Re-visiting the Third Pillar of Science

Wu Feng
wfeng@vt.edu

QCD at the Femtoscale in the Era of Big Data
Institute of Nuclear Theory

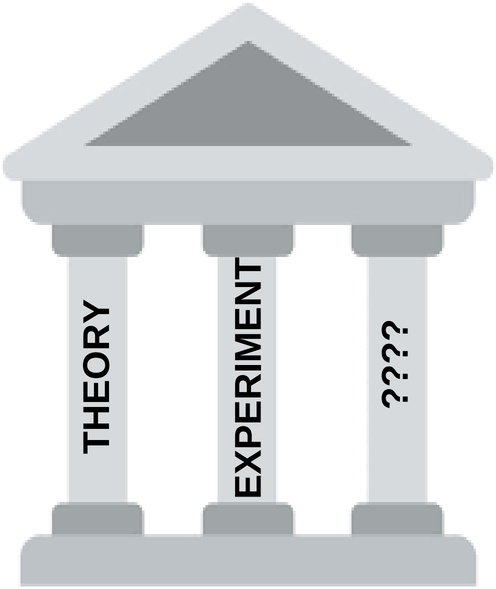

SyNeRGy
http://synergy.cs.vt.edu

VIRGINIA TECH

A Little Bit About Me …

- Education
  - Ph.D., Computer Science,
    U. Illinois at Urbana-Champaign, 1996
- Professional
  - Current Appointments
    - Professor and Elizabeth & James Turner Fellow; Departments of Computer Science, Electrical & Computer Engg., and Health Sciences; Virginia Tech
    - Director, SyNeRG Laboratory (http://synergy.cs.vt.edu/) → SEEC Center
    - Site Director, Center for Space, High-performance, and Resilient Computing (SHREC)
  - Previous Appointments & Professional Stints
    - *Academia:* Ohio State U. ('00-'03), Purdue U. ('98-'00), U. of Illinois at Urbana-Champaign ('96-'98)
    - *Government:* Los Alamos Nat'l Lab ('98-'06), NASA Ames Research Ctr ('93)
    - *Industry:* IBM T.J. Watson Rsch ('90), Vosaic ('97), Orion Multisystems ('04-'05), EnergyWare ('08-'10)

The Three Pillars of Science

"Together with ***theory*** and ***experimentation***, ***computational science*** now constitutes the third pillar of scientific inquiry, enabling researchers to build and test models of complex phenomena … that cannot be replicated in the laboratory."

– Presidential Information Technology Advisory Committee (PITAC), 2005

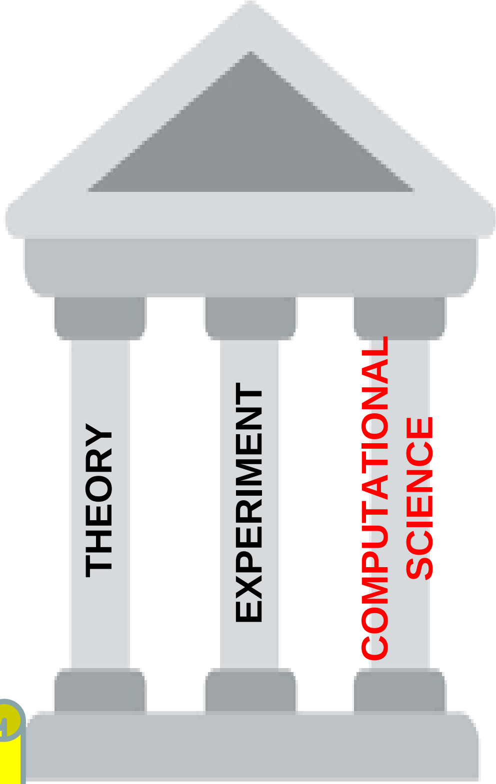

The third pillar of science is viewed as a way to ***simulate physical reality*** to design, predict, and optimize natural and engineered systems.

Computational Science → Berkeley Dwarfs → OpenDwarfs

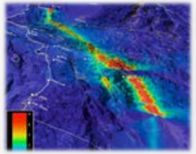
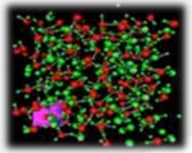
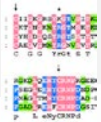
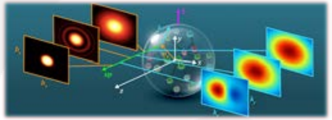


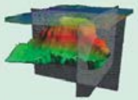
Computational Geoscience

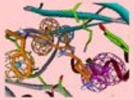
Computational Chemistry

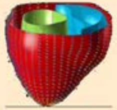
Computational Medicine

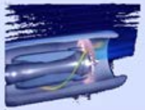
Computational Modeling


Computational Physics

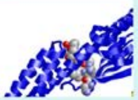
Computational Biology

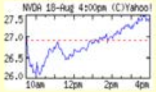

Computational Finance


Image Processing

⋮

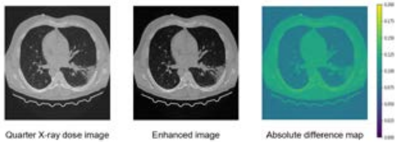

⋮




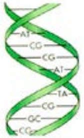
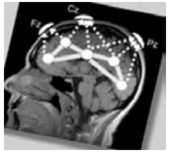
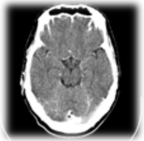


