



When High-Performance Computing Meets Bioinformatics

Wu Feng, Ph.D.

Department of Computer Science
Department of Electrical & Computer Engineering



www.mpiblast.org



A Little Bit About Me ...

- Professional
 - Current Appointments
 - Associate Professor, CS & ECE, Virginia Tech
 - Director, **SyNeRG** Laboratory (<http://synergy.cs.vt.edu/>)
 - Member, Center for High-End Computing Systems (CHECS)
 - Faculty Co-Director, Center for High-Performance Reconfigurable Computing (CHREC)
 - Previous Appointments & Professional Stints
 - Academia: The Ohio State U., Purdue U., Univ. of Illinois
 - Government: Los Alamos National Laboratory, NASA Ames Research Center
 - Industry: Orion Multisystems, Vosaic, IBM T.J. Watson Research Center



A Little Bit More About Me ...

- Research Areas
 - High-Performance Computing (<http://synergy.cs.vt.edu/>)
 - Systems Software & Adaptive Run-Time Systems
 - Network Protocols & Interconnects
 - Monitoring & Measurement
 - Virtual Machines for K-8 Pedagogy (<http://myvice.cs.vt.edu/>)
 - Accelerators (<http://accel.cs.vt.edu/>)
 - Green Supercomputing
 - Low-Power & Power-Aware Supercomputing (<http://sss.cs.vt.edu/>)
 - The Green500 List (<http://www.green500.org/>)
 - Bioinformatics
 - Sequence Search, e.g., mpiBLAST (<http://www.mpiblast.org/>)
 - Short-Read Mapping Software, e.g., RMAP, MUMmer
 - Molecular Dynamics, e.g., GEM and NAB
 - Neuroinformatics
 - ParaMEDIC: “Data Teleportation” of Bio Data



Mission

- Deliver solutions at the research intersection of
 - Commoditization / Personalization / Ease of Use
 - High-Performance Computing
 - Bioinformatics
 - ... or any application area that deals with large-scale datasets and data processing



Goal

- Empower scientists to realize the untapped potential of commodity multi-core and many-core processors.
 - **Traditional and heterogeneous multi-core**
 - AMD (8) and Intel (6 → 8) and Cell in Sony PlayStation 3 (1+8)
 - **Reconfigurable fixed multi-core**
 - Tiler TILE64 (64) and Tiler TILE-GX-100 (100)
 - **Field programmable gate array (FPGA)**
 - Xilinx Virtex5 → Convey HC-1
 - **GPU many-core**
 - ATI Radeon HD5870, NVIDIA Tesla → NVIDIA Fermi
 - *Other platforms*
 - IBM BlueGene/L, IBM BlueGene/P
- Desktop Supercomputing “In-a-Box”



How Did I Come into Bioinformatics?

- From an infrastructure perspective ...



Green Destiny Supercomputer


(circa December 2001 – February 2002)



- A 240-Node Cluster in Five Sq. Ft.
- Each Node
 - 1-GHz Transmeta TM5800 CPU w/ High-Performance Code-Morphing Software running Linux 2.4.x
 - 640-MB RAM, 20-GB hard disk, 100-Mb/s Ethernet
- Total
 - 240 Gflops peak (LINPACK: 101 Gflops in March 2002.)
 - 150 GB of RAM (expandable to 276 GB)
 - 4.8 TB of storage (expandable to 38.4 TB)
 - **Power Consumption: Only 3.2 kW (diskless)**
- Reliability & Availability
 - **No unscheduled downtime in 24-month lifetime.**
 - Environment: A dusty 85-90°F warehouse!

Equivalent LINPACK to a
256-CPU SGI Origin 2000
(On the TOP500 List at the time)

Featured in *The New York Times*, *BBC News*, and *CNN*.
Now in the Computer History Museum.



Michael S. Warren, Los Alamos National Laboratory

Green Destiny → mpiBLAST?

