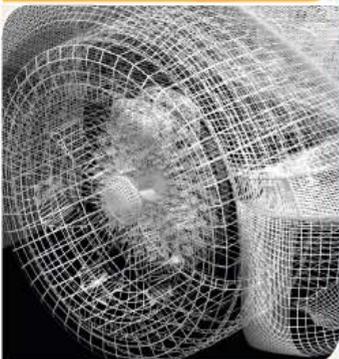


# Achieving sustainability through exascale analysis of complex biological networks

David A. Bader



**Georgia  
Tech**



College of  
Computing

Computational Science and Engineering

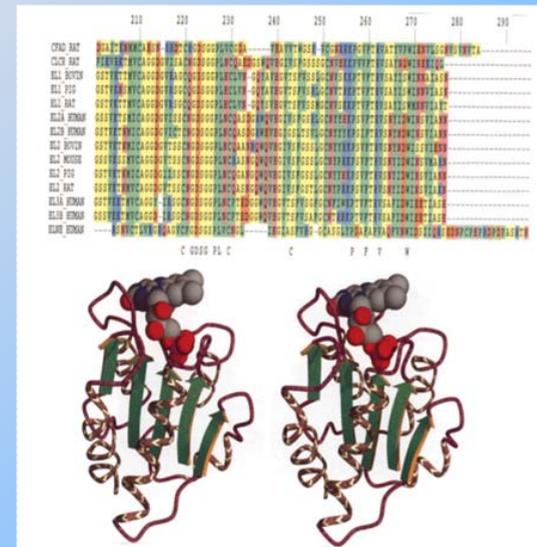
# NSF Workshop on Petascale Computing in the Biological Sciences



Co-organized by David A. Bader (Georgia Tech), Allan Snavey, and Gwen Jacobs, this workshop held in August, 2006:

- Jumpstarted collaborations between computer scientists and computational biological scientists thereby enabling efficient and early use of petascale computing resources as they come online.
- Promoted meaningful information exchange between the two domains; it was structured as a tutorial with bioscientists describing their applications along with the computational challenges they present and computer scientists describing their methods and tools for improving application performance and capability.

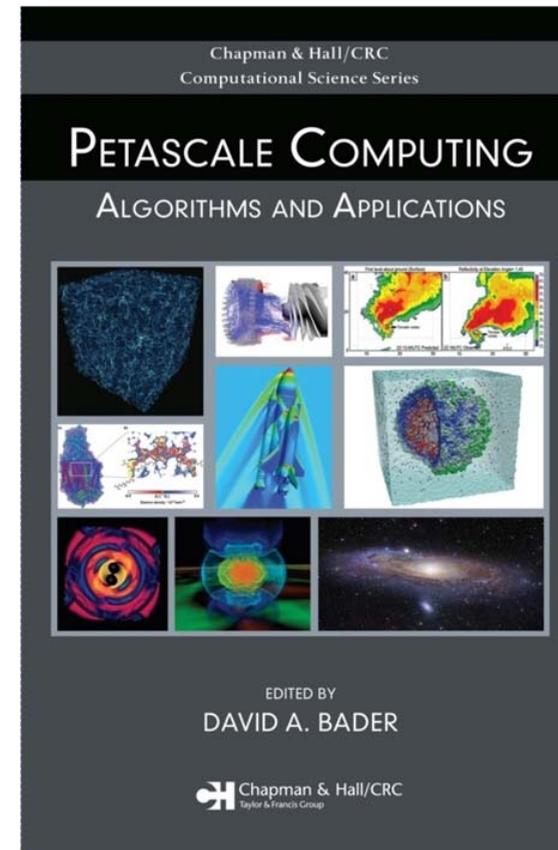
## *NSF Workshop Report on Petascale Computing in the Biological Sciences*