Race to Sequence the Human Genome

- Theory & Experiment (Collins@NIH)
  - Goal: Complete in 15 years
    1990 – 2005
  - Cost: $3,000M (1990-2000/2003)

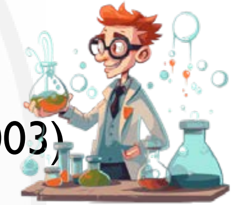

- Computing (Venter@Celera)
  - Goal: Complete in 3 years & cheaper
    1998 - 2001
  - Cost: $300M (1998-2000/2003)

mpiBLAST

ancestry

The third pillar of science is viewed as a way to ***simulate physical reality*** to design, predict, and optimize natural and engineered systems.

Bat Wing Simulation on CPU and GPU

(Source: A. Amritkar, W. Feng, D. Tafti, Virginia Tech)

Funding Acknowledgements:
AFOSR FA9550-12-1-0442 and
NSF CNS-0960081

CPU

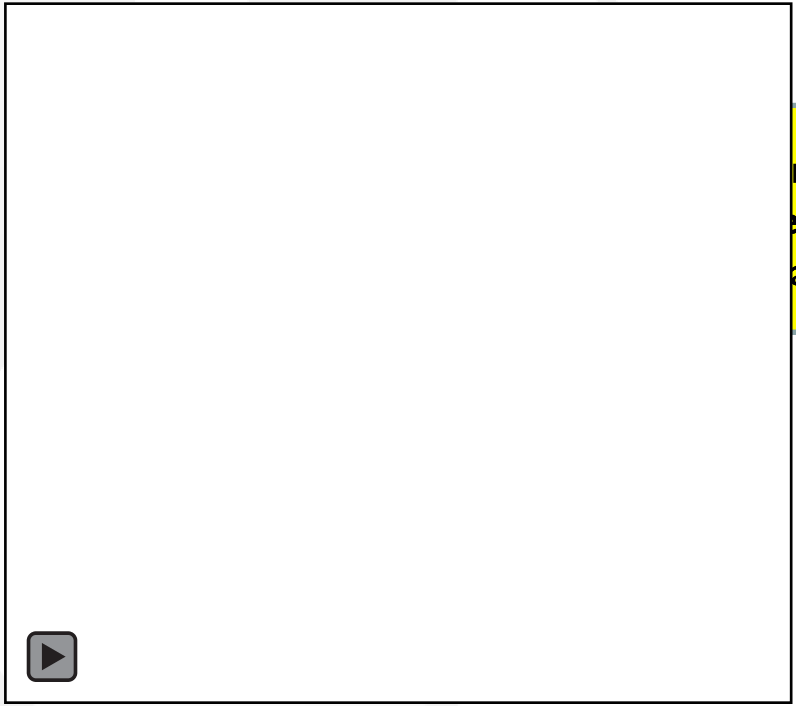

GPU

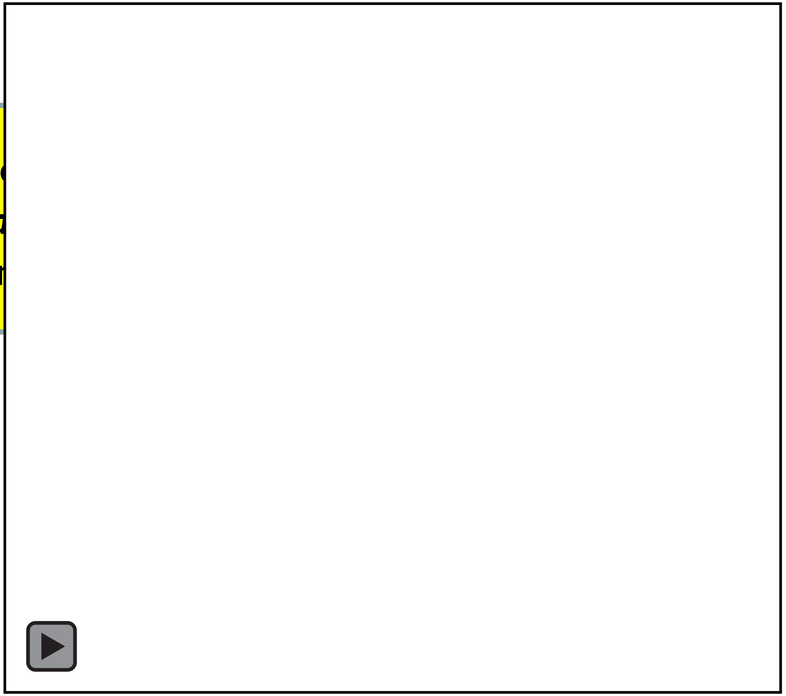

W. Feng, wfeng@vt.edu, 540.231.1192

Evolution of the Third Pillar of Science

- Computational Science: Genomics, AccuWeather, and Real Madrid
  - **Simulate physical reality** to design, predict, and optimize systems

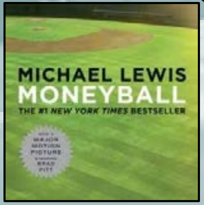

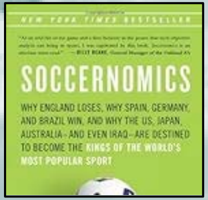

Subtle broadening …
From ***simulating physical reality***
to ***"computing on the data"*** too!

