




# Convey Application Performance

Steve Wallach  
swallach "at" conveycomputer "dot" com


Convey $\leq$  Convex++



# What is BIG DATA?




Dilbert.com DilbertCartoonist@gmail.com



Venture capitalist

# Hybrid-Core Architecture



- x86 host system runs Linux and x86\_64 binaries
- Coprocessor runs application specific personalities
- Globally shared memory presents same view to threads on x86 cores and Coprocessor

# Market Segments

## Data Intensive Computing

### Bioinformatics

- Bioinformatics Suite
- Genome Assembly
- 3rd Party Applications

### High-Performance Analytics

- Predictive Analytics
- ID Resolution
- Internet Search
- Risk Analytics
- Know your Customer
- Fraud Detection

### Government/Security

- Predictive Analytics
- ID Resolution
- [REDACTED]
- Know your customer
- [REDACTED]

### Education/Research

- Heterogeneous Computing
- Algorithms
- FPGA Research

# Bioinformatics on Convey



	Organization	Application	Approx Server Speed Up	Usage
Convey Bioinf Suite	Convey Bioinformatics Suite	BWA	10x	Reference Mapping
	Convey Bioinformatics Suite	Velvet/CGC	6x / memory reduction	De Novo Assembly
	Convey Bioinformatics Suite	Kmer Counter	fast, multiple kmer, long kmer	Read Analysis for Assembly
	Convey Bioinformatics Suite	SWSearch	15x	Smith-Waterman Search
	Convey Bioinformatics Suite	BLAST(p,x)	6x	Protein Database Search
Available	CLC bio	CLC Genomics platform	-	Analysis, workflows, visualization
	Michigan Technological University	PCAP	12x	Overlap-based Assembly
	Jacquard	FHAST (BowTie)	11x	Burrows-Wheeler Aligner
	University of California San Diego	InsPect	13x	Protein Assembly with PTMs
Presented	Iowa State University	Sheppard	25x	Short-read Mapping
	Technical University Crete	BLASTn	3x	Nucleotide Sequence Search
	University of California Los Angeles	Fluid Registration	4x	Medical Imaging
	University of South Carolina	BEAGLE, MrBayes, BEAST, Garli	8x	Bayesian and Maximum Likelihood phylogenetics library
	University of Washington	BFAST	30x	Short-read Mapping
Project Initiated	Bielefeld University / Paderborn Uni	ENHANCE		Taxonomic Profiling; Suffix Tree Read Mapping
	Boston University	BLAST(p,x)		Protein Database Search
	Chinese Academy of Sciences	PerM aligner		Periodic Spaced Seeds Read Mapping,
	CHREC	Bowtie2, CLUSTALW		Read Mapping, Multiple Sequence Aligner
	CSIR-ACE	Unafold		RNA Structure Prediction
	Free University of Berlin	SeqAn		Sequence Analysis Library
	The Broad / Green Mountain	GATK		Locus walker, HMM; UnifiedGenotyper, HaplotypeCaller
	Iowa State University	RMAP		Short-read Mapping
	Synective	GROMACS		Molecular Dynamics
	University of Washington	Infernal		RNA Secondary Structure
	Virginia Bioinformatics Institute	Various		Molecular Dynamics, Bioinformatics

# Smith-Waterman Performance



*1000 proteins vs nr database*

**14.5x speedup over  
fastest software  
approach**


- SIMD approach 15-20x faster than naïve implementation

12c x86: 2 Intel X5670 2.93GHz processors (12 cores),  
192GB DDR3, stripe 4 @ 600GB SATA disks

HC-2: 96GB DDR3 (host), 16GB SG (coproc), stripe 4 @ 600GB SATA disks

HC-2ex: 192GB DDR3 (host), 64GB SG (coproc), stripe 4 @ 600GB SATA disks


Smith-Waterman Performance



# De Bruijn Graph Algorithm

Popular for short-read *de novo* assembly

“Memory size is the biggest difficulty... and that can be quite a constraint for users” - Daniel Zerbino, author of Velvet



- Implemented in memory as hash table-based binary tree
- Requires random access to large amounts of memory
- Memory bandwidth is limiting factor

<sup>1</sup> “Velvet: algorithms for *de novo* short read assembly using de Bruijn graphs”, Daniel R. Zerbino and Ewan Birney, EMBL-European Bioinformatics Institute, Genome Res. 18 (2008) 821.

# Convey Graph Constructor

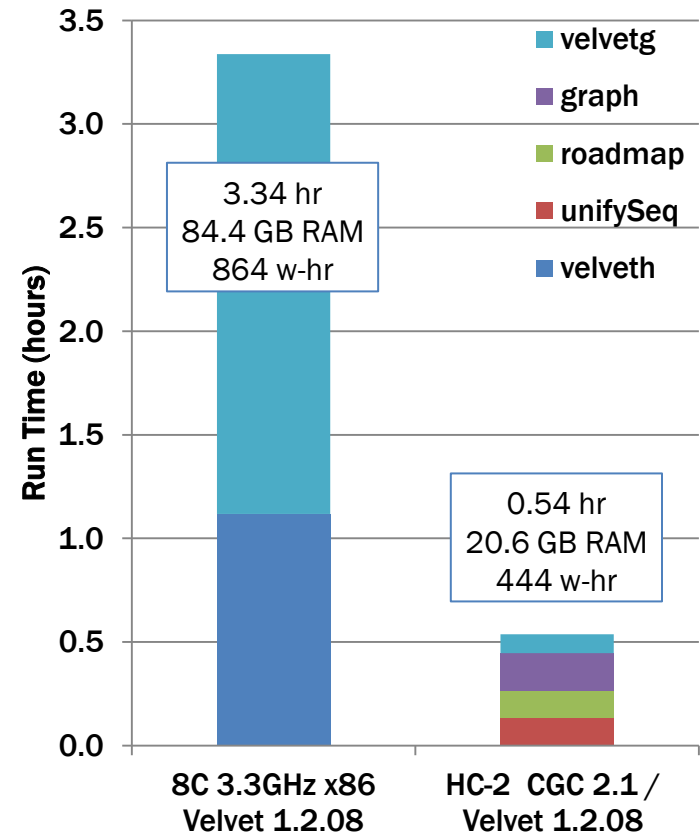


## Enables Large Genomes

- **6.2x speed up depends on**
  - Data set size
  - Kmer space complexity
- **RAM reduced 76%**
  - Data types / structures
  - Automated roadmap partitioning
- **1.9x Power Performance**

HC-2: 2 Intel E5-2643 processors (8 cores total; 3.33GHz), stripe 4@1TB 10KRPM disks, 384GB DDR3 (host), 64GB SG (coprocessor)  
X86: As above, host only

## JGI 40 Gbp Cow Rumen



*“Convey’s GraphConstructor offers a new approach to help researchers ... to achieve better assemblies or look at bigger jobs such as metagenomic or mammalian genome samples”*



# Grand Challenge Metagenomes

*DOE Joint Genome Institute*

## HC Capability

- 800 Gbp assembly possible
- Reasonable time
- No prefiltering

Data Set	Bases (B)	Contigs (M)	Span (Gbp)	% Mapped	Max RAM	Runtime (hrs)
Wisconsin Restored Prairie	53	0.6	0.3	11.5	60	19.25
Wisconsin Switchgrass	58	7.9	1.9	4.9	60	1.72
Wisconsin Continuous Corn	192	22.1	5.1	9.0	112	15.40
Iowa Continuous Corn	197	2.7	0.9	7.6	215	17.92
Wisconsin Native Prairie	199	7.1	2.2	11.3	445	35.73
Kansas Continuous Corn	271	5.0	1.3	3.7	239	20.25
Iowa Native Prairie	367	10.7	3.0	12.9	534	39.90
Kansas Native Prairie	598	27.8	9.3	8.8	503	107.47
Wetland Surface Sediment	772	323.2	76.3	22.8	716	129.33
Rumen (in progress)	1153					457.17

*“My plan was to go through them in order of size until the system broke, but I was able to run through all the Great Prairie samples without any problems.”*

**Dr. Alex Copeland**  
**Joint Genome Institute**

CGC run on HC-1: 128GB (host), 64GB (coproc), stripe 2@1TB SATA for all except for WI Switchgrass run on HC2. Velvet run on 1 TB SMP server


“Assembly of Large Metagenomes Using a Hybrid Core Computer” Alex Copeland, Sequencing, Finishing and Analysis in the Future Conference, Santa Fe, NM, June 7 2012.

“Efficient Assembly of Grand Challenge Metagenomes on a Hybrid-Core Architecture” Alex Copeland, et al., Advances in Genome Biology and Technology, Marco Island, FL, Feb 20 2013

# Memcached

## key-value store

- **Implemented as network service**
  - binary or ascii commands sent over sockets
  - values stored in memory for fast access
- **Typically used as a cache**
  - no built-in redundancy or resilience
  - throughput within latency tolerance constraints




Ricky Ho, DZone

# Convey Memcached Throughput



- **Convey HC-2ex compared to 16-core SandyBridge**
  - Dual 10GBe
  - Three clients running workload generators
- **Get Throughput**
  - Servers preloaded with data (36byte keys, 100byte values)
  - 100% hit (data returned for each request)



# Levenshtein Algorithm

- **Measures difference between strings**
  - number of changes required to change one string to another
- **Difference between ‘frog’ and ‘princess’ is 7**
  - frog -> prog
  - prog -> prig
  - prig -> prin
  - prin -> princ
  - princ -> prince
  - prince -> princess
  - princess -> princess

$$\text{lev}_{a,b}(i,j) = \begin{cases} \max(i,j) & , \min(i,j) = 0 \\ \min \begin{cases} \text{lev}_{a,b}(i-1,j) + 1 \\ \text{lev}_{a,b}(i,j-1) + 1 \\ \text{lev}_{a,b}(i-1,j-1) + [a_i \neq b_j] \end{cases} & , \text{ else} \end{cases}$$

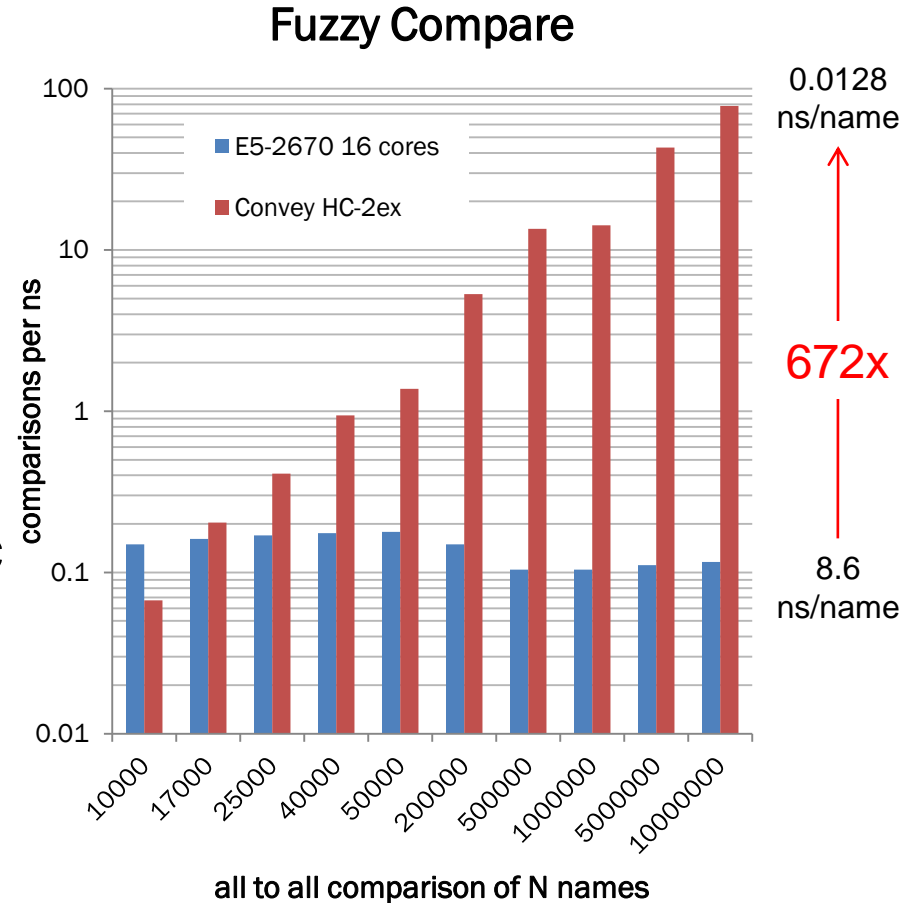
formula from Wikipedia “Levenshtein distance”

	p	r	i	n	c	e	s	s
f	1	2	3	4	5	6	7	8
r	2	1	2	3	4	5	6	7
o	3	2	2	3	4	5	6	7
g	4	3	3	3	4	5	6	7


# Convey Fuzzy String Compare



- Fuzzy String compare personality and application
- Levenshtein personality
  - 2048 10x10 comparators
  - up to 307B results per sec
- **672x speedup over a SandyBridge server**
  - 10M names all-to-all



# MX - 100 Single-node Block Diagram



2013 - IDC - TUCSON



# Graph 500 Results: Performance/Power

## Single Node, All Scale (by GTEPS)

Rank	Installation Site	Machine	Number of nodes	Number of cores	Problem scale	GTEPS
★	Convey Computer Corporation	fox6 / Convey MX-100	1	12	29	14.56
★	Convey Computer Corporation	thunder4 / Convey HC-2ex	1	12	27	11.47
★	Argonne Nat. Lab	HC-2ex, host-210	1	8	27	11.44
★	Univ. Calif. Riverside	yosemite / HC-2ex, host-210	1	8	27	11.44
★	Convey Computer Corp.	coconino / HC-2ex, host-210	1	8	27	11.44
25	Chuo University	Intel(R) Xeon(R) CPU E7- 4870 @ 2.4	1	40	26	8.15
27	Convey Computer Corporation	thunder4 / Convey HC-2ex	1	12	27	7.85
31	Convey Computer Corporation	Vortex / Convey HC-1ex	1	4	27	6.64
33	Virginia Bioinformatics Institute	convey-ex01/Convey HC-1ex	1	4	27	6.00
38	Chuo University	Intel Xeon E5-2690 2.90GHz (2 socke	1	16	29	4.22
40	Institute of Computing Technology, C	I840	1	32	30	3.69
41	Chuo university	AMD Opteron(tm) Processor 6174 2.	1	48	29	3.24
43	Chuo University	GraphCREST-M48 AMD Opteron 617	1	48	29	2.42
44	Chuo University	Intel(R) Xeon(R) CPU E7-4870 @ 2.40	1	40	30	2.16
45	Imperial College London	Convey HC-1 Server	1	6	24	1.72
46	SNL	Dingus	1	4	28	1.72
46	SNL	Wingus	1	4	27	1.72
47	Convey Computer Corporation	Vortex	1	4	28	1.61
47	Bielefeld University, CeBiTec	Convey01	1	4	28	1.61
47	Convey Computer Corporation	Hc1-d	1	4	28	1.61
48	LBL/NERSC	Convey2	1	4	28	1.50
53	Chuo University, Tokyo Japan	GraphCREST-W12	1	12	0	1.18
63	Chuo University, Tokyo Japan	GraphCREST-M48	1	48	27	0.83
64	Convey Computer Corporation	Convey XC-1ex	1	4	27	0.76
65	Intel Research	Westmere E7-4870 2.4GHz	1	40	27	0.70
77	LLNL	Leviathan + FusionIO	1	40	36	0.05
77	LLNL	Appro	1	32	34	0.05

2013 - IDC - TUCSON  
★ = Convey