- **Green Destiny** (<http://sss.cs.vt.edu/>)
 - ... resulted in tremendous media **over-exposure**
- ✓ “Efficient Supercomputing with Green Destiny,” *slashdot.org*, Nov. 2003.
 - ✓ “Green Destiny: A ‘Cool’ 240-Node Supercomputer in a Telephone Booth,” *BBC News*, Aug. 2003.
 - ✓ “Servers on the Edge: Blades Promise Efficiency and Cost Savings,” *CIO Magazine*, Mar. 2003.
 - ✓ “Developments to Watch: Innovations,” *BusinessWeek*, Dec. 2002.
 - ✓ “Craig Venter Goes Shopping ...,” *GenomeWeb*, Oct. 2002.
 - ✓ “Not Your Average Supercomputer,” *Communications of ACM*, Aug. 2002.
 - ✓ “At Los Alamos, Two Visions of Supercomputing,” *The New York Times*, Jun. 25, 2002.
 - ✓ “Bell, Torvalds Usher Next Wave of Supercomputing,” *CNN*, May 2002.



Q&A: Pharmaceuticals + J. Craig Venter

- Q&A Exchange with Pharmaceutical Companies (2002)
 - Pharmaceutical: “Can you get the same type of results for bioinformatics applications?”
 - Wu: “What is your primary application?”
 - Pharmaceutical: “BLAST ...”
- J. Craig Venter in *GenomeWeb* (October 2002)

“... to build **something that is replicable so any major medical center around the world can have a chance to do the same level of computing** ... interested in IT that doesn’t require massive air conditioning. The room at Celera cost \$6M before you put the computer in. [Thus, I am] **looking at these new green machines being considered at the DOE that have lower energy requirements & therefore produce less heat.**”

What is BLAST?

Basic Local Alignment Search Tool

```
>gb|AE008729.1| Salmonella typhimurium LT2, section 31
of 220 of the complete genome
Length = 22294

Score = 46.1 bits (23), Expect = 0.054
Identities = 32/35 (91%)
Strand = Plus / Plus

Query: 157   ctgctgcatggcggcgcatcggtagcgctggcgga 191
            ||||| | ||||| ||||| ||||| |||||
Sbjct: 21920 ctgctacacggcggcgcatcggtagcgctggcgga 21954
```

- The Basic Idea
 - Given a query DNA or amino-acid sequence,
 - BLAST heuristically finds similar sequences in database.
 - BLAST reports the statistical significance of similarities between the query and database.
- The Challenge for BLAST
 - Growth of the GenBank database *outstripping* the growth of main memory. (Bottleneck: Disk I/O)
 - 67,000-fold growth in GenBank database (1982-2005).
- Solution
 - Parallelize BLAST or a variant thereof.




Approaches to Sequence Search

- *Smith-Waterman (SW) Algorithm*
 - Optimal local alignment of sequences
 - To align two sequences of lengths m and n ,
 - $O(mn)$ in time and space
- *FASTA Algorithm* by Lipman & Pearson
 - Heuristic alignment of sequences
 - Marks potential matches before doing an optimal SW type of algorithm
- *BLAST Algorithm* by Altschul et al.
 - Faster heuristic alignment that approximates SW




Becoming an Accidental Bioinformaticist

- Premise
 - Genomic databases growing faster than our ability to compute (e.g., sequence search) on them.
 - Discover *missing genes* in genomes via sequence search
- Approach
 - Use a parallel aggregate map-reduce distributed framework to leverage the database in a way that can scale up to petabyte capacity
- Results
 - At modest scale
 - At large scale
 - Missing Genes



mpiBLAST: Parallel Sequence Alignment

- NCBI BLAST: The 'Gold Standard' of Sequence Search
 - Can consume up to 90% of the compute cycles in bioinformatics
- mpiBLAST: Parallel BLAST on cluster supercomputers
 - Scalability
 - Ability to support unlimited database sizes
 - Tested to 32,768 processing cores
 - Performance
 - Dramatically improves both response time and throughput




Performance Comparison on 128-Node Cluster

	Throughput	Response Time
State of the Art	128x	1x
mpiBLAST	305x	305x



Scalability Results: mpiBLAST 1.6

- Scalability
 - Tens of thousands of nodes
- Test Setup
 - Reference Database: microbial genome
 - Test Query: 0.25 million queries from the same database
 - Search program: blastn
- Results
 - Initially, super-linear ...
 - 93% efficiency to 32,768 processor cores



“Massively Parallel Genomic Sequence Search on the Blue Gene/P Architecture”
 ACM/IEEE SC'08 (Supercomputing '08), Nov. 2008.