Takeaway #1: The third pillar of science is simply **COMPUTING**, encompassing ***simulating physical reality*** and ***computing on the data.***

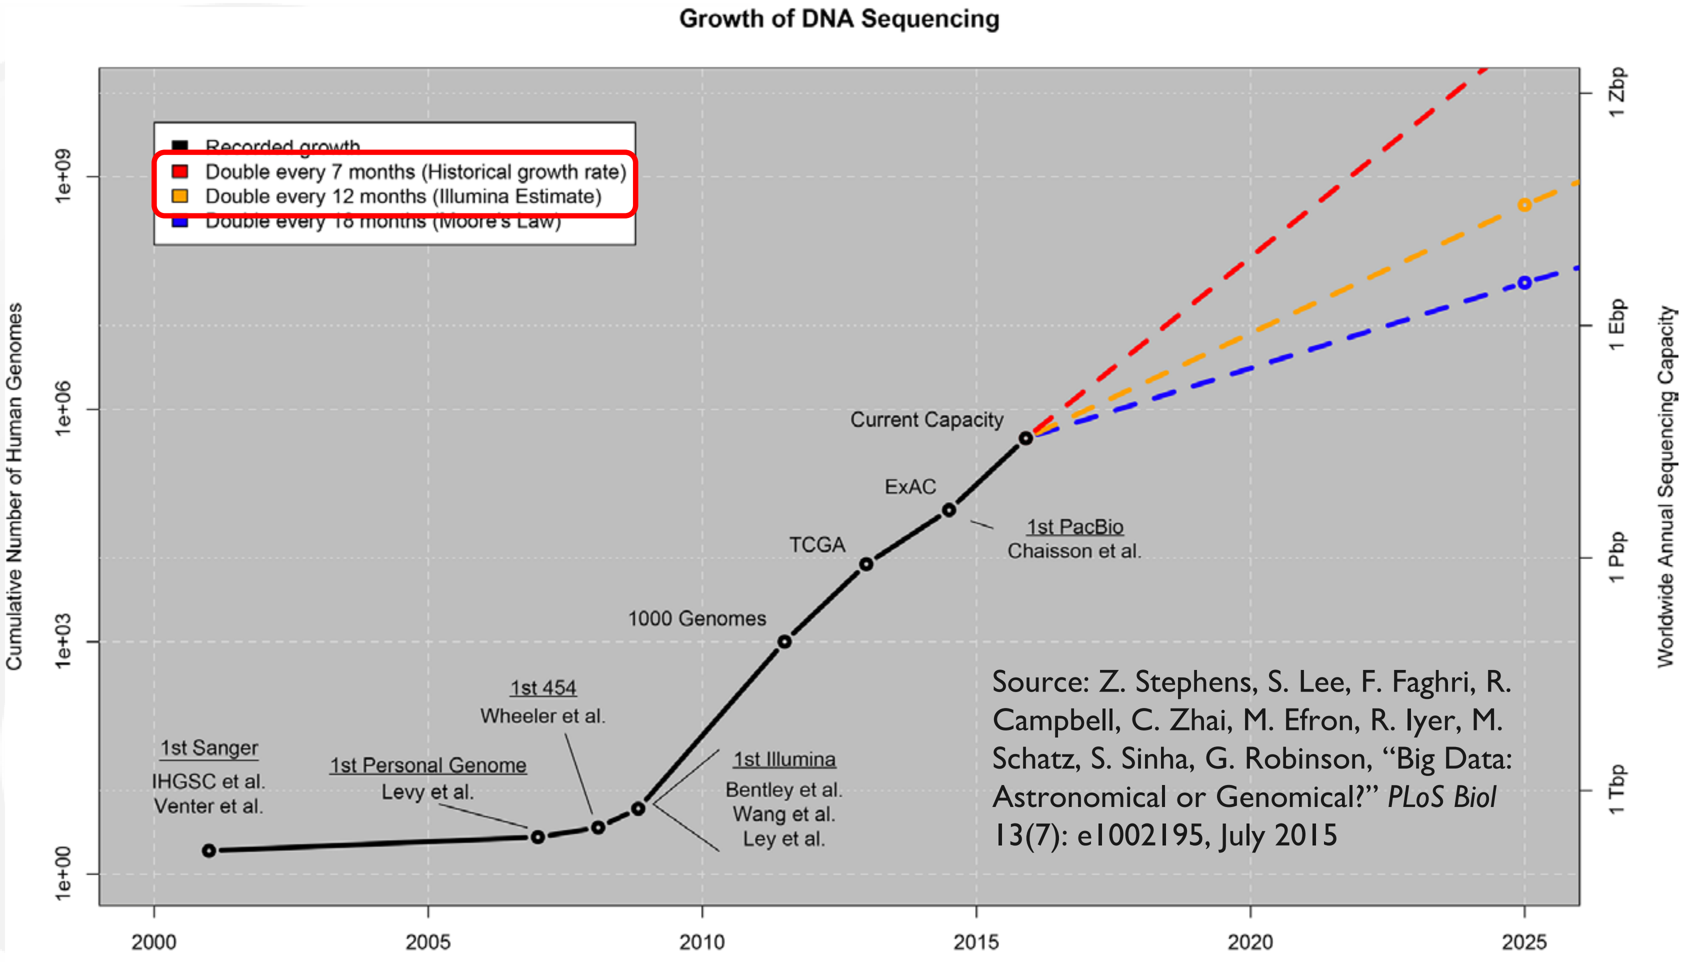

Source: Z. Stephens, S. Lee, F. Faghri, R. Campbell, C. Zhai, M. Efron, R. Iyer, M. Schatz, S. Sinha, G. Robinson, "Big Data: Astronomical or Genomical?" *PLoS Biol* 13(7): e1002195, July 2015

