Exploring Accelerating Science Applications with FPGAs

Dr. Olaf O. Storaasli
Future Technologies Group
Computer Science & Mathematics Division
Oak Ridge National Laboratory

RSSI 2007 NCSA UIUC

Research Team

Olaf Storaasli

Dave Strenski

Acknowledgment

This research was sponsored by the Laboratory Directed Research & Development Program of ORNL managed by UT-Battelle for the U. S. Department of Energy under Contract No. DE-AC05-00OR22725. The U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

Contents

Background: FPGAs, Genome Sequencing

Results: FASTA for 3 openfpga.org Cases

Goal: Speed Supercomputers with FPGAs

Context

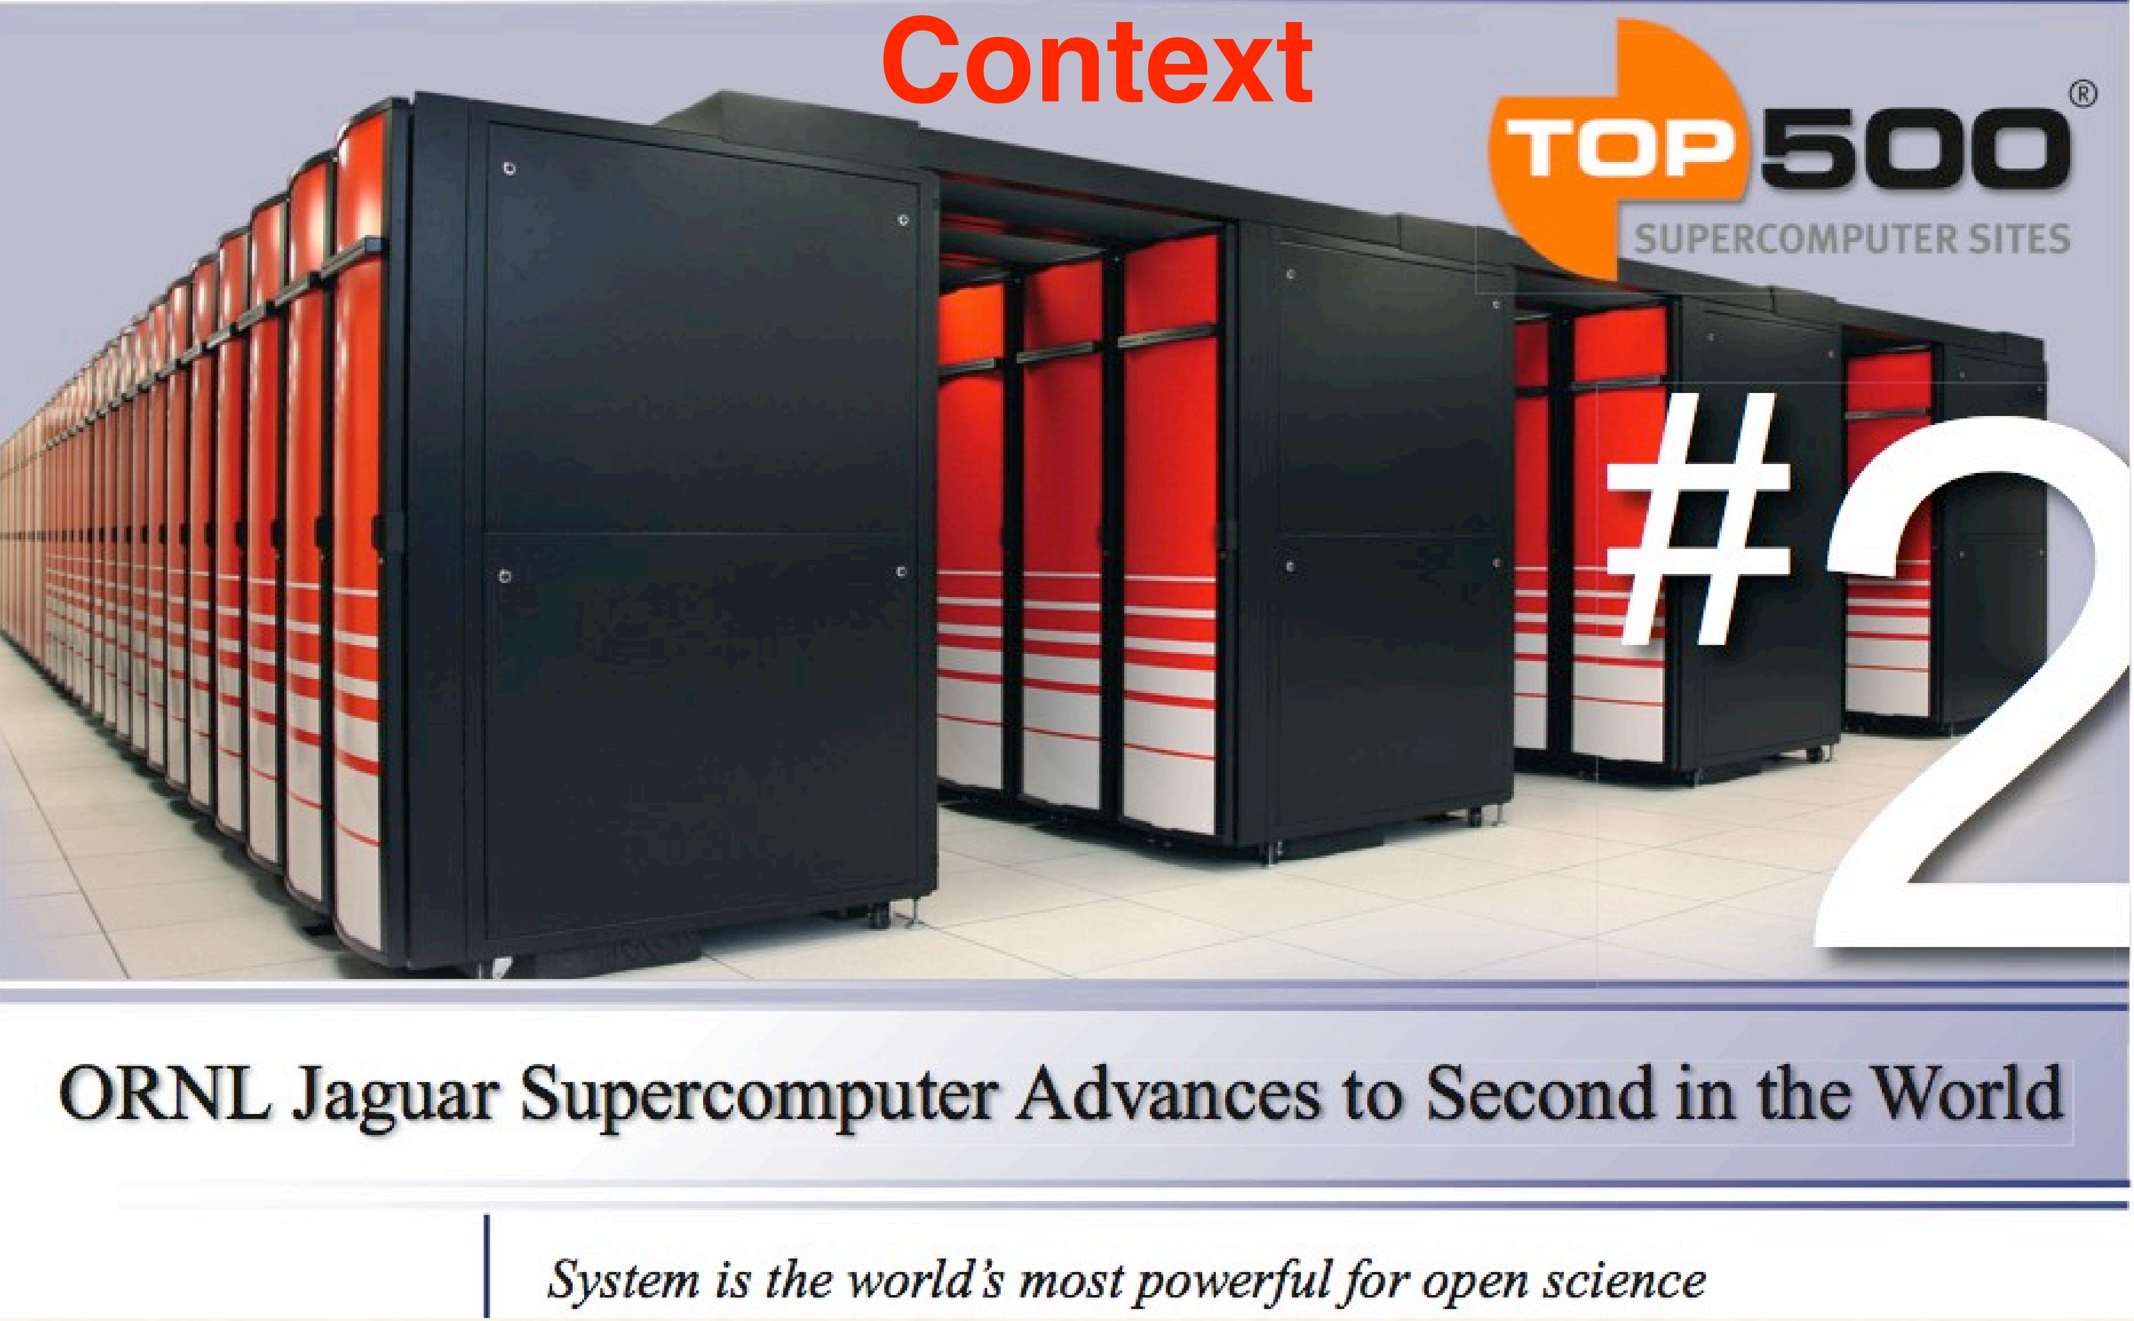

102TF 2Q07 => 250TF 4Q07 => 1PF 4Q08

Future Technologies Group

Oak Ridge National Laboratory

Future Supercomputer Technologies

Commodity: 2^n core 2 GHz chips

Special: El Dorado, Cyclops, PiM

Accelerators

- **FPGA: DSP => HPEC => HPC** <==

- **Cell: IBM, Sony, Toshiba**

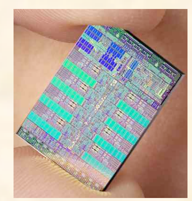

- **GPUs: on μP?**

- **Array: ClearSpeed**

Xilinx Virtex4 FPGA

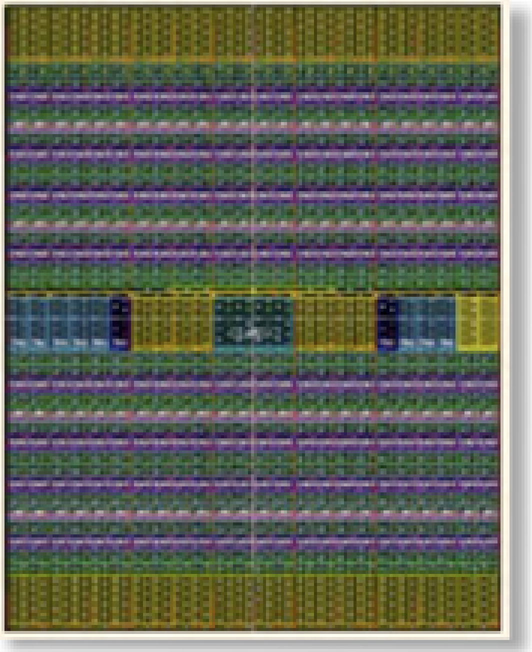

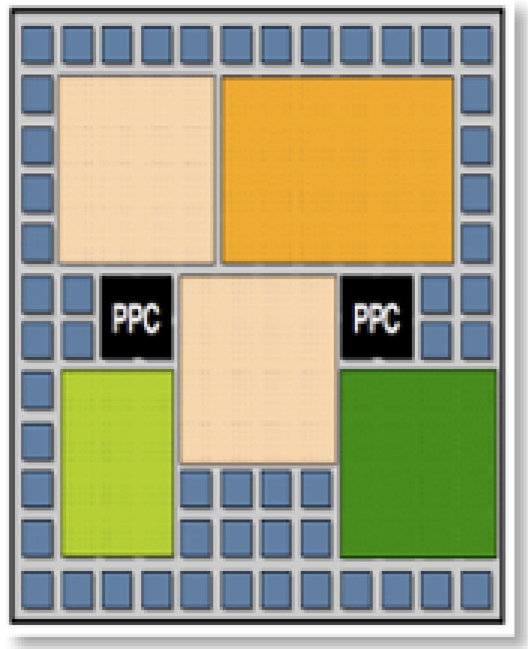

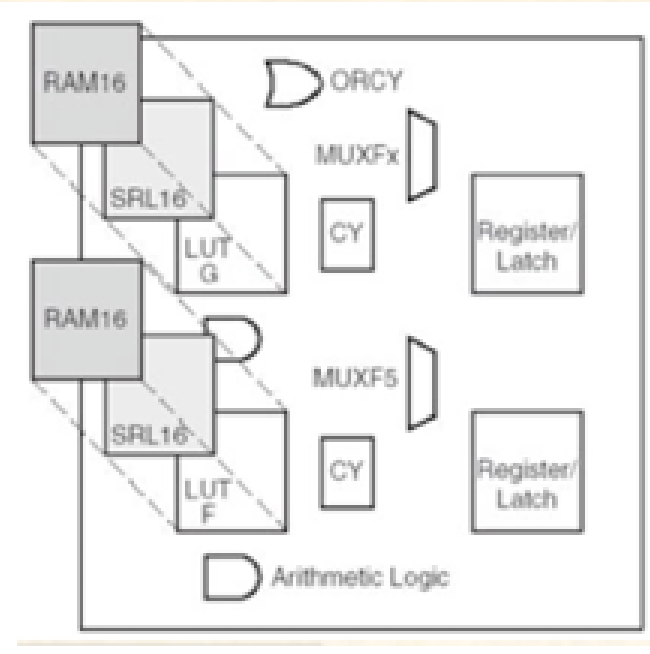

Logic array **PPC Processors** **67,584 slices***

*LX-160 (89,088 on LX-200)

Future Technologies Group ornl

OAK RIDGE National Laboratory

Why FPGAs?

- **Performance**—optimal silicon use (maximize parallel ops/cycle)
- **Rapid growth**—Cells, Speed, I/O
- **Power**—1/10th CPUs
- **Flexible**—tailor to application

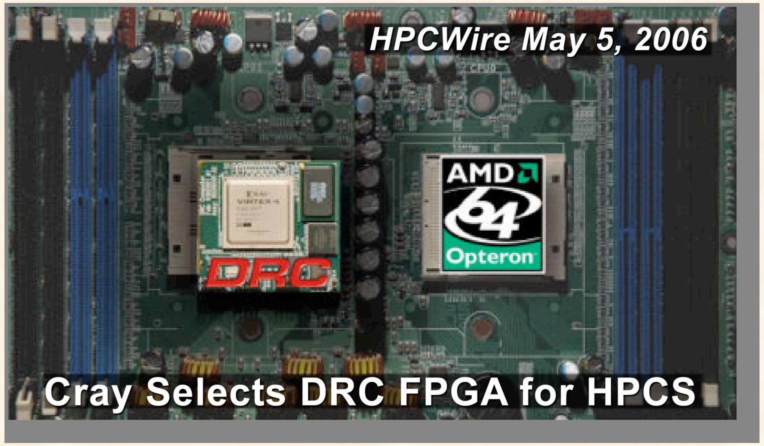

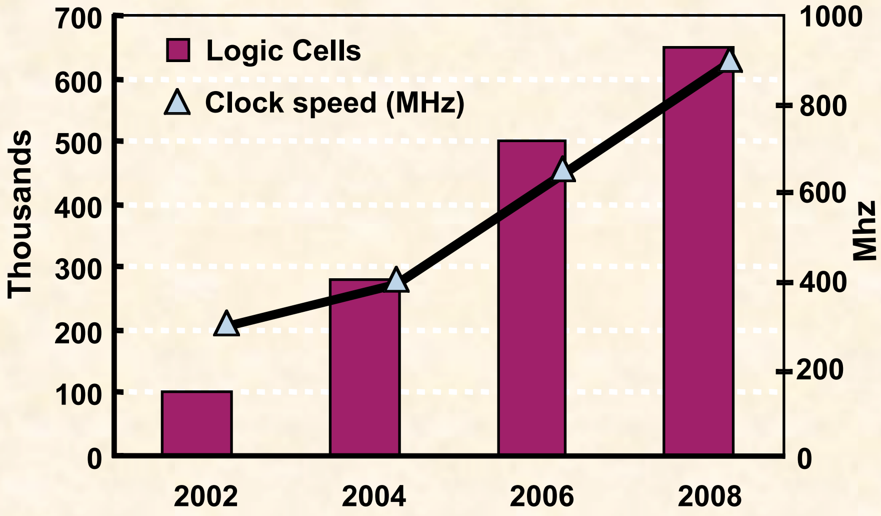

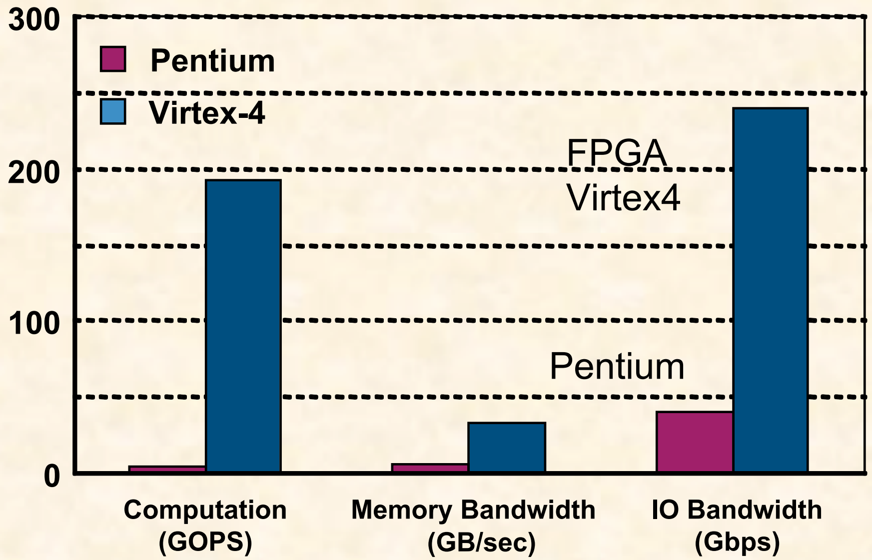

Future Technologies Group ornl

OAK RIDGE National Laboratory

Growing HPC Industry Interest

CRAY THE SUPERCOMPUTER COMPANY

"After exhaustive analysis, Cray concluded that, although multi-core commodity processors will deliver some improvement, exploiting parallelism through a variety of processor technologies using scalar, vector, multithreading & hardware accelerators (e.g., FPGAs or ClearSpeed co-processors) creates the ***greatest opportunity*** *for application acceleration"*

HPC Future**, Steve Scott, Cray CTO:**

HPCWire **24 March '06**

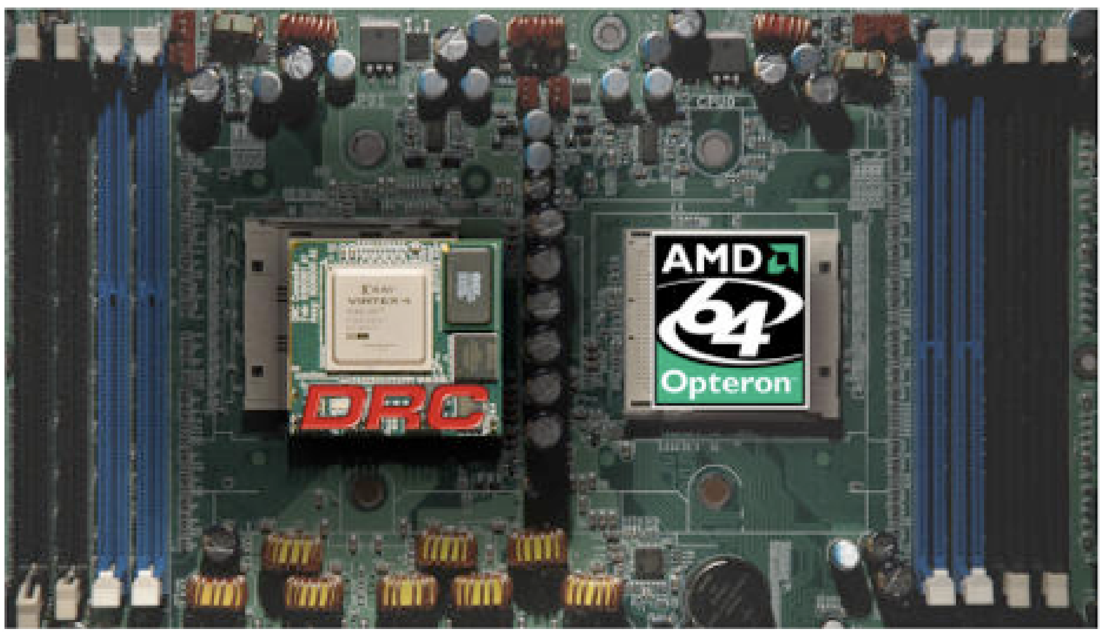

Cray Selects DRC FPGA Coprocessors for HPCS & future Supercomputers (*HPCWire*, 5 May '06)

sgi Virtex4 FPGA blades to:

"Accelerate mission-critical applications by over 100x"

SRC FPGAs + FP chip

Potential: Petaflops/Exaflops at reduced power

Future Technologies Group ornl

Oak Ridge National Laboratory

Smith-Waterman Algorithm Scoring

Query Sequence

Database Sequence

	0	A	C	G	T	...	C
0	0	0	0	0	0	0	0
A	0	2	0	0	0	2	0
C	0	0	4	2	1	0	2
G	0	0	2	6			
A	0						
A	0						
C	0						
...	0						
G	0						

1. Initialize row & column 1 to 0
2. Score matches from upper left
3. Add to above-left score *(2+4=6)*

Future Technologies Group ornl

Oak Ridge National Laboratory

Search34 Computation Profile

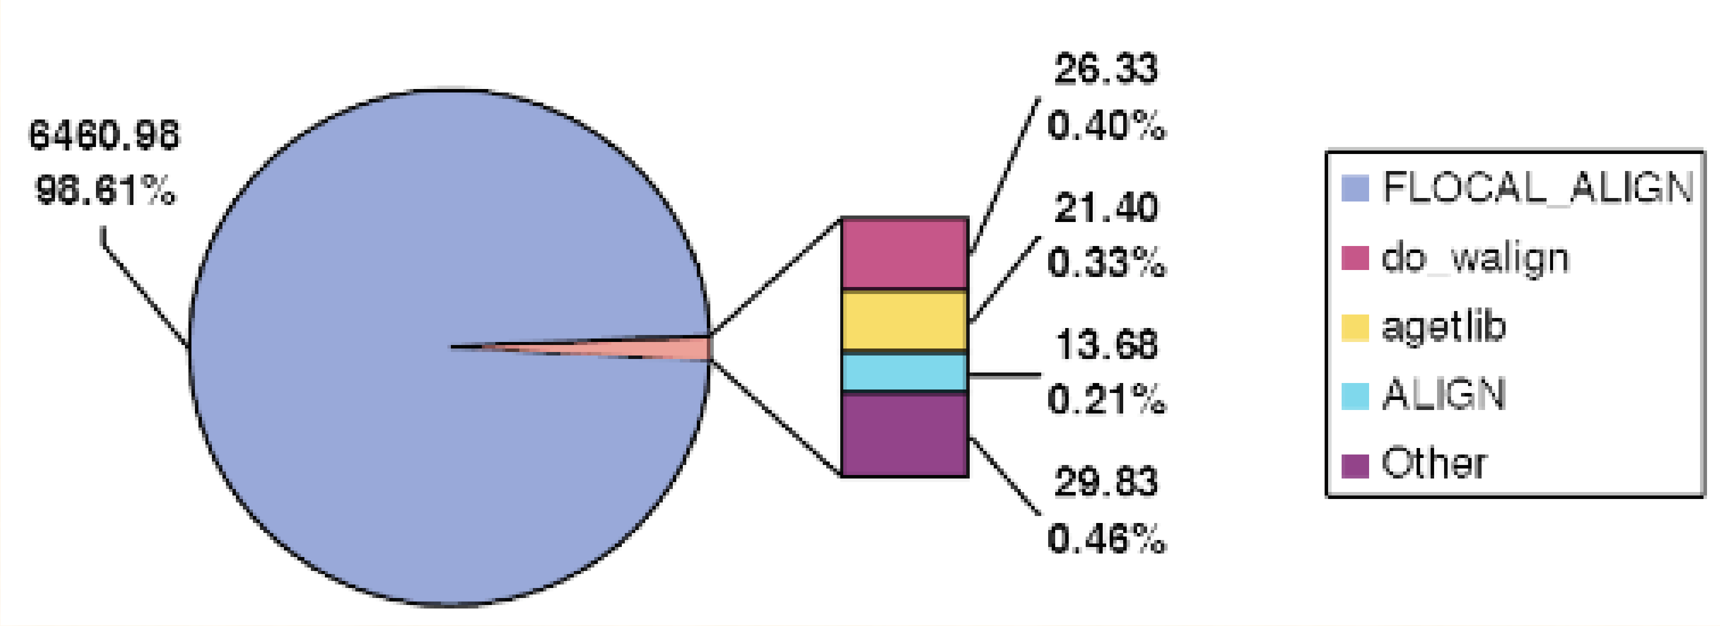

98.61% is FLOCAL_ALIGN

Smith-Waterman Pipeline

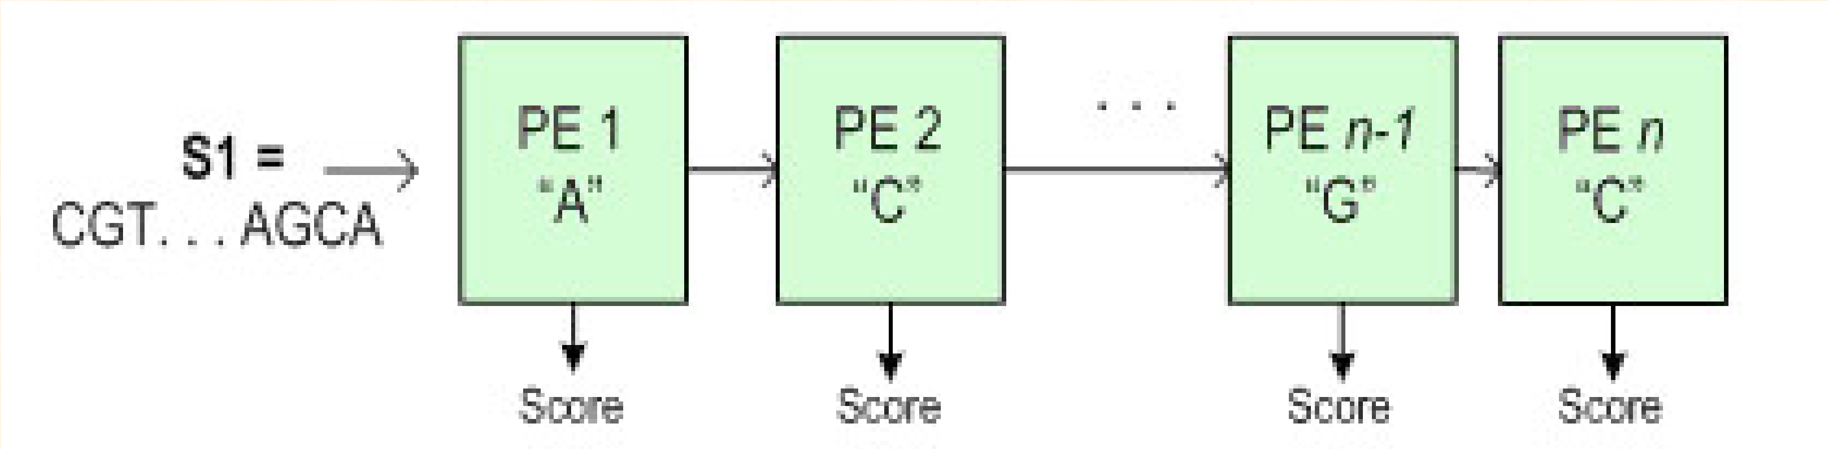

1. Query character preloaded into each PE
2. String S1 shifted thru pipe to compare
3. Score generated

Future Technologies Group ornl

OAK RIDGE National Laboratory

Smith-Waterman

Parallel Score Calculation

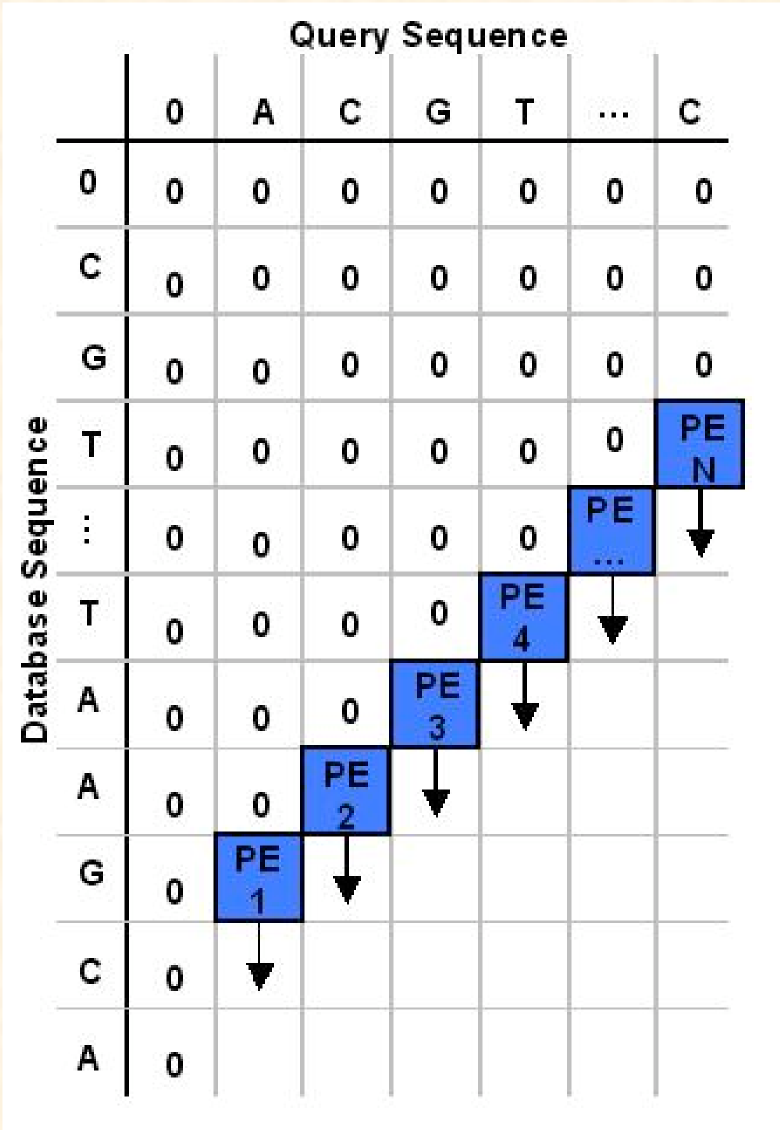

Overall Algorithm

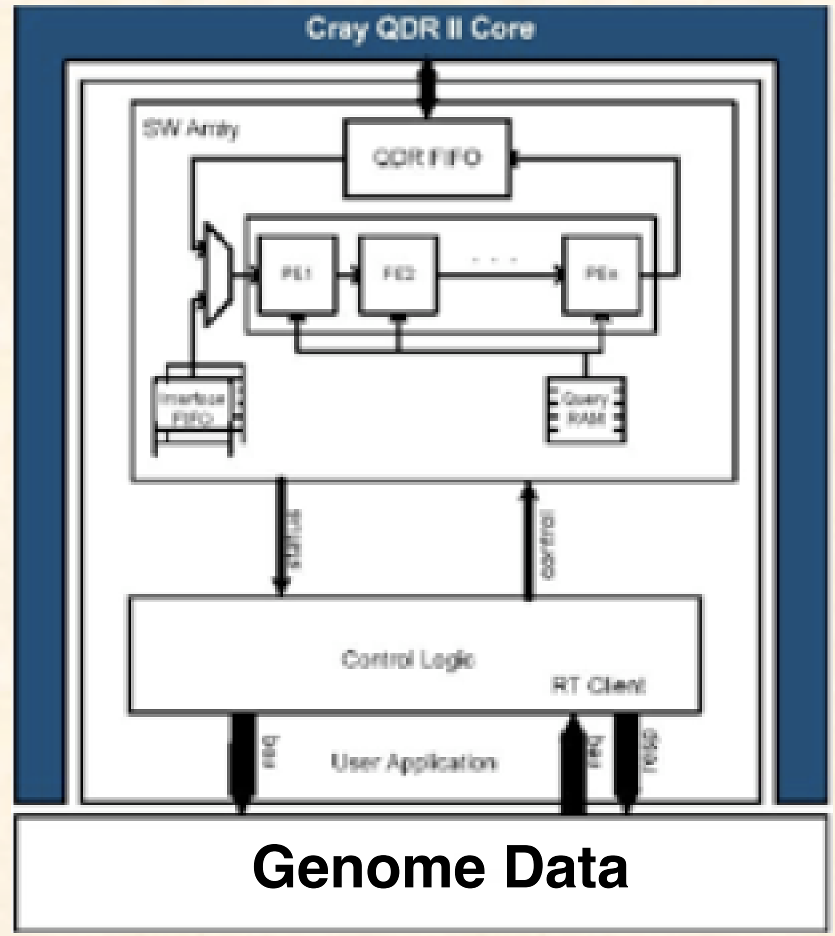


IBM Cell BladeCenter

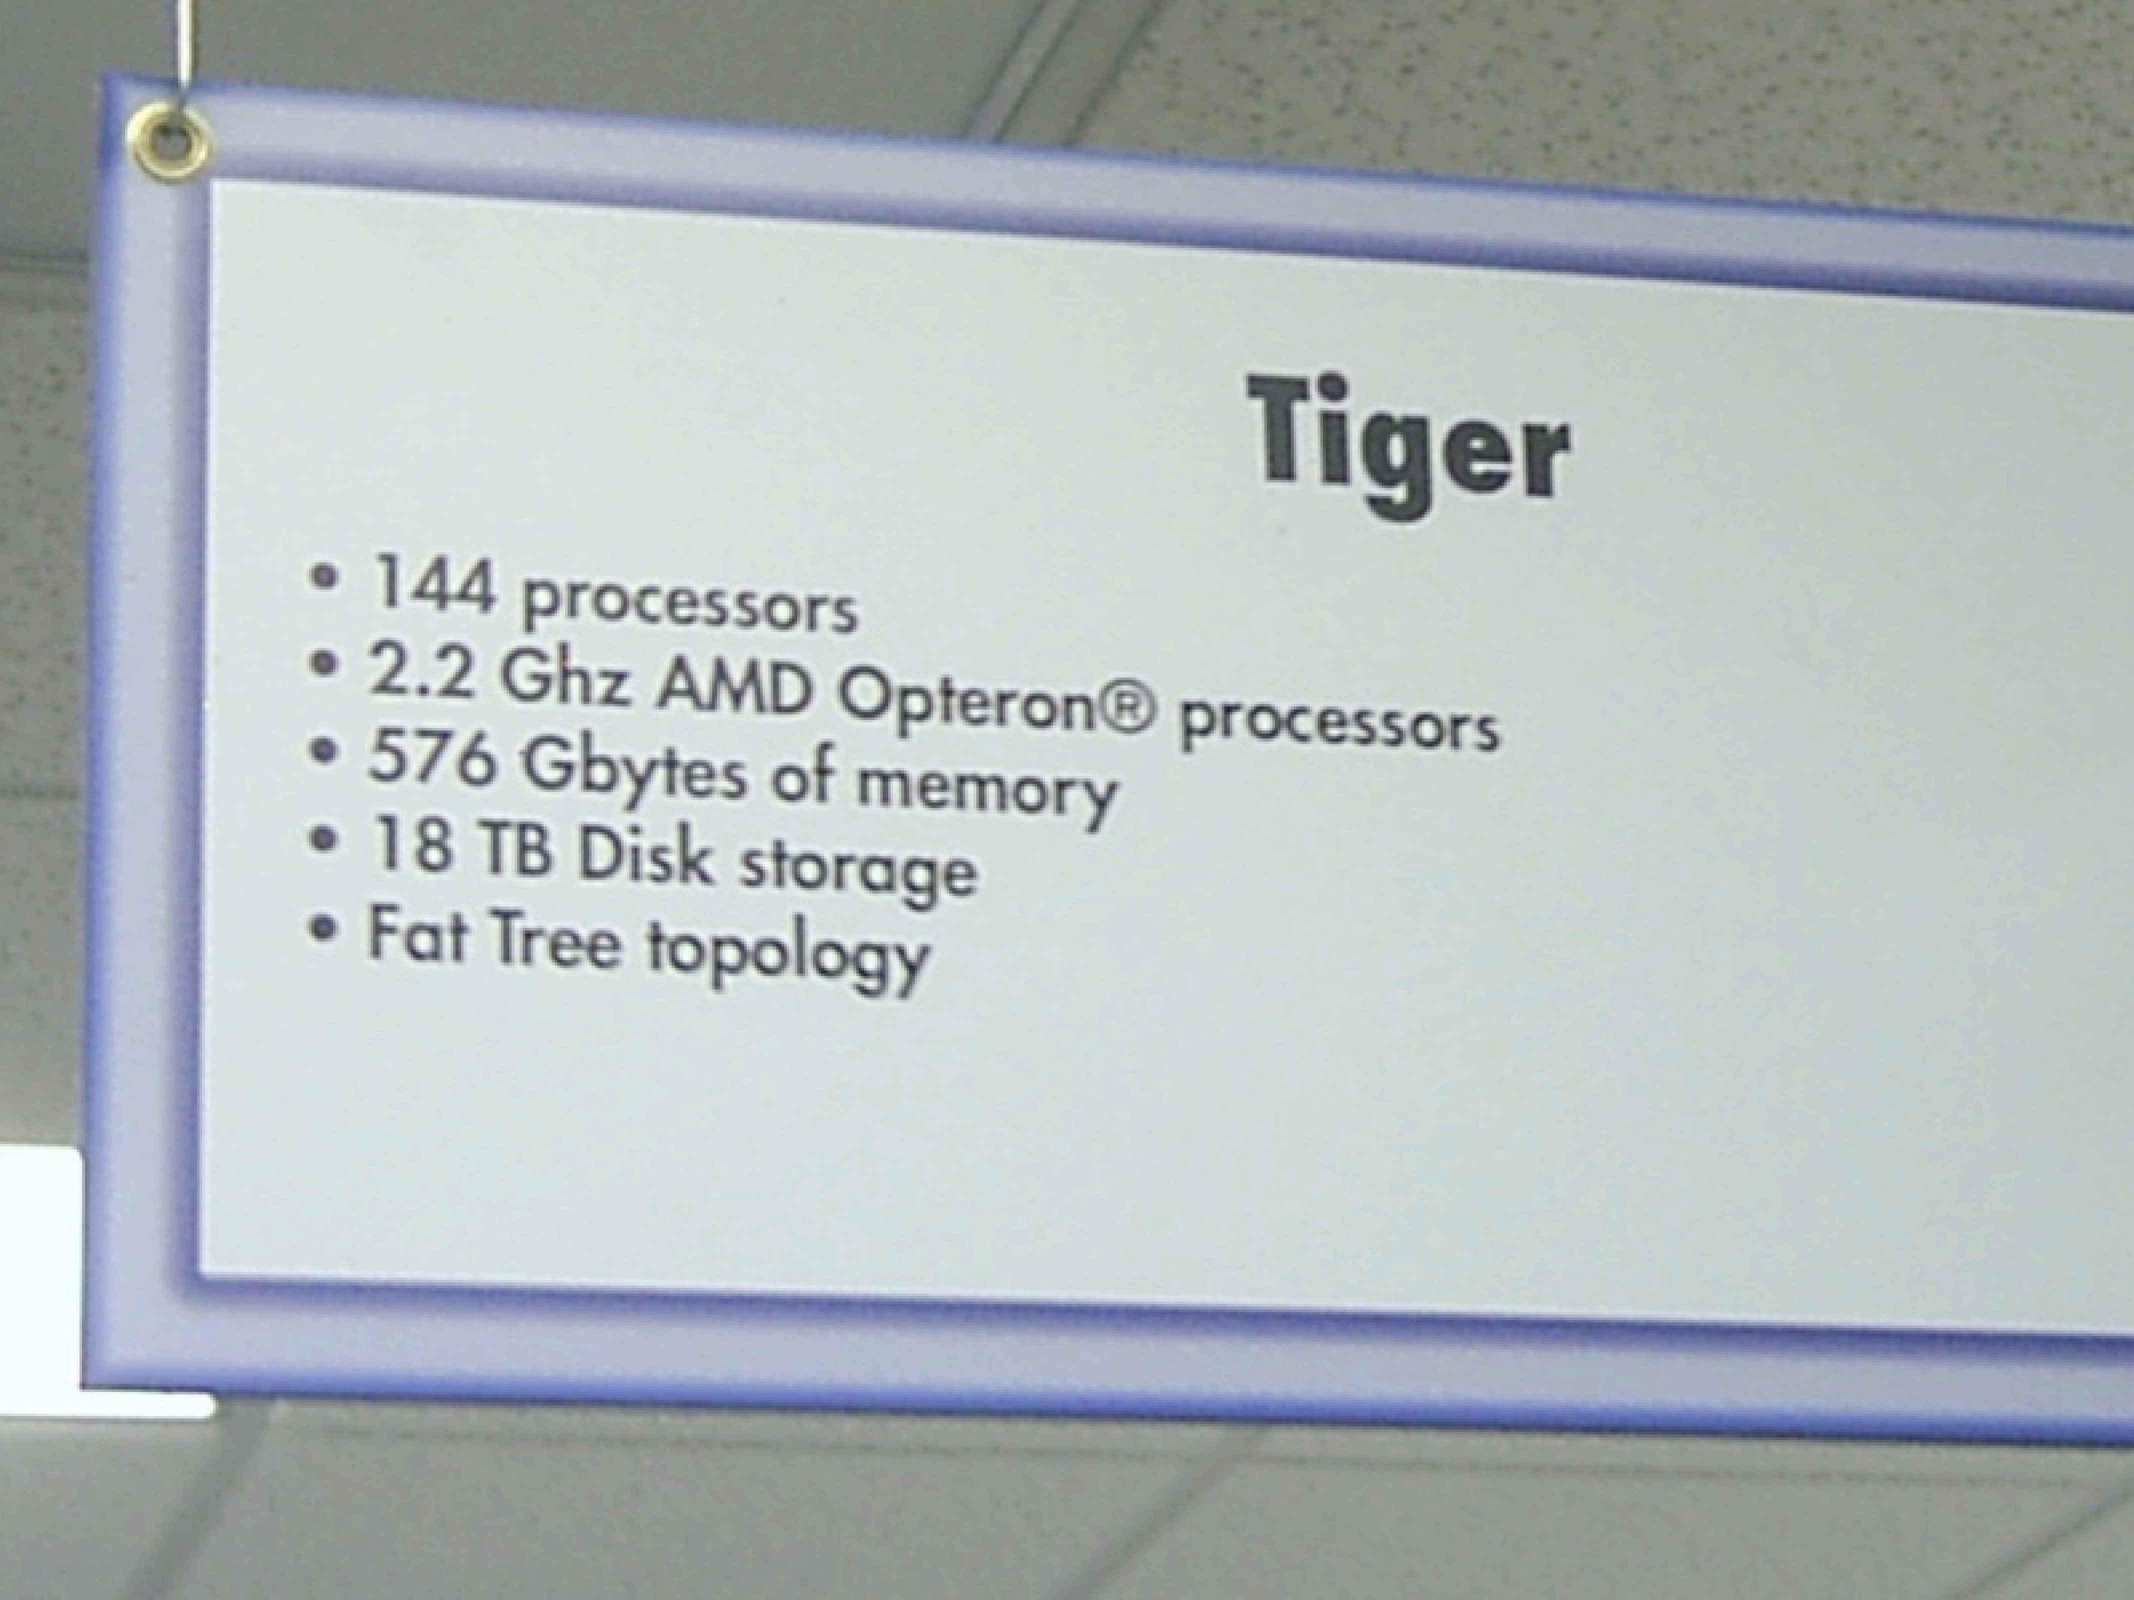
Tiger
• 144 processors
• 2.2 Ghz AMD Opteron® processors
• 576 Gbytes of memory
• 18 TB Disk storage
• Fat Tree topology

Cray XD1

- 144 processors => 633 GFLOPS peak

  (12 chassis 2.2GHz AMD Opterons)

- 6 Xilinx Virtex2Pro 5M gate FPGAs
- 576 GB Memory, 18TB disk
- Cray terabit backplane
- Porting applications

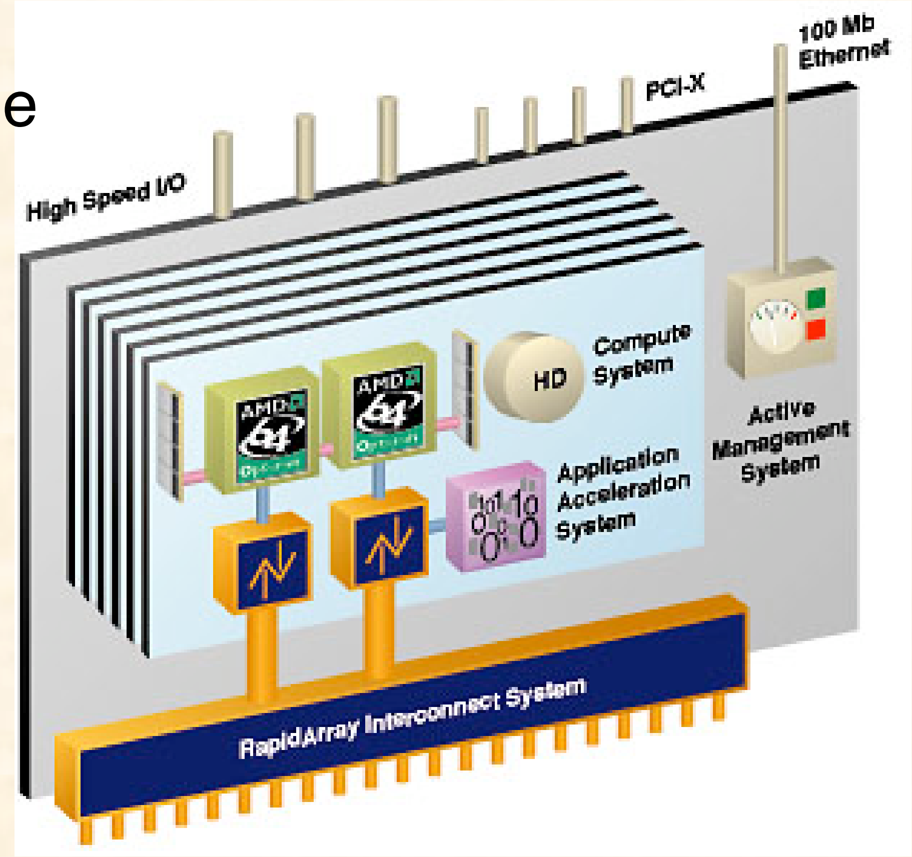

FPGA Coding: Graphical vs. Text

Gauss matrix solver

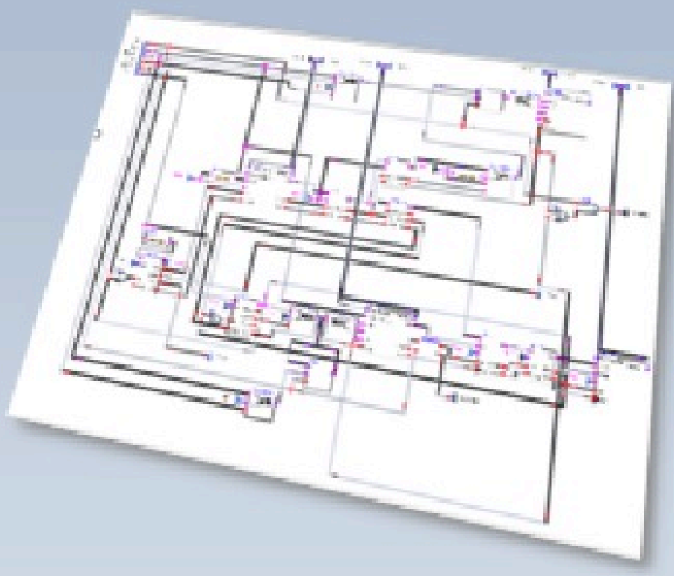

Viva: Graphical Icons—3-dimensional

Compiler, simulator, and debugger

MitrionC: Text/flow—1-dimensional

+ Carte/SRC, CHiMPS-VHDL/Xilinx , DSPlogic

Future Technologies Group ornl

Oak Ridge National Laboratory

FPGA Tool Comparison*