Accelerated Bioinformatics Software


- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - GPU-RMap
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner

GPU-BLAST

- Raw Prototype of GPU-BLAST
 - Uses GPU to select sequences with high similarities
 - Searches high similar sequences with CPU serial BLAST
 - **6x speedup over serial CPU BLAST**



GPU-BLAST: *Initial Results* → 6x Now ...



Query sequence: 1,115 characters
 Database: 159,639,626 characters




Accelerated Biocomputing Software

- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - GPU-RMap
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner




GPU-Smith-Waterman

- Existing Implementations
 - *Fast but brittle..*
- Our implementation
 - **Fast ... 14x speed-up**
 - CPU: 3,427 seconds = 0.952 hours
 - GPU: 238 seconds = 0.066 hours
 - **Robust**
 - Previous implementations may deliver better speed-up *but*
 - **Do not implement the backtrace, thus artificially inflating the speed-up, or**
 - **Can only take queries that fit in memory on the GPU, up to 8150 x 8150 on a 1GB GPU**
 - We address both the above limitations.



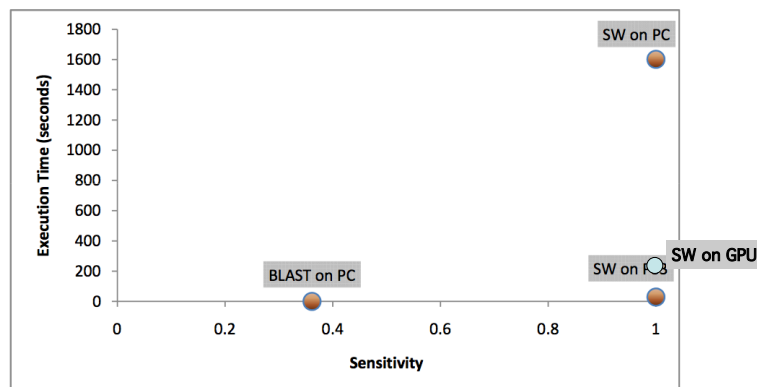
Perspective on Smith-Waterman

- Smith-Waterman (SW): Optimal but SLOW



Perspective on Smith-Waterman

- SW on Cell (i.e., PlayStation 3) and GPU



Accelerated Biocomputing Software

- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - GPU-RMap
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner

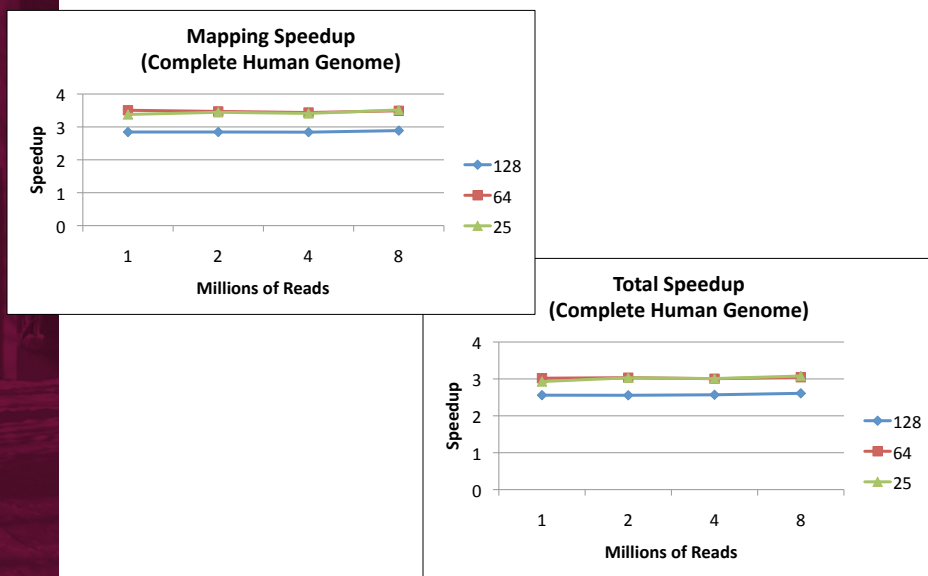


Accelerated Biocomputing Software

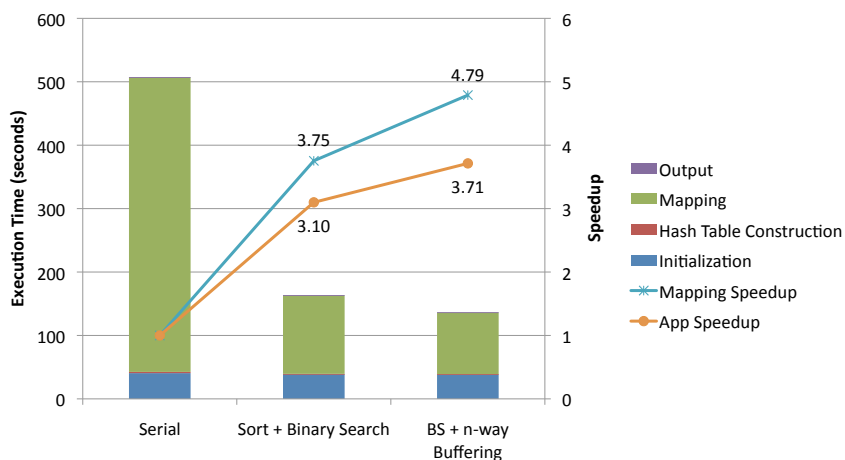
- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - **GPU-RMap**
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner



GPU-RMAP: Initial Results



CUDA-RMAP Profiling



Chr1 Vs 1 Million Reads (64 bp width)


Accelerated Biocomputing Software

- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - GPU-RMap
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner



GPU-Accelerated Molecular Dynamics

- *Fundamental Application in Computational Biology*
 - Simulate interactions between atoms & molecules for a period of time by approximations of known physics




- *Example Usage*
 - Understand mechanism behind the function of molecules
 - Catalytic activity, ligand binding, complex formation, charge transport
- GPU Acceleration
 - Performance: 172,340x speed-up on one GPU card
 - 22.4 hours → 0.47 seconds



Electrostatic Potential for Molecular Dynamics

Processor + Optimization	Execution Time	Speed-Up
CPU	80690.20	-
GPU	355.64	227x
GPU + Kernel Split	219.20	368x
GPU + Multi-Level 1	51.43	1569x
GPU + Kernel Split + Multi-Level 1	35.04	2303x




Viral Capsid

Now up to 172,340x speed-up in a single server.



Electrostatic Potential for Molecular Dynamics



Processor + Optimization	Power (W)	Execution Time (s)	Energy (J)	Energy-Delay Product (J · s)
CPU serial	98.53	2,037.53	200,757.83	409,050,101
CPU multi-core	98.53	1,135.61	111,891.65	127,065,277
GPU + kernel split	228.07	9.50	2,166.67	20,583

This has dropped by another two orders of magnitude with the latest results from previous slide.



Gratuitous YouTube Video

- (Near) Real-Time In-Situ Molecular Dynamics
 - See this at <http://www.youtube.com/watch?v=zPBFenYg2Zk> or Google 'YouTube Feng NVIDIA Virginia Tech'





Accelerated Biocomputing Software

- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
- Short-Read Mapping
 - GPU-RMap
- Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
- Neuroinformatics
 - GPU-Miner




GPU-Accelerated Neuroinformatics

- *Fundamental Application in Neuroinformatics*
 - Perform temporal data mining of neural activity
 - Analyze sequential & time-stamped symbolic data to find “interesting” patterns
- *Example Usage: Reverse-engineer the brain*
- GPU Acceleration
 - Performance: 12x speed-up on one GPU card




External Probe



Electrode Array

Computationally expensive!