LHC Data Analysis

Taking a closer look at LHC

The LHC produces at design parameters over 600 millions collisions (~ 10^9 collisions) proton-protonper second in ATLAS or CMS detectors. The amount of data collected for each event is around 1 MB (1 Megabyte).

10^9 collisions/s x 1 Mbyte/collision = 10^{15} bytes/s = **1 PB/s** (1 Petabyte/second)

Since 1 DVD ~ 5 GB : 200000 DVDs per second would be filled or about 6000 IPods (ones with 160 GB of storage) per second!

This is several orders of magnitude greater than what any detector data acquisition system can handle.

A **trigger** is designed to reject the uninteresting events and keep the interesting ones (more information about trigger in this page).

For example, the ATLAS trigger system is designed to collect about 200 events per second.

200 events/s x 1 Mbyte = 200 MB/s (200 Megabyte/second)

Taking two shifts of ten hours per day, and about 300 days per year:

200 MB/s x 2 x 10 x 3600 x 300 ~ $4 \cdot 10^{15}$ bytes/year = **4 PB/year**

Doubling every two years?

Collectively, the **LHC experiments produce about 15 petabytes of raw data each year** that must be stored, processed, and analyzed.

A **three level trigger** is used to select events that show signs of interesting physics processes.

Serial Performance (SPECintCPU)