# Petascale Computing: Algorithms and Applications (David A. Bader)



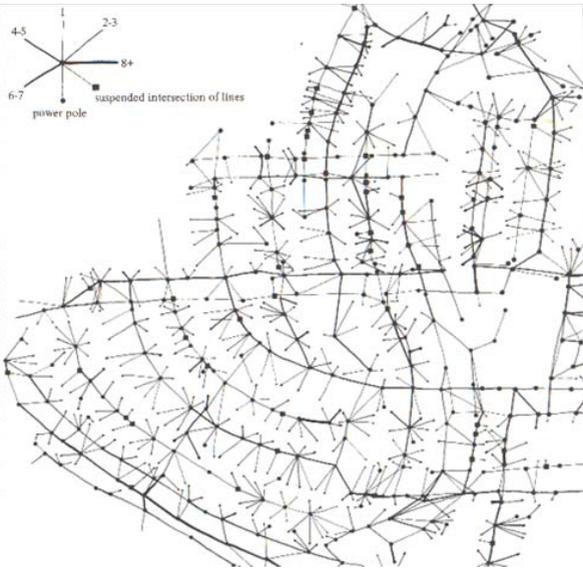
- Provides the first collection of articles on petascale algorithms and applications for computational science and engineering
- Covers a breadth of topics in petascale computing, including architectures, software, programming methodologies, tools, scalable algorithms, performance evaluation, and application development
- Discusses expected breakthroughs in the field for computational science and engineering
- Includes contributions from international researchers who are pioneers in designing applications for petascale computing systems



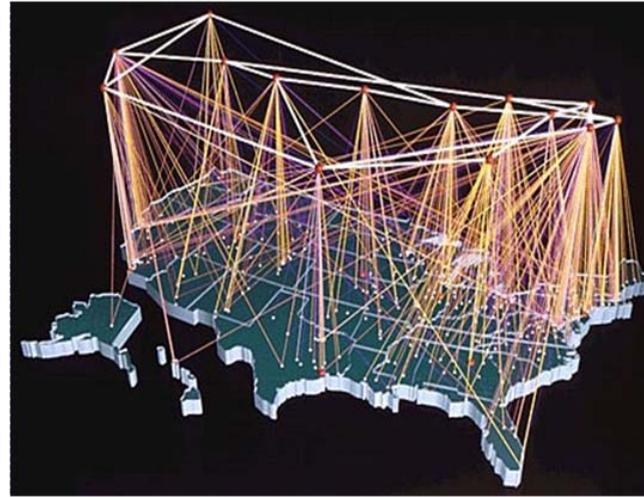
Chapman & Hall/CRC Computational Science  
Series, © 2007

