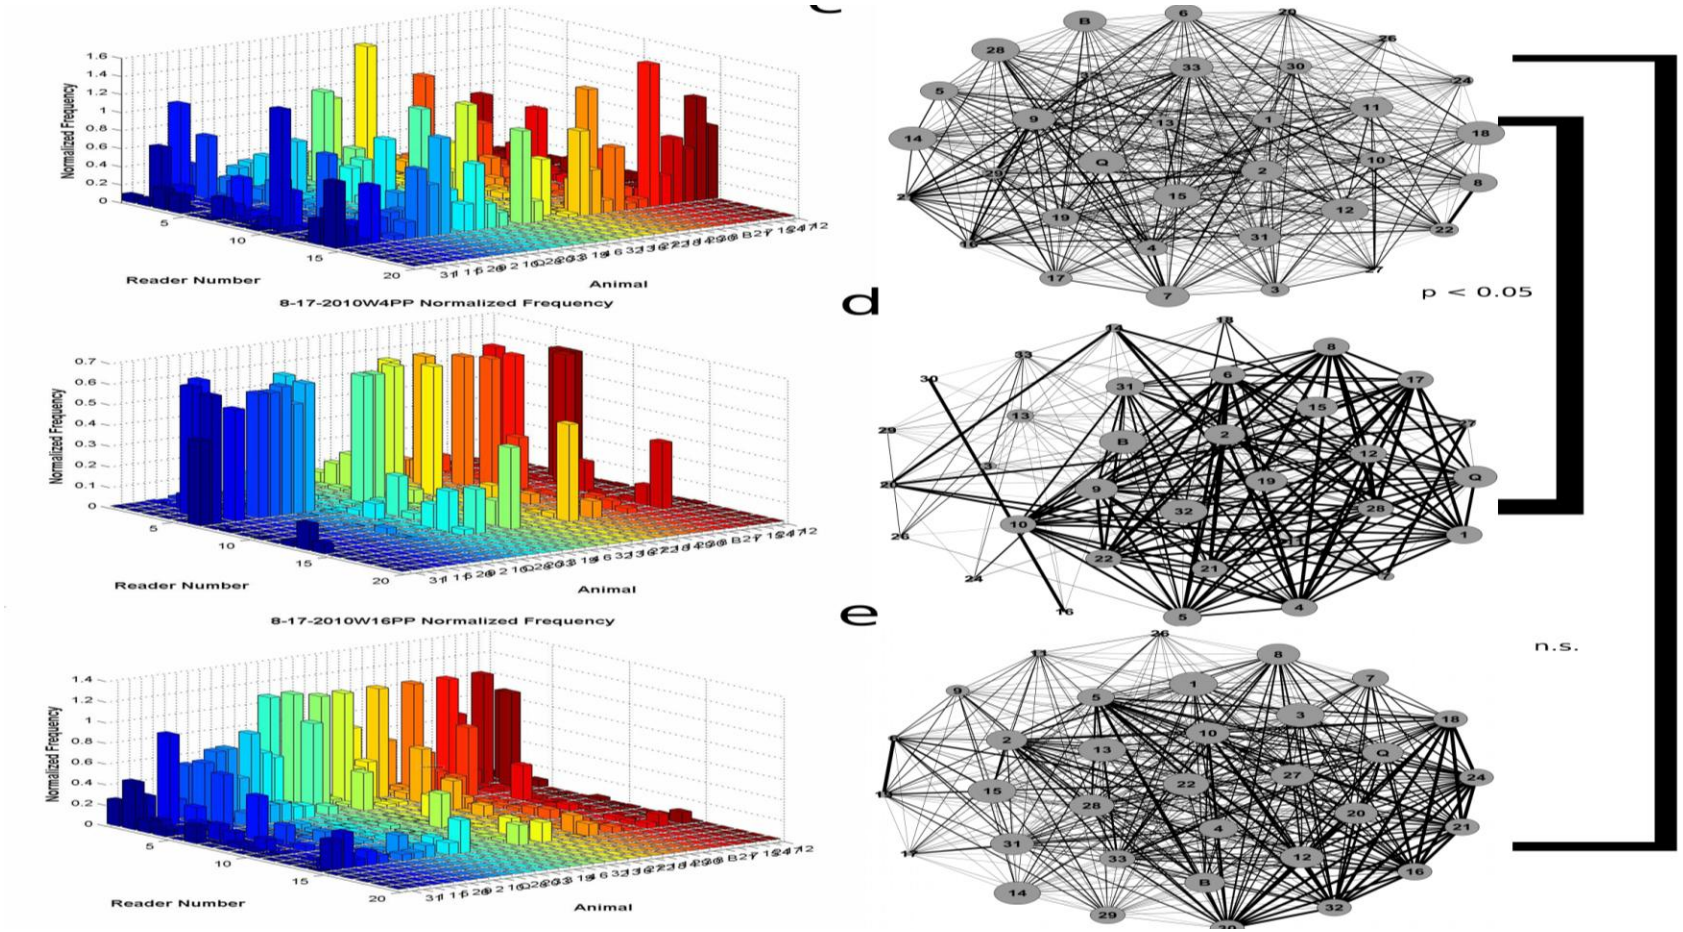

Hippocampal slice electrophysiology

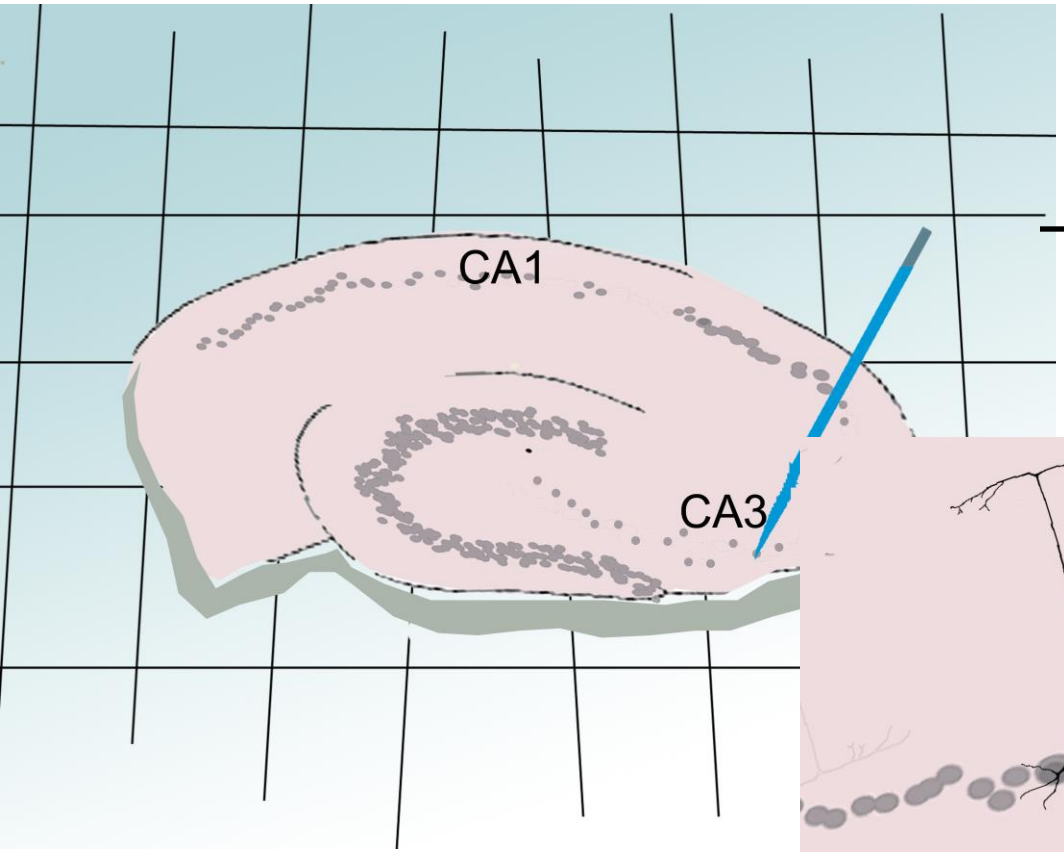

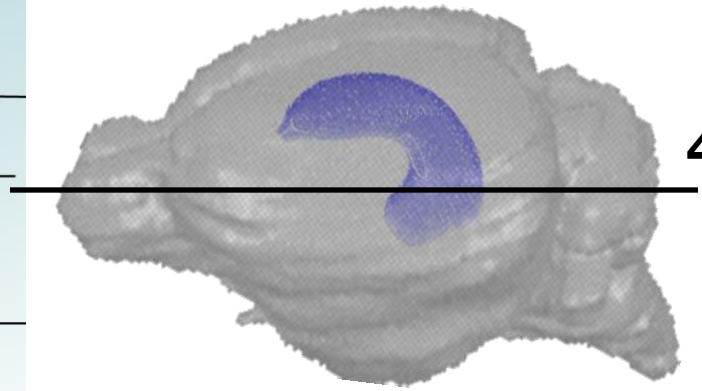

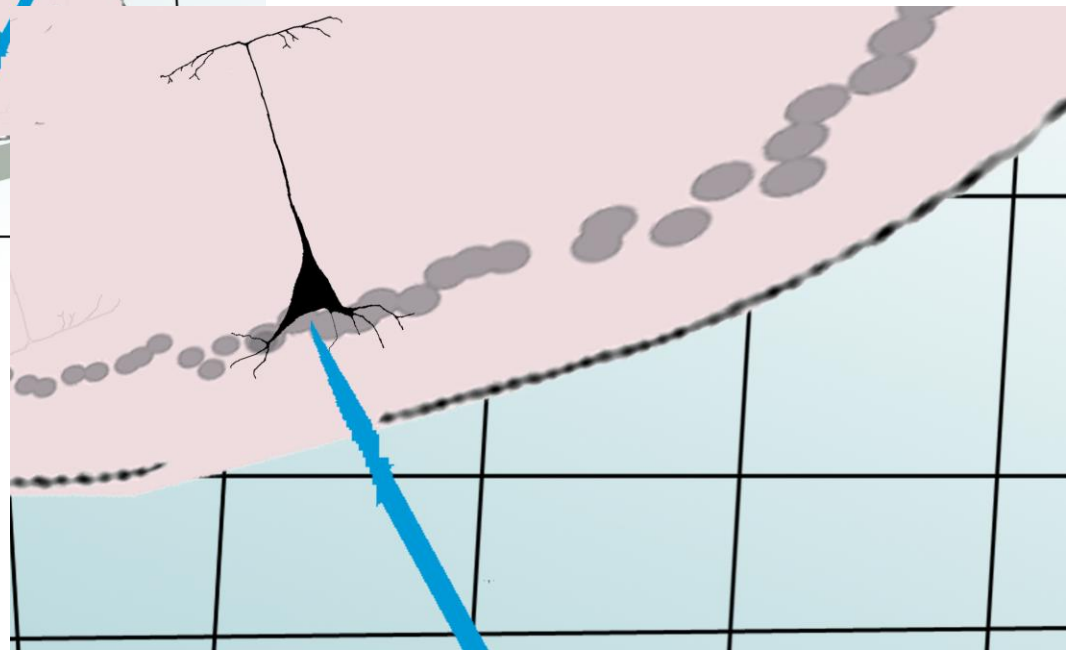

Epileptiform activity in the hippocampus

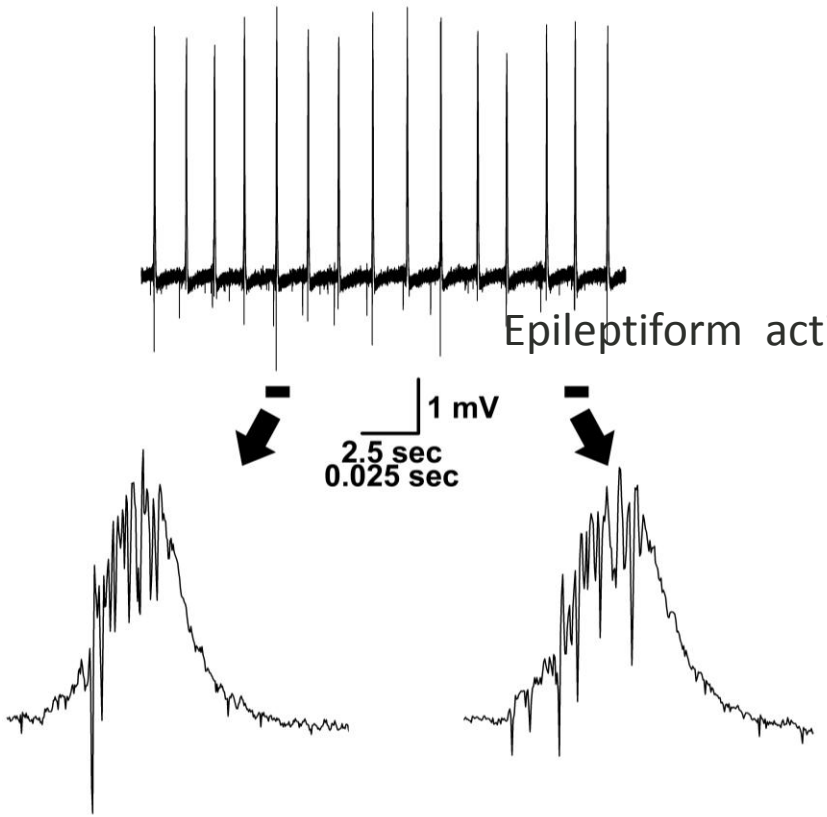

Epileptiform activity in the hippocampus

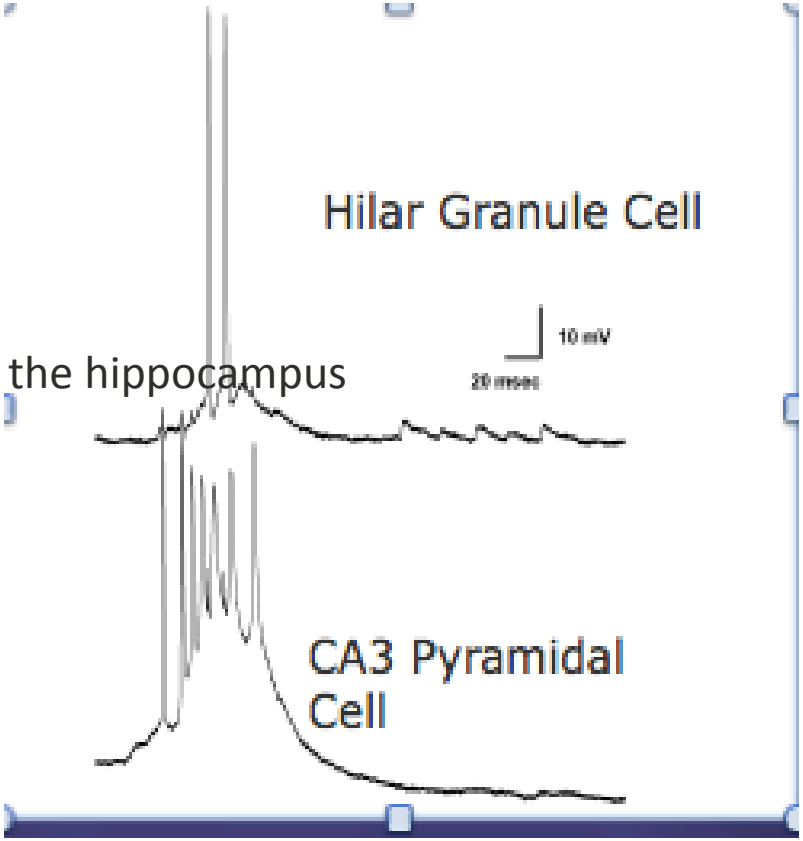

Issues

- Data collected is mainly semantic in nature
- Incomplete data
- Need to extract meaning from unstructured data (audio/video etc)
  - Entity extraction, automated tagging, text analytics
  - Natural Language Processing Technologies
- Limited software analysis tools
- Searching for relationships
- Graphs and combinatorics
- Large data, but also np-complete

Source: Andy Wilson, Sandia National Laboratory, 15 Sept 2010

LDRD
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

Do We Need New HPC Algorithms?

- Generally involves analyses of non-numeric data
  - Number of combinatorial possibilities grows rapidly.
  - Find in the data, meaningful relationships.
  - Convergence tests when not evaluating all possible combinations
  - Test for statistical significance--when the data is non-numeric?
- Application parallelism is abundant and asynchronous – e.g., concurrent graph search
- If possible, equally effective faster methods (ensembles, neural networks, etc.)
- Generally more about reasonable/ appropriate answers than statistical provability
- Often involves the use of incomplete data.
  - How does this affect the analysis process?

Do We Need New HPC Architectures?

- Build a Bigger Computer with Better Algorithms
- Problem is in N-space
  - Requires Global Machine Reach (in communication and naming)
  - Generally extremely sparse problems with very little local work
- Dense techniques will not work
  - There is no hope for accelerators, GPGPUs, etc. – Worry about how many GUPS your machine can sustain
- Even a PGAS Load/Store architecture is probably insufficient

Richard Murphy, Sandia National Laboratory, 15 September 2010

Do We Need New HPC Architectures?

- US Department of Energy
  - X-Caliber
  - Cray XMT
- Others?
  - Convey
  - ET International, Inc.
    - IBM Cyclops-64 chipset-160 cores/chip
    - ET Many-Core Software Systems parallel run-time execution model

Large-Scale Graph Analytics: Use Cases, Target Audiences, and Knowledge Discovery Toolbox (KDT) Technology

Aydin Buluc, LBL (abuluc@lbl.gov)
John Gilbert, Adam Lugowski and Drew Waranis, UCSB ({gilbert,alugowski,awaranis}@cs.ucsb.edu)
David Alber and **Steve Reinhardt**, Microsoft ({david.alber,steve.reinhardt}@microsoft.com)

Graph Analytics

- Knowledge Discovery Toolbox enables domain experts to deeply analyze large graphs quickly
- Graphs arise from
  - Social networks (human or animal)
  - Transaction networks (*e.g.,* Internet, banking)
  - Molecular biological interactions *(e.g.,* protein-protein interactions)
- Many queries are
  - Ranking
  - Clustering
  - Matching
- Graphs are not all the same
  - Directed simple graphs, hypergraphs, bipartite graphs, …

Courtesy: Steve Reinhardt, 15 September 2010

Use Cases

- Homeland security / Understand roles of members of terrorist groups based on known links between them / “Looking just at cell-phone communications, who are the leaders?”
- International banking / Detect money laundering / “Find instances of money being transferred at least 5 times and coming back to its source.”

- Common thread: Enabling the domain specialist to analyze graphs directly gets to the “right” answer faster and possibly at all

Courtesy: Steve Reinhardt, 15 September 2010

KDT Overview

- Target audiences
  - Primarily, (non-graph-expert) domain specialists needing to analyze large graphs
  - Secondarily, graph-algorithm researchers and developers needing access to highly performant scalable graph infrastructure
- Target use cases
  - Broadly, problems needing the detail of algorithms that traverse the graph extensively
  - Social-network-based ranking and search
  - Homeland security
- Current practicalities
  - Python interface layered on Combinatorial BLAS
    - Delivers full scaling of CombBLAS with negligible Python overhead
  - v0.2 release expected in August
  - Open-source code available at `kdt.sourceforge.net` under New BSD license

Courtesy: Steve Reinhardt, 15 September 2010

Large Graphs / Hairballs

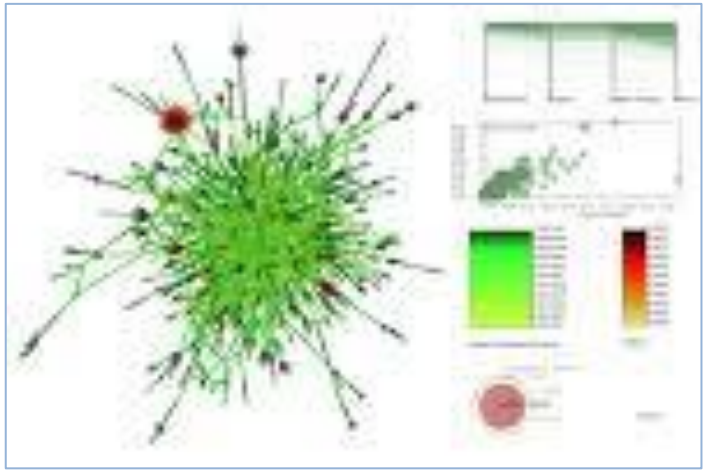

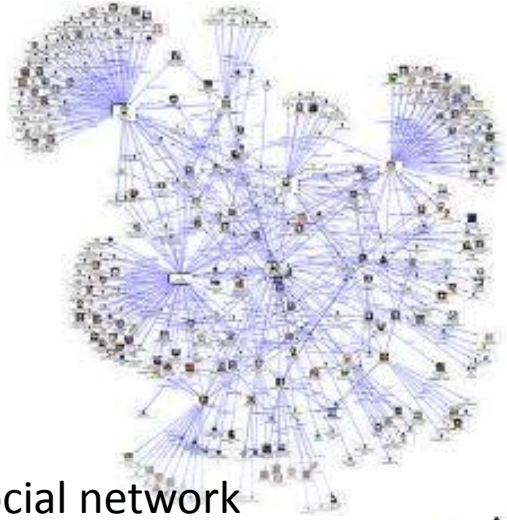

Social network

fmsasg.com

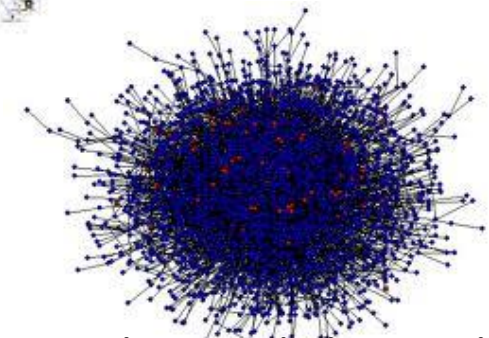

Protein-protein interaction network estradalab.org

Wikipedia cross-links

urbagram.net

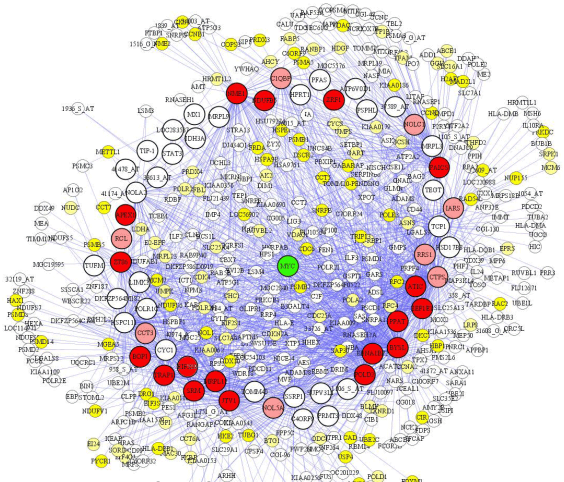

metabolic network

di.unipi.it

Untangling the Hairball

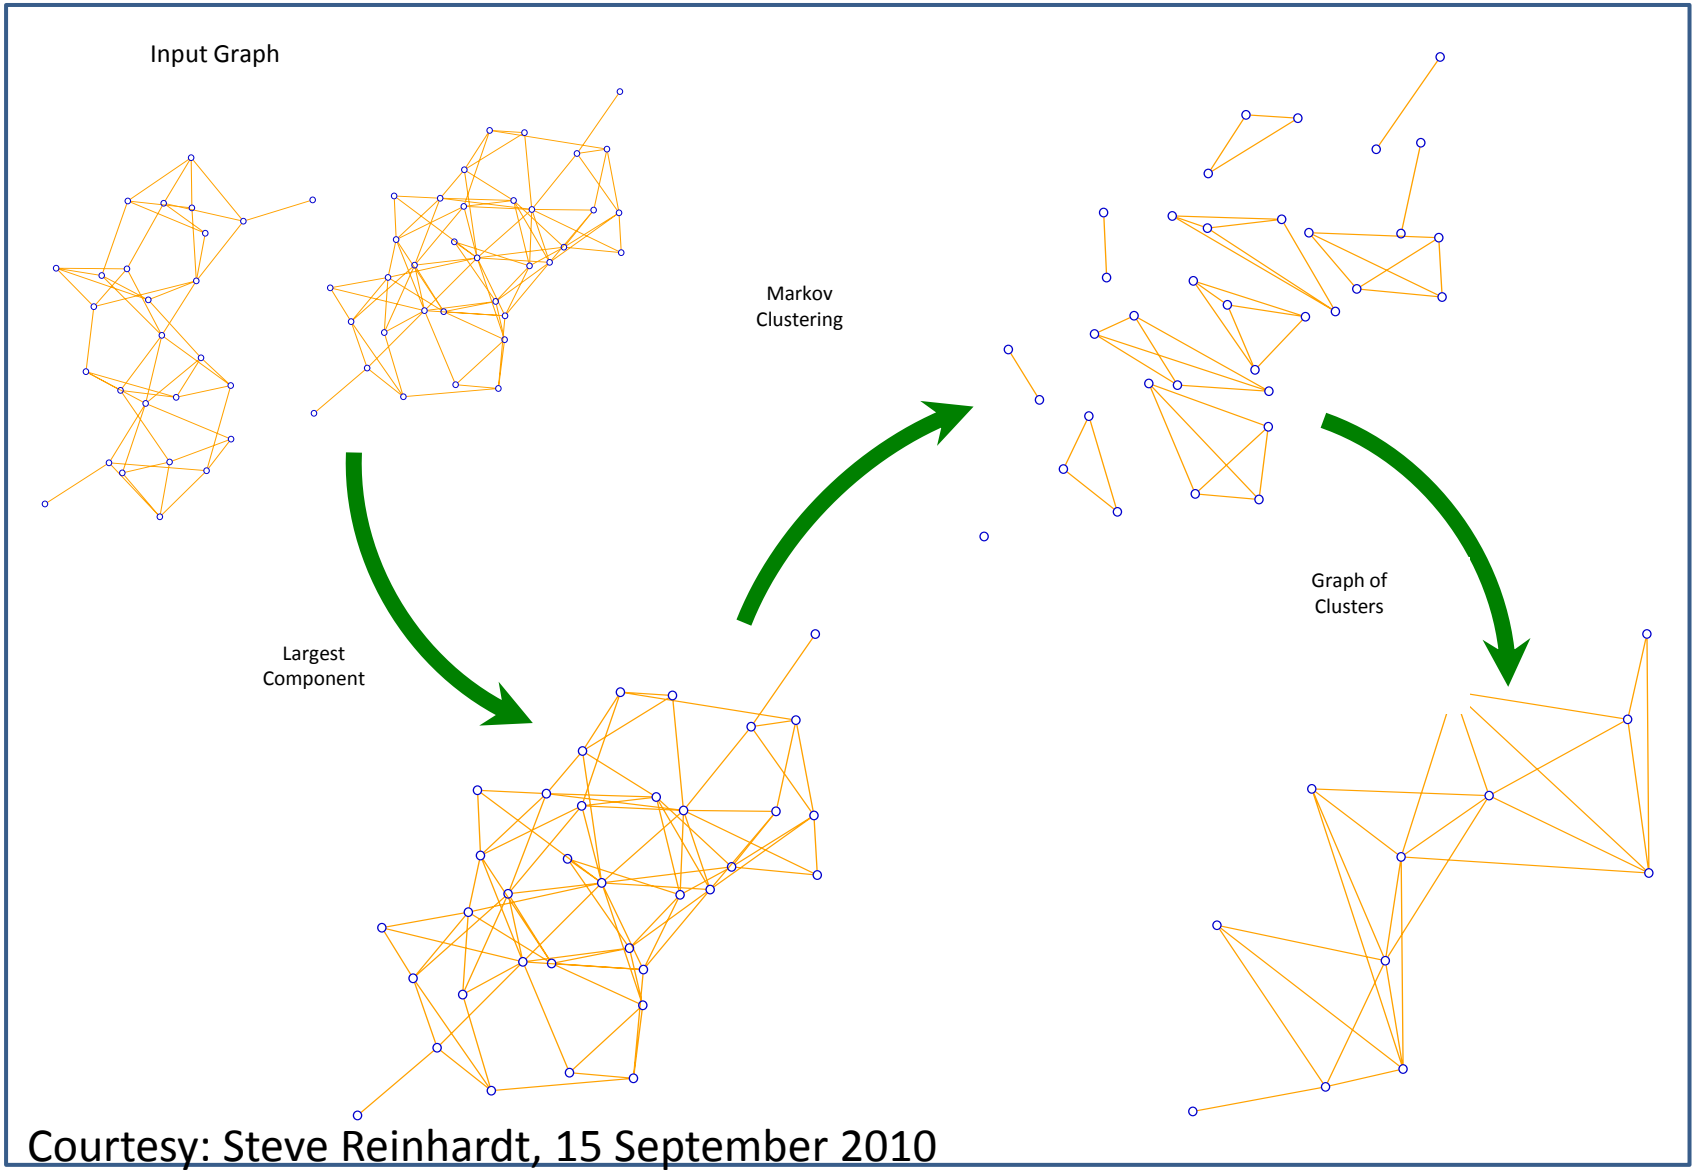

Courtesy: Steve Reinhardt, 15 September 2010

KDT Use Case: Find Influential People in a Social Network

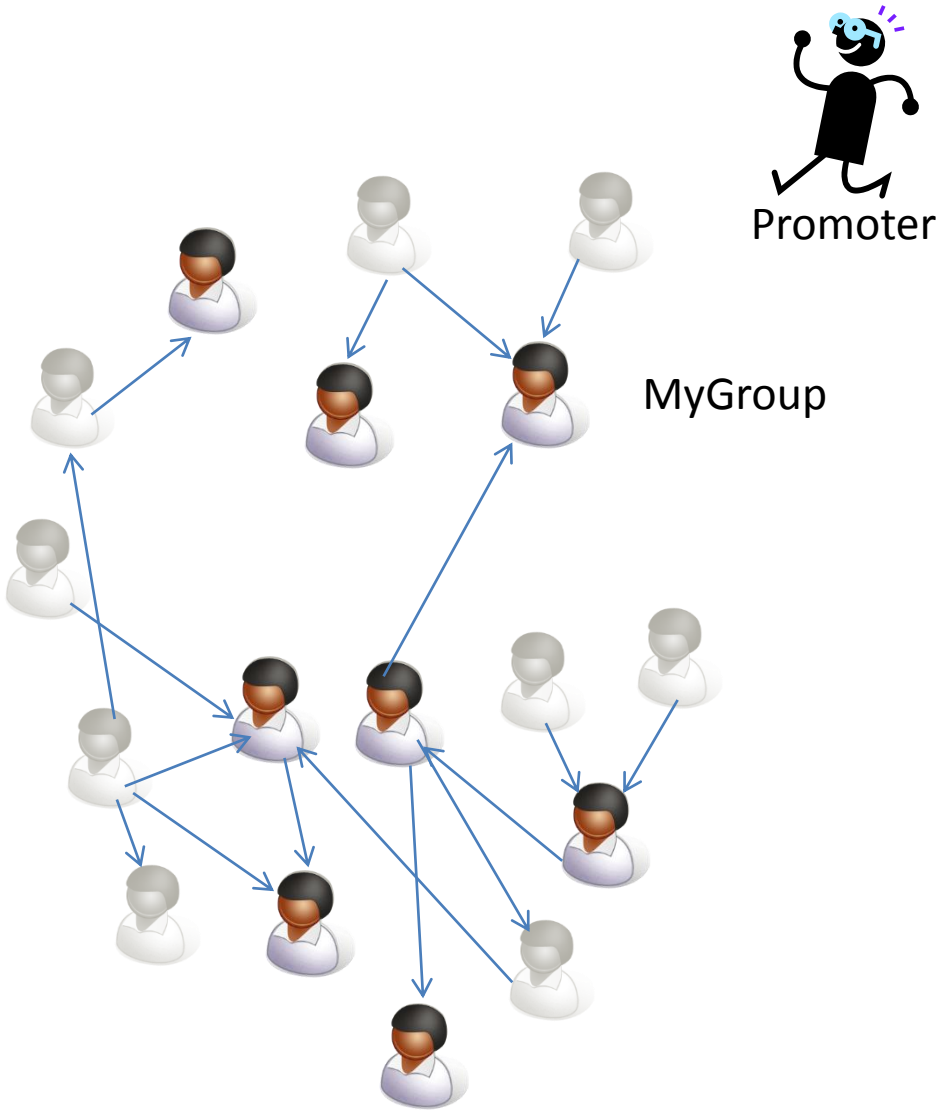

- Promoter has a SN group
- Wants to identify influencers on which to focus marketing efforts so as to maximize viral effect of the group
- Calls KDT with group name, gets back top N influencers
- Useful for (*e.g.*) viral marketing, public health

Courtesy: Steve Reinhardt, 15 September 2010

Comparison to Other Parallel Packages

Package	Target users		Interface	Supported memory*
	Graph-alg devs	Domain experts		
Pegasus	X		Hadoop	Distributed on-disk
Pregel	X		C++	Distributed on-disk
PBGL	X		C++	Distributed in-memory
MTGL	X		C++	Shared
SNAP (GA Tech)	X		C	Shared
SNAP (Stanford)	X	X	C++ / NodeXL	Shared
GraphLab	X		C++	Shared
CombBLAS	X		C++	Shared or distributed, in-memory
KDT	X	**X**	Python	Shared or **distributed**, in-memory

*"Shared" meaning either cache-coherent or Cray XMT-style

Courtesy: Steve Reinhardt, 15 September 2010

Local v. Global Metrics

Degree Centrality v. Betweenness Centrality

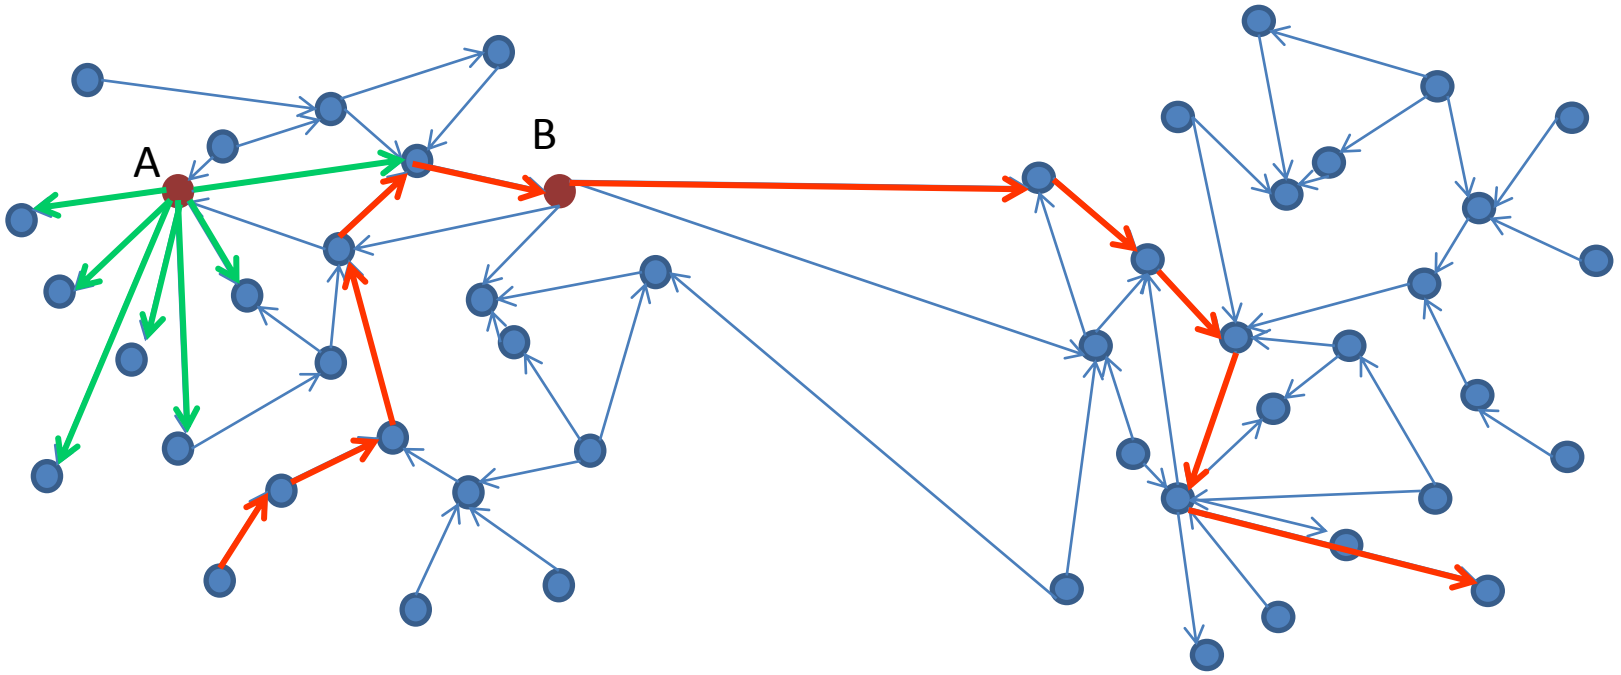

- Is vertex A or B most central?
  - A has directed edges to more vertices (degree centrality)
  - B is on more shortest paths between vertex pairs (betweenness centrality)

Courtesy: Steve Reinhardt, 15 September 2010

The case for sparse matrices

Many irregular applications contain coarse-grained parallelism that can be exploited by abstractions at the proper level.

Traditional graph computations	Graphs in the language of linear algebra
Data driven, unpredictable communication.	Fixed communication patterns
Irregular and unstructured, poor locality of reference	Operations on matrix blocks exploit memory hierarchy
Fine grained data accesses, dominated by latency	Coarse grained parallelism, bandwidth limited

UCSB

Performance

Graph500 in KDT or Combinatorial BLAS

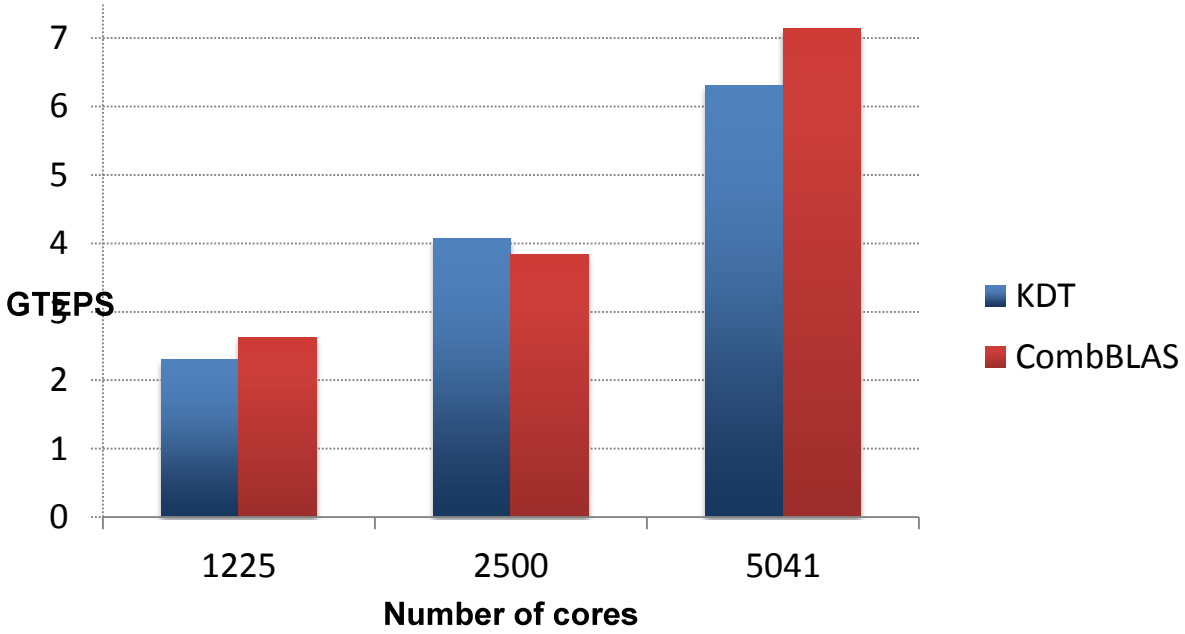

- Graph500 benchmark on 8B edges, C++ or KDT calling CombBLAS
- NERSC “Hopper” machine (Cray XE6)
- [Buluç & Madduri]: New hybrid of CombBLAS MPI + OpenMP gets 25 GTEPS on 2T edges (scale 37) on 43,200 cores of Hopper

Courtesy: Steve Reinhardt, 15 September 2010

Performance of Betweenness Centrality

- With a few hundred cores, can do even a complex graph analysis in near-interactive time
- 2M edges, approximate betweenness centrality sampling at 3%

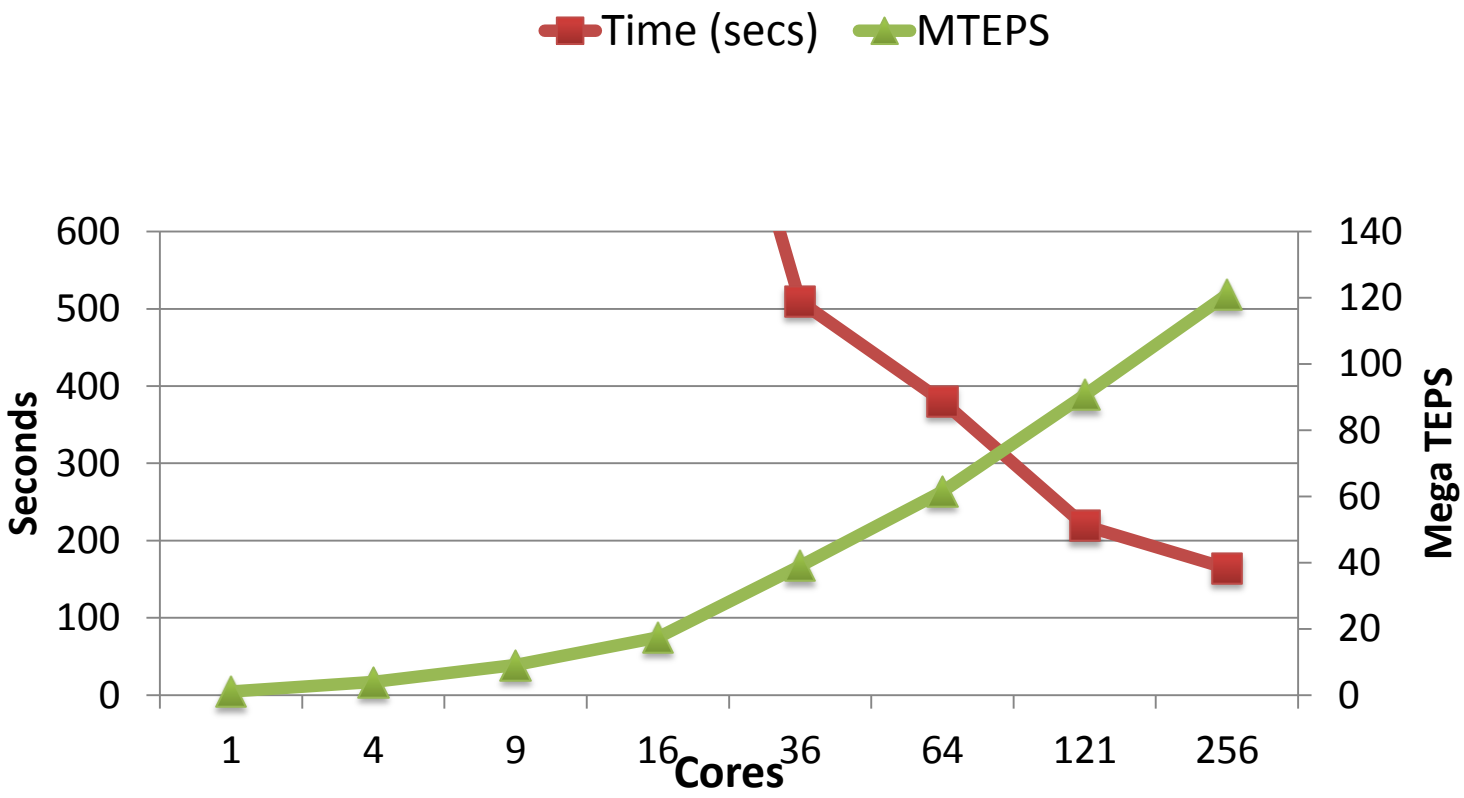

Courtesy: Steve Reinhardt, 15 September 2010

KDT Summary

- Open-source toolbox targeted at domain specialists
- Scalable to 10B-edge graphs and thousands of cores
- Limited set of methods, no graph viz yet
- kdt.sourceforge.net for details

Courtesy: Steve Reinhardt, 15 September 2010

Data Curation

- National Science Foundation requirement for all new proposals
- Types of data, samples, physical collections, software, curriculum materials, and other materials to be produced in the course of the project;
- Standards to be used for data and metadata format and content (where existing standards are absent or deemed inadequate, this should be documented along with any proposed solutions or remedies);
- Policies for access and sharing including provisions for appropriate protection of privacy, confidentiality, security, intellectual property, or other rights or requirements;
- Policies and provisions for re-use, re-distribution, and the production of derivatives; and
- Plans for archiving data, samples, and other research products, and for preservation of access to them.

Thank you