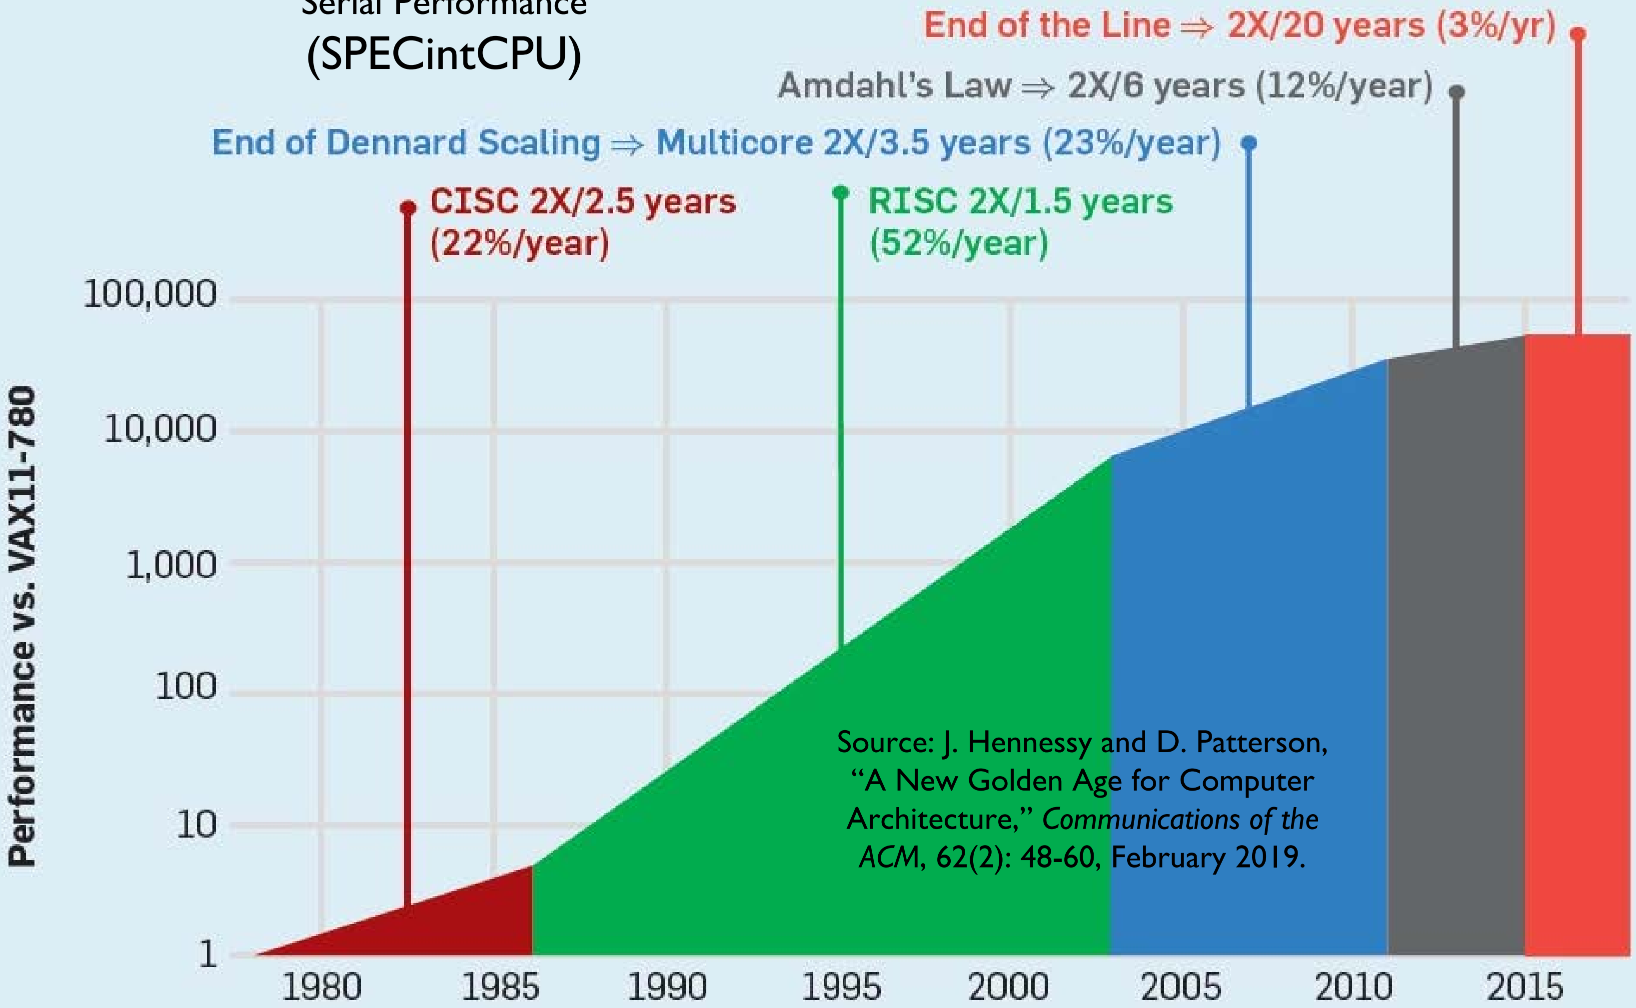

Source: J. Hennessy and D. Patterson, "A New Golden Age for Computer Architecture," *Communications of the ACM*, 62(2): 48-60, February 2019.

Challenge

- The rate of growth in **big data** is ***far outstripping*** the rate at which computing can (brute-force) **compute** on the data.

Approach

- Synergistic co-design of architecture, software, and in particular, algorithms to more ***efficiently*** and ***intelligently*** compute on the data.

Re-Visiting the Third Pillar of Science:

Beyond Simulating Physical Reality → COMPUTING

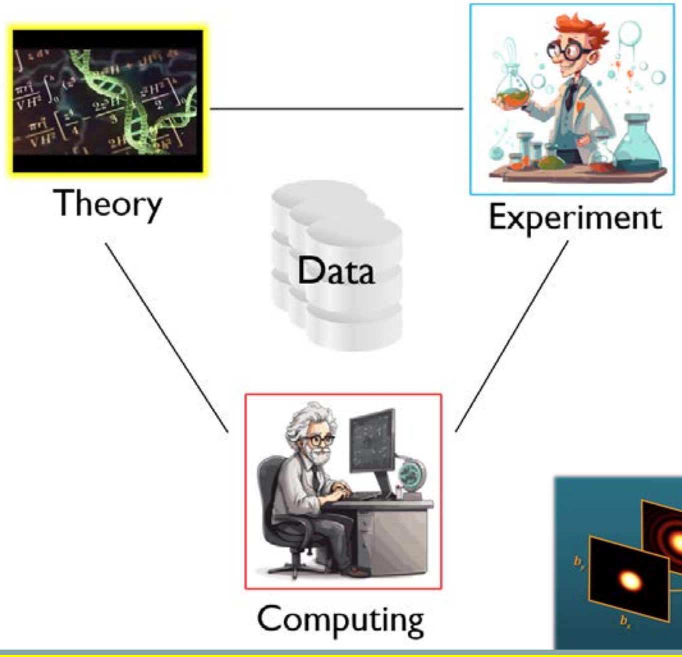

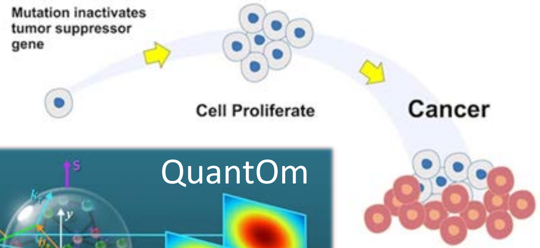

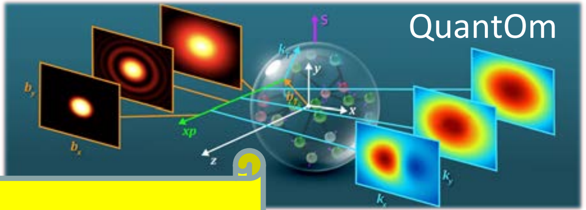

Takeaway #1: The third pillar of science is simply **COMPUTING**, encompassing ***simulating physical reality*** and ***computing on the data.***

Re-visiting the Third Pillar of Science

Takeaway #2: Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery, e.g., rational drug design

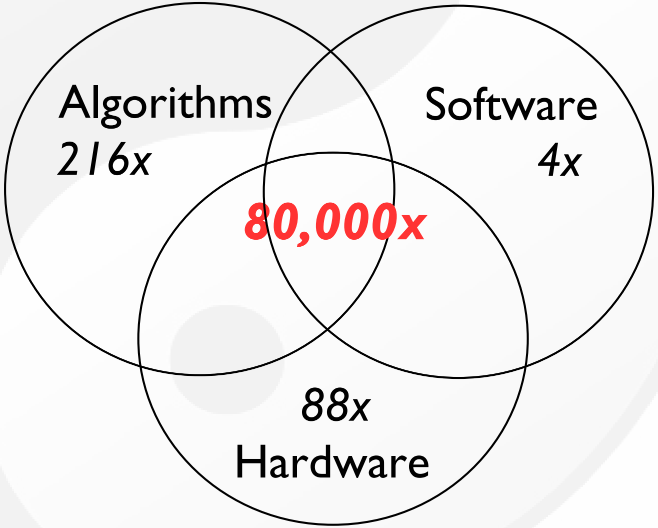

- Daga, Scogland, & Feng, "Architecture-Aware Mapping and Optimization on a 1600-Core GPU," In *Proc. IEEE Int'l Conf. on Parallel & Distributed Systems*, 12/2011.
  - > 88x speedup relative to serial CPU implementation
  - > 4x speedup via software optimization

http://www.youtube.com/watch?v=zPBFenYg2Zk

Re-visiting the Third Pillar of Science

Takeaway #2: Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery, e.g., rational drug design

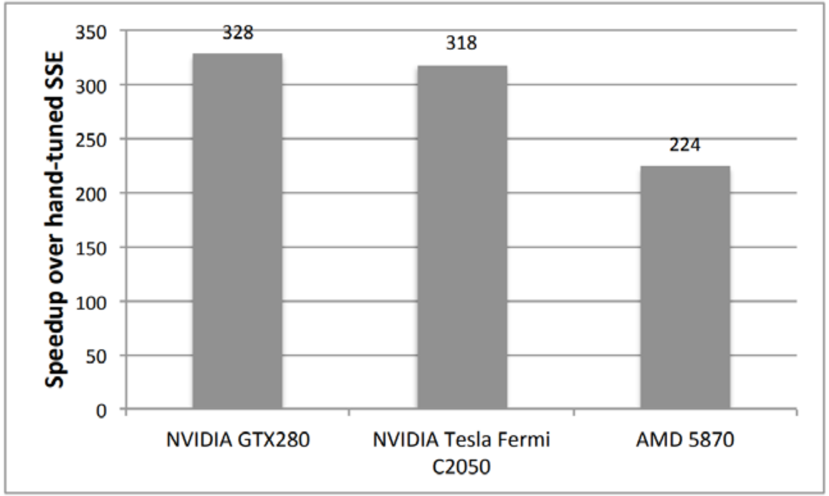

(a) Speedup on GPU Platforms with NVIDIA-Specific Optimizations

Daga, Scogland, & Feng, "Architecture-Aware Mapping and Optimization on a 1600-Core GPU," In *Proc. IEEE Int'l Conf. on Parallel & Distributed Systems*, 12/2011.

> 88x speedup relative to serial CPU implementation

> 4x speedup via software optimization

Observation

– Optimized GPU molecular modeling code delivered higher speedup (328x) on an *OLDER* GPU! Why?

Re-visiting the Third Pillar of Science

Takeaway #2: Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery, e.g., rational drug design

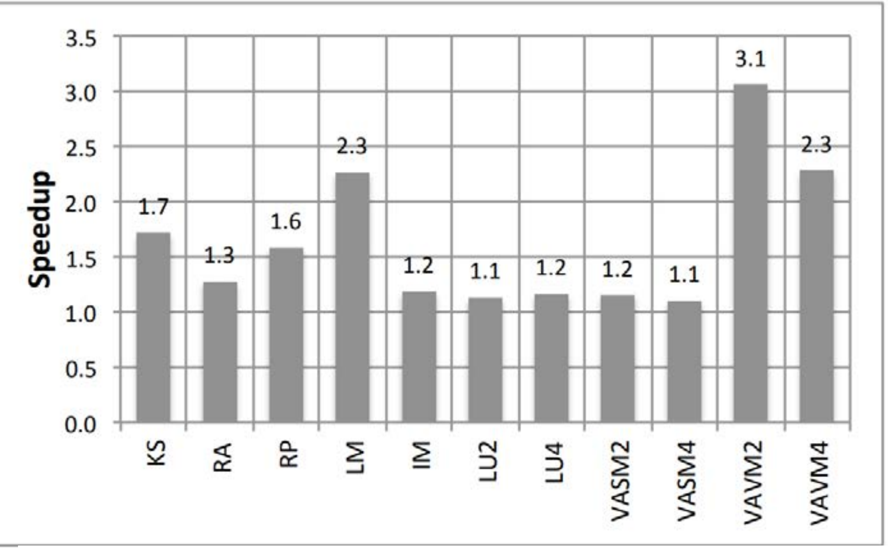

KS: Kernel Splitting, RA: Register Accumulator, RP: Register Preloading, LM: Local Memory, IM: Image Memory, LU{2,4}: Loop Unrolling{2x,4x}, VASM{2,4}: Vectorized Access & Scalar Math{float2, float4}, VAVM{2,4}: Vectorized Access & Vector Math{float2, float4}

Daga, Scogland, & Feng, "Architecture-Aware Mapping and Optimization on a 1600-Core GPU," In *Proc. IEEE Int'l Conf. on Parallel & Distributed Systems*, 12/2011.

> 88x speedup relative to serial CPU implementation

> 4x speedup via software optimization

Observation

- Multiplicative speedup with the right (or compatible) combination of optimizations → 4.22x

Who Cares? Dated Results?!

- Yes, but that's *not* the point. Applicable to next-gen HPC.

Re-visiting the Third Pillar of Science

Takeaway #2: Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery, e.g., rational drug design

> 88x speedup relative to serial CPU implementation

> 4x speedup via software optimization

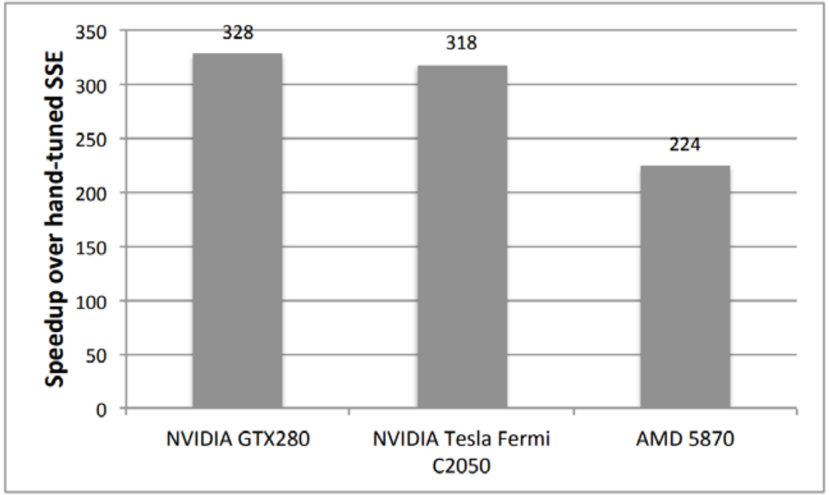

(a) Speedup on GPU Platforms with NVIDIA-Specific Optimizations

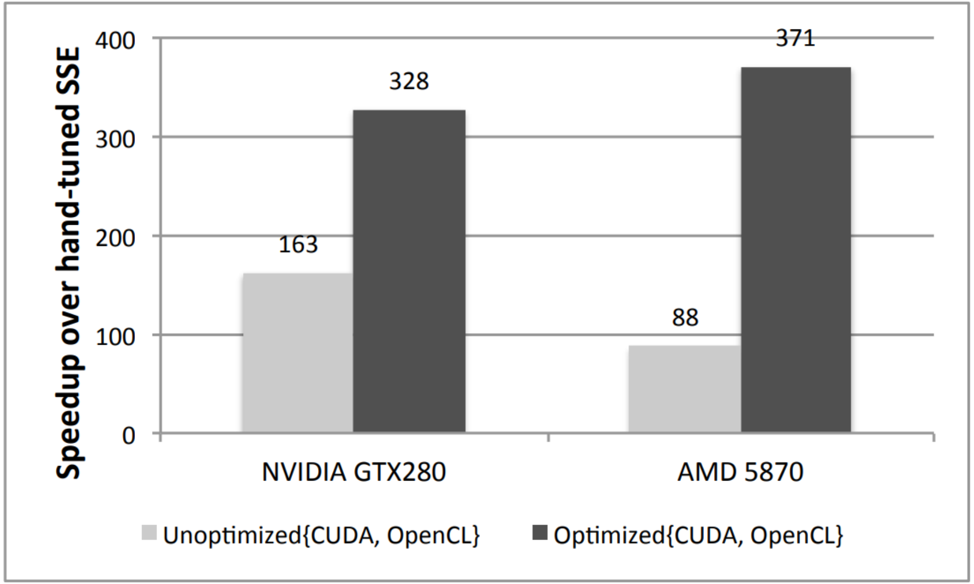

Re-visiting the Third Pillar of Science

Takeaway #2: Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery.

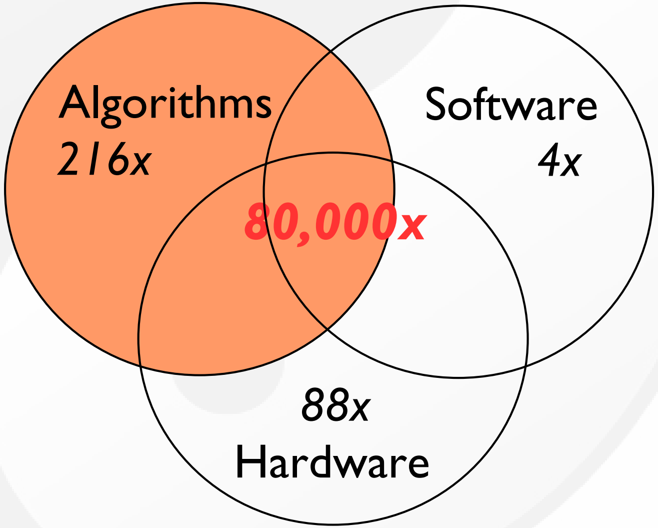

http://www.youtube.com/watch?v=zPBFenYg2Zk

- Daga, Scogland, & Feng, "Architecture-Aware Mapping and Optimization on a 1600-Core GPU," In *Proc. IEEE Int'l Conf. on Parallel & Distributed Systems*, 12/2011.
  - > 88x speedup relative to serial CPU implementation
  - > 4x speedup via software optimization
- Daga & Feng, "Multi-Dimensional Characterization of Electrostatic Surface Potential Computation on Graphics Processors," *BMC Bioinformatics*, 4/2012.
- Anandakrishnan et al., "Multiscale Approximation with Graphical Processing Units for Multiplicative Speedup in Molecular Dynamics," In *Proc. ACM Int'l Conf. on Bioinformatics, Comp. Biology, & Health Informatics*, 10/2016.

Re-visiting the Third Pillar of Science

- Computing is a transformative "third pillar of science" to accelerate discovery, which is the purpose of the DOE SciDAC program, i.e., Scientific Discovery through Advanced Computing

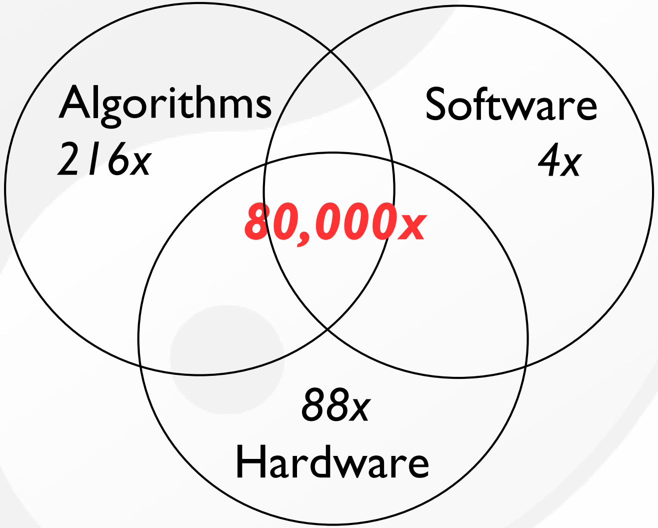

Takeaway #3: Be careful!
Don't fool yourself and, in turn, fool the masses.

http://www.youtube.com/watch?v=zPBFenYg2Zk

12 Ways to Fool the Masses
(David H. Bailey, NASA & LBL, 1991)

1. Quote only 32-bit performance results, ***not*** 64-bit results. → 2x speedup
2. Present performance figures for an inner kernel and then represent these figures as the performance of the entire application.
3. Quietly employ assembly code and other low-level language constructs.
4. Scale up the problem size with the number of processors but omit any mention of this fact → WEAK SCALING
5. Quote performance results *projected* to a full system.
6. Compare your results against scalar, unoptimized code on conventional systems.
7. When direct run-time comparisons are required, compare with an old code on an obsolete system.
8. If Mflop/s rates must be quoted, base the operation count on the parallel [version], ***not*** on the best sequential [version].
9. Quote performance in terms of processor utilization, parallel speedups, or Mflop/s per dollar.
10. Mutilate the algorithm used in the parallel implementation to match the architecture.
11. Measure parallel run-times on a dedicated system but measure conventional run times in a busy environment.
12. If all else fails, show pretty pictures and videos, and don't talk about performance.

Summary: Re-visiting the Third Pillar of Science

1. The third pillar of science is simply COMPUTING, encompassing simulating physical reality and computing on the data.
2. Synergistic co-design of algorithms, software, and hardware can massively accelerate discovery.
3. Don’t fool yourself and, in turn, fool the masses.

Aspirational Goal: A visionary report that addresses the trends and grand challenges for the femtoscale inverse problem to access information on quarks and gluons

Meeting report: the terabase metagenomics workshop and the vision of an Earth microbiome project — 315 — 2010
JA Gilbert, F Meyer, D Antonopoulos, P Balaji, CT Brown, CT Brown, ...
Standards in genomic sciences 3, 243-248

What's Next?

- Case Studies on Synergistic Co-Design of Algorithms, Software, and Hardware
  - Brain Tomography on GPU. Carcinogenesis: Weighted Set Cover vs. Graph Cluster. […]
- HPC Systems
  - IterML: Iterative Machine Learning (AFOSR & DOD)
    - Context: Computational fluid dynamics (CFD) → OpenDwarfs, i.e., fundamental "DNA" building blocks for scientific computing
  - CoreTSAR: Core Task-Size Adapting Runtime System (DOE & NSF)
    - Context: Initially, discrete CPU+GPU systems w/ discrete memory
      Now, also "fused" co-located CPU+GPU systems w/ shared memory
      See Aurora @ ANL with PVC & El Capitan @ LLNL with MI-300a
  - Scalable Deep Learning (with ANL → Meta & Llama-3)
    - Context: Caffe, Caffe2, and Tensorflow
    - Takeaway: Large-scale multi-node DL does *NOT* scale.