# Complex Networks arise from a variety of sources

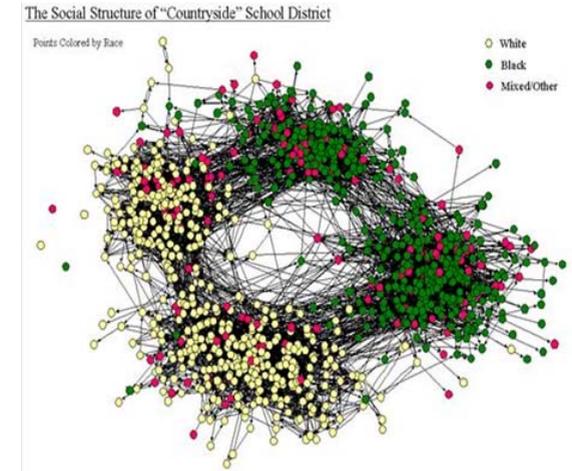
Power Distribution Networks



Internet backbone



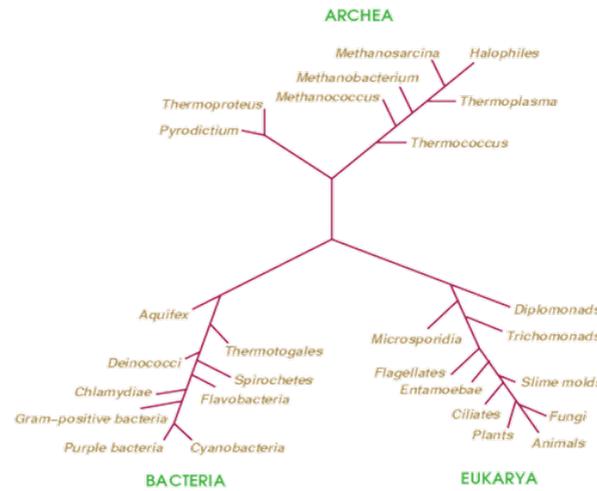
Social Networks



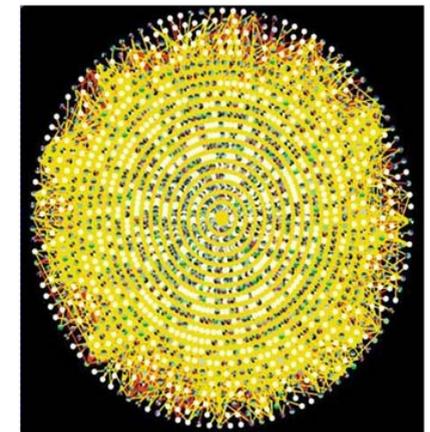
Graphs are everywhere!



Ground Transportation



Tree of Life



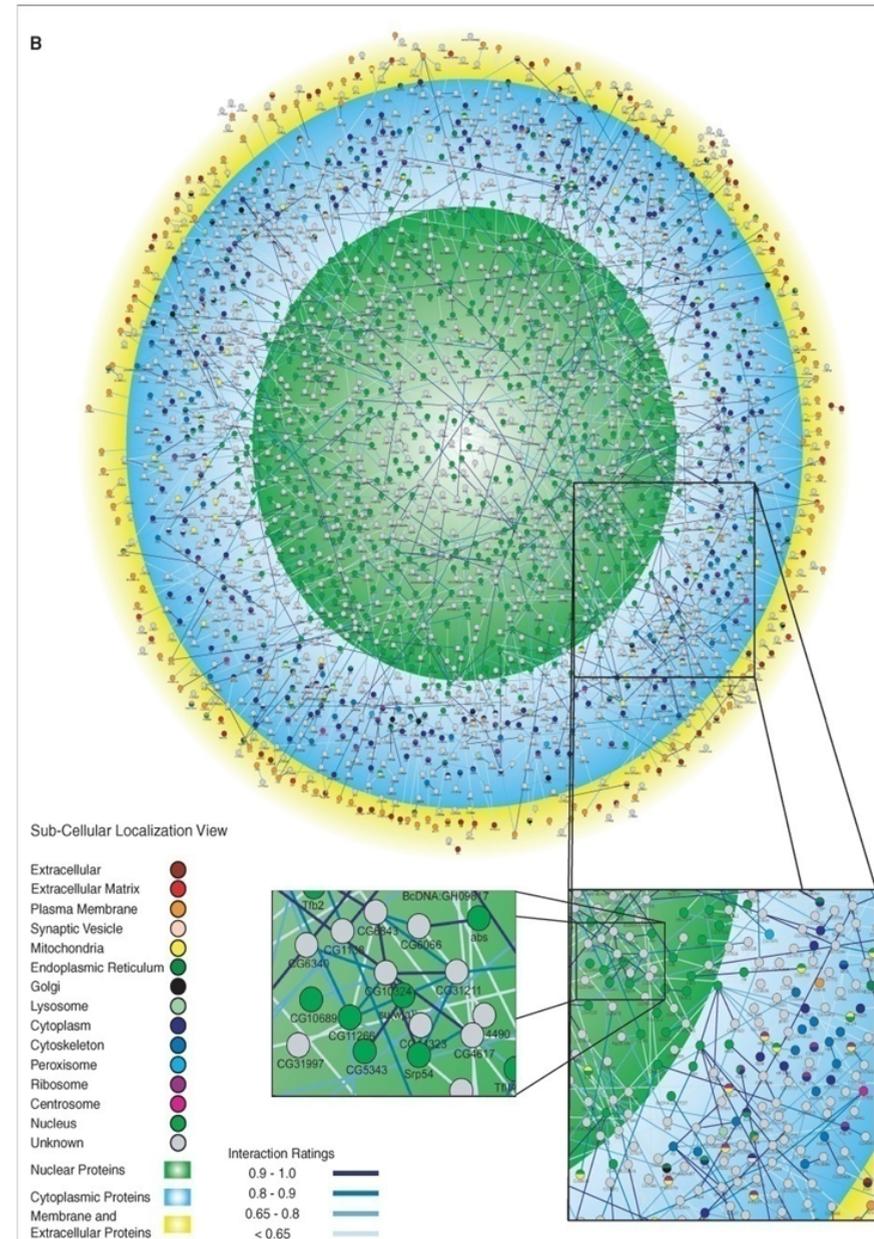
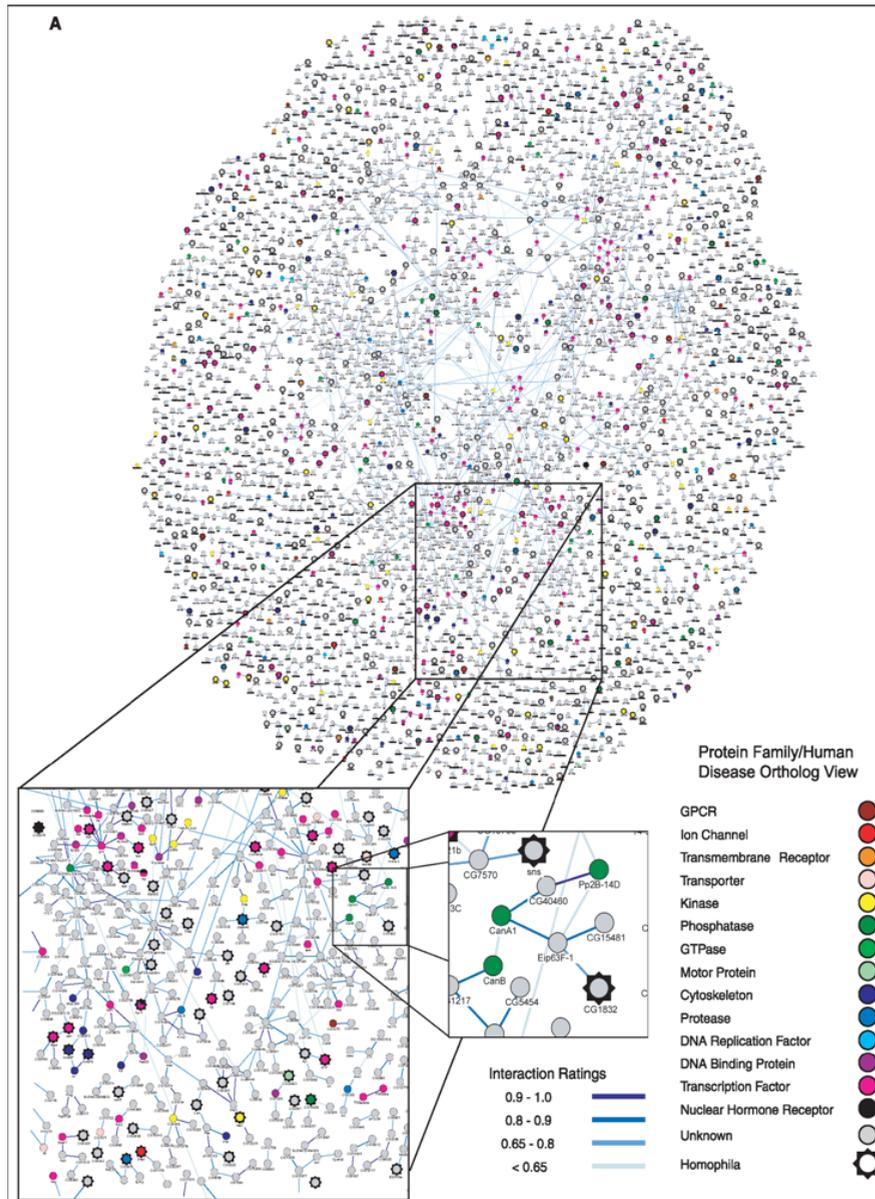
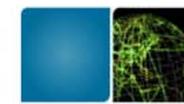
Protein-interaction networks



# Biological Complex Networks

- Protein-interaction networks (PINs), signal transduction networks, biological pathways, food-webs
- PIN analysis: Novel protein function prediction, identification of critical nodes
- High-throughput experimental techniques → massive volumes of biological data
- Protein-interaction datasets are available for yeast (high-confidence), human, fly

# Giot L, Bader JS, ..., Rothberg JM, A protein interaction map of *Drosophila melanogaster* Science 302: 1727-1736, 2003.

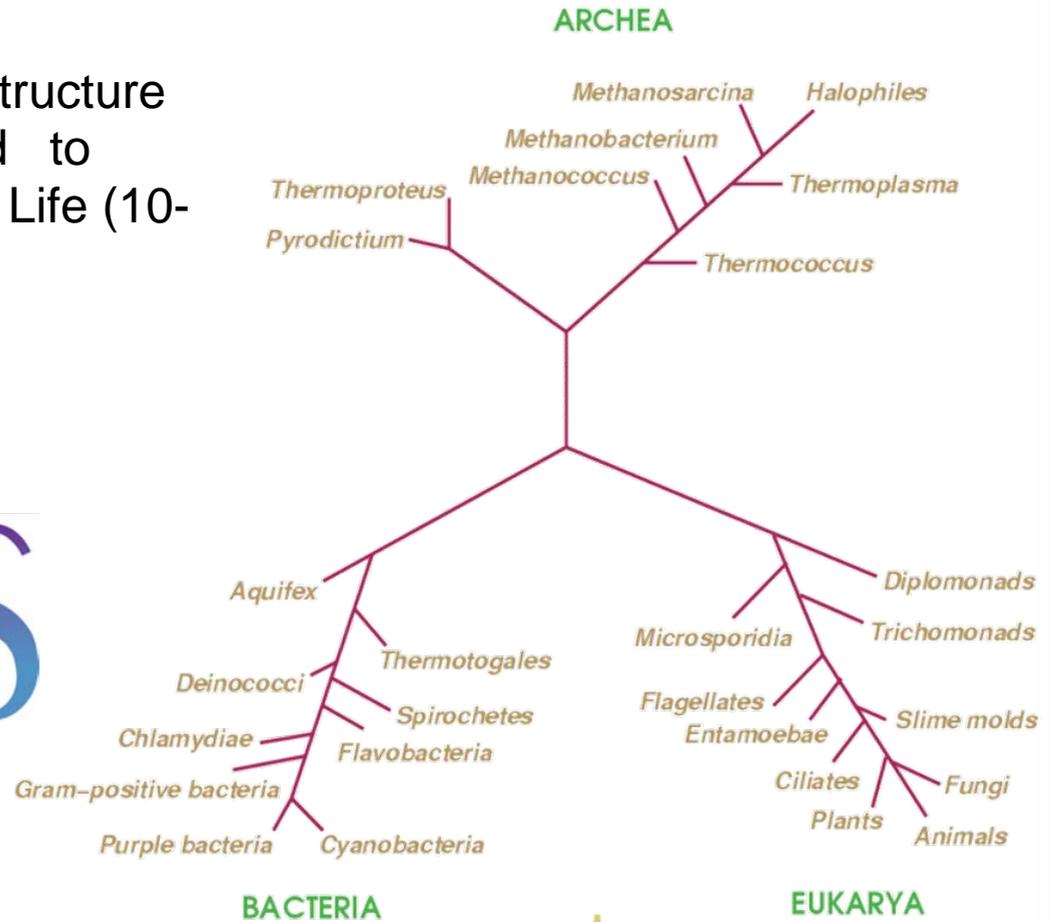


# NSF Cyberinfrastructure for Phylogeny Research (CIPRES)



CIPRES aims to establish the cyber infrastructure (platform, software, database) required to attempt a reconstruction of the Tree of Life (10-100M organisms)

NSF ITR Large award, \$12M



**FACULTY**  
**David A. Bader**, Director of HPC



[www.phylo.org](http://www.phylo.org)

David A. Bader



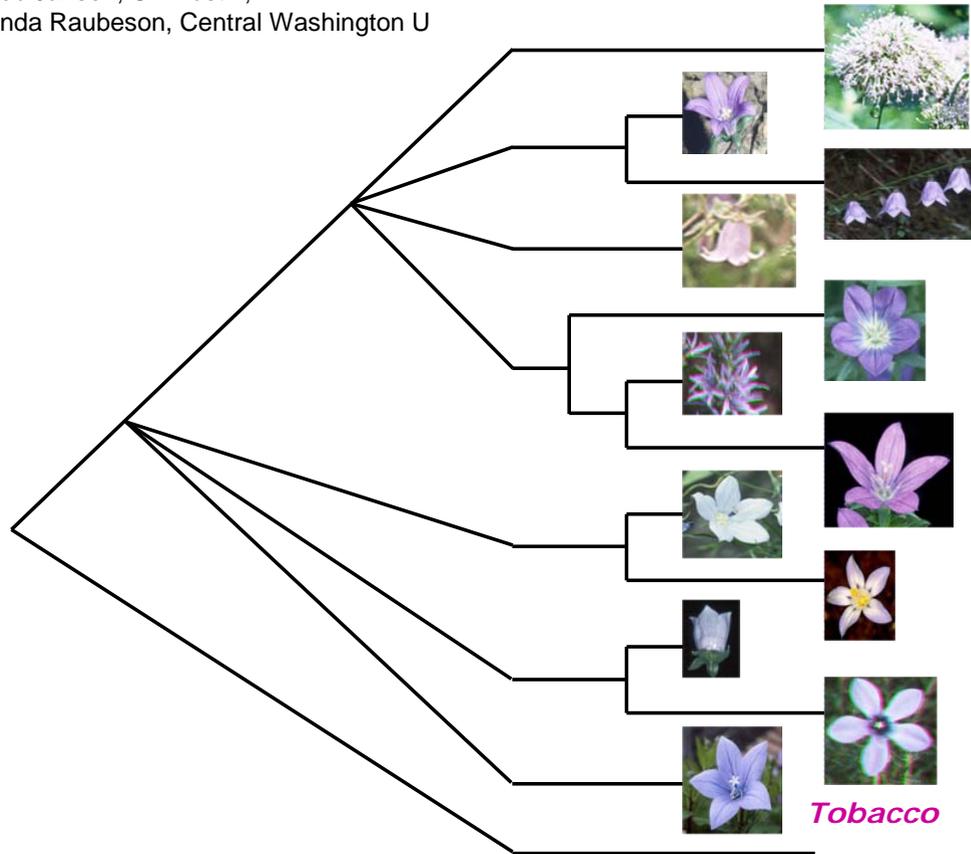
# Computational Phylogeny



## GRAPPA

### *Campanulaceae*

- Bob Jansen, UT-Austin;
- Linda Raubeson, Central Washington U



- Genome Rearrangements Analysis under Parsimony and other Phylogenetic Algorithm
  - Freely-available, open-source, GNU GPL
  - already used by other computational phylogeny groups, Caprara, Pevzner, LANL, FBI, Smithsonian Institute, Aventis, GlaxoSmithKline, PharmCos.
- Gene-order Phylogeny Reconstruction
  - Breakpoint Median
  - Inversion Median
- over one-billion fold speedup from previous codes
- Parallelism scales linearly with the number of processors



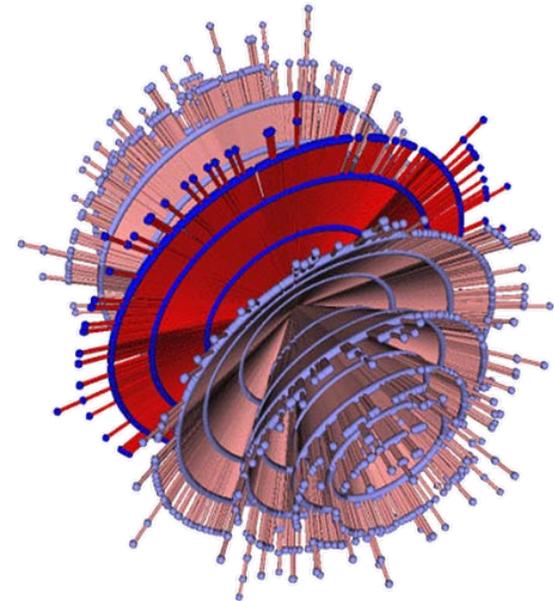
Cyber Infrastructure for Phylogenetic Research





# Informatics Graphs are Tough

- **Very different from graphs in scientific computing!**
  - Graphs can be enormous
  - Power-law distribution of the number of neighbors
  - Small world property – no long paths
  - **Very limited locality, not partitionable**
  - Highly unstructured
  - Edges and vertices have types



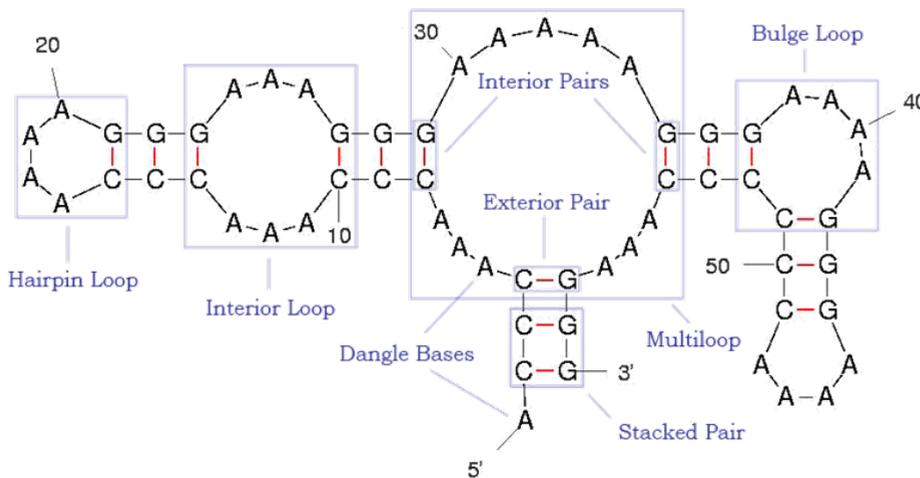
Six degrees of Kevin Bacon  
Source: Seokhee Hong

- Experience in scientific computing applications provides only limited insight.

# RNA Secondary Structure Prediction



- RNA is composed of smaller building blocks called bases (Adenine, Cytosine, Guanine, Uracil)
- Pairing and non-pairing of bases is called “folding”
- Result of folding called secondary structure



## Program Goals

Accurate structure of large viruses such as:

- Influenza
- HIV
- Polio
- Tobacco Mosaic
- Hanta



## FACULTY

**Christine Heitsch** (Mathematics)

**David A. Bader**

**Steve Harvey** (Biology)



# Cyberinfrastructure Challenges

- Current HPC systems are designed for physics-based simulations that use
  - Floating-point, linear algebra
    - Top 500 List measures Linpack!
  - Regular operations (high-degrees of locality)
    - e.g., Matrices, FFT, CG
  - Low-order polynomial-time algorithms
- Focus of current HPC/petascale systems:
  - Dense linear algebra
  - Sparse linear algebra
  - FFT or multi-grid
  - Global scatter-gather operations
  - Dynamically evolving coordinate grids
  - Dynamic load-balancing
  - Particle-based or lattice-gas algorithms
  - Continuum equation solvers
- Computational biology and bioinformatics often require
  - Integer performance
    - Strings, trees, graphs
  - Combinatorics
    - Optimization, LP
    - Computational geometry
  - Irregular data accesses
  - Dynamic programming, backtracking
  - Heuristics and solutions to NP-hard problems
- **Exascale cyberinfrastructure must address these problems!**



# SUMMARY RECOMMENDATIONS

- Support the design and analysis of exascale algorithms for large-scale problems in biological sciences
- Design exascale cyberinfrastructure to meet the requirements of whole genome and systems biology applications
- Example exascale challenge problems:
  - Compute the Tree of Life: phylogenetic relationships, whole genome comparisons
  - Simulate life through systems from atoms to organisms to populations
  - Understand sustainable ecosystems through modeled interactions between organisms, environment, climate, etc.



# Collaborators

- Virat Agarwal
- Kamesh Madduri
- Seunghwa Kang
- Sulabh Patel (EA)
- Aparna Chandramowlishwaran
- Vipin Sachdeva (IBM Austin Research Lab)
- Guojing Cong (IBM TJ Watson Research Center)
- Bruce Hendrickson (Sandia National Laboratories)
- Jon Berry (Sandia National Laboratories)
- John Feo (Microsoft)
- Michael Perrone (IBM)
- Sid Manning (IBM)
- Bob Szabo (IBM)



# Acknowledgment of Support

- National Science Foundation

- PETA APPS
- CSR: A Framework for Optimizing Scientific Applications (06-14915)
- CAREER: High-Performance Algorithms for Scientific Applications (06-11589; 00-93039)
- ITR: Building the Tree of Life -- A National Resource for Phyloinformatics and Computational Phylogenetics (EF/BIO 03-31654)
- DEB Comparative Chloroplast Genomics: Integrating Computational Methods, Molecular Evolution, and Phylogeny (01-20709)
- ITR/AP(DEB): Computing Optimal Phylogenetic Trees under Genome Rearrangement Metrics (01-13095)
- DBI: Acquisition of a High Performance Shared-Memory Computer for Computational Science and Engineering (04-20513).



- National Institutes of Health

- IBM PERCS / DARPA High Productivity Computing Systems (HPCS)

- DARPA Contract NBCH30390004

- IBM Shared University Research (SUR) Grant

- Sony-Toshiba-IBM (STI)

- Microsoft Research

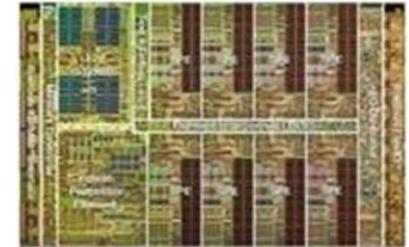
- Sun Academic Excellence Grant



# Sony-Toshiba-IBM Center of Competence for the Cell/B.E. at Georgia Tech



- ❖ **Mission:** grow the community of Cell Broadband Engine users and developers
- **Fall 2006:** Georgia Tech wins competition for hosting the STI Center
- First publicly-available IBM QS20 Cluster
- 200 attendees at 2007 STI Workshop
- Multicore curriculum and training
- **Demonstrated performance** on
  - Multimedia and gaming
  - Scientific computing
  - Medical applications
  - Financial services



David A. Bader, Director



**TOSHIBA**



<http://sti.cc.gatech.edu>

David A. Bader

# NSF Computing Research Infrastructure: Development of a Research Infrastructure for Multithreaded Computing Community Using Cray Eldorado Platform



- The Cray XMT system serves as an ideal platform for the research and development of algorithms, data sets, libraries, languages, tools, and simulators for applications that benefit from large numbers of threads, massively data intensive, *sparse-graph* problems that are difficult to parallelize using conventional message-passing on clusters.
  - A shared community resource capable of efficiently running, in experimental and production modes, complex programs with thousands of threads in shared memory;
  - Assembling software infrastructure for developing and measuring performance of programs running on the hardware; and
  - Building stronger ties between the people themselves, creating ways for researchers at the partner institutions to collaborate and communicate their findings to the broader community.



**CRAY**

## FACULTY

**David A. Bader**, PI (CSE)

**Jeffrey Vetter**, co-PI (CSE)

Collaborators include: Univ of Notre Dame, Univ. of Delaware, UC Santa Barbara, CalTech, UC Berkeley, Sandia National Laboratories