Accelerated Biocomputing Software


- Pairwise Sequence Alignment
 - mpiBLAST
 - GPU-BLAST
 - GPU-Smith-Waterman
 - *Next Up? Multiple Sequence Alignment*
 - Short-Read Mapping
 - GPU-RMap
 - Molecular Dynamics
 - GPU-GEM
 - GPU-NAB → GPU-AMBER
 - Neuroinformatics
 - GPU-Miner
- Application of Accelerated Biocomputing Software:
- “Finding Missing Genes” ...

Discovering Missing Genes: A Preview

(Source: Dr. João Setubal, Virginia Bioinformatics Institute)

- Premise: *The Case of the Missing Genes*
 - Most genomes completed to date have been detected by a gene-finder program, which can miss real genes.
- Approach
 - Discover *missing genes* in genomes via sequence-similarity computations (i.e., mpiBLAST, <http://www.mpiblast.org/>).
- Infrastructure
 - Create a worldwide supercomputer.
 - 12,000+ processors across six U.S. supercomputing centers.
 - 0.5-petabyte storage system at the Tokyo Institute of Technology.
 - Created *ParaMEDIC: Parallel Metadata Environment for Distributed I/O and Computing* to run on the above worldwide supercomputer.
 - Decouples computation and I/O and drastically reduces I/O overhead.
 - Delivers an additional 27-fold speed-up over mpiBLAST.

Approach to “Missing Genes”

- All-to-all sequence search of 567 microbial genomes:
 - 2.63×10^{14} sequence searches**
 - CompuMatrix, powered by ParaMEDIC
 - 12K+ cores across 7 U.S. supercomputers and a 0.5-petabyte store in Tokyo, Japan: **10 days (~240 hours)**
 - International Storage Challenge Award at 



- IBM Blue Gene/P with 32,768 cores
 - **Less than 12 hours!**

Approach to “Missing Genes” ... Part 2

- How are we finding these missing genes?
 - Obtain 1,297 prokaryotic replicons from NCBI RefSeq.
 - Generate and *categorize* ORFs > 99 bp, using specific start codons.
 - Interested in those ORFs that do *not* share genomic space with any annotated entity.
 - Translate ORFs into amino-acid sequences and construct a BLASTP sequence-similarity search, using mpiBLAST.
 - Resulting alignments are then scanned for evidence of ORFs that are *not* represented in any of the current genome annotations.



Recent Publications

- A. Warren, J. Archuleta, W. Feng, and J. Setubal, “Missing Genes in the Annotation of Prokaryotic Genomes,” *BMC Bioinformatics*, 2009.

Missing genes in the annotation of prokaryotic genomes

Andrew S. Warren^{*1,2}, Jeremy Archuleta², Wu-chun Feng² and João Carlos Setubal^{*1,2}

- P. Balaji, W. Feng, H. Lin et al., “Global-Scale Distributed I/O with ParaMEDIC,” *Concurrency and Computation: Practice and Experience*, 2009.
- P. Balaji, W. Feng, H. Lin et al., “Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer,” *International Supercomputing Conference (ISC '08)*, Dresden, Germany, June 2008. Distinguished Paper Award.



Acknowledgments

- This research was supported in part by
 - **AMD** via the NSF Center for High-Performance Reconfigurable Computing (CHREC)
 - **IBM** and **NVIDIA** via Faculty Awards
 - Eli Lilly & Company *Lilly*
 - National Science Foundation 
- Heshan Lin, Ph.D., Technical Lead on 
www.mpiblast.org
- Collaborators
~50 researchers worldwide + 15 students



Wu FENG, Ph.D.

feng@cs.vt.edu

SyNeRG  Laboratory

<http://synergy.cs.vt.edu/>



SUPERCOMPUTING
in SMALL SPACES

<http://sss.cs.vt.edu/>

THE GREEN
500 

<http://www.green500.org/>



<http://www.mpiblast.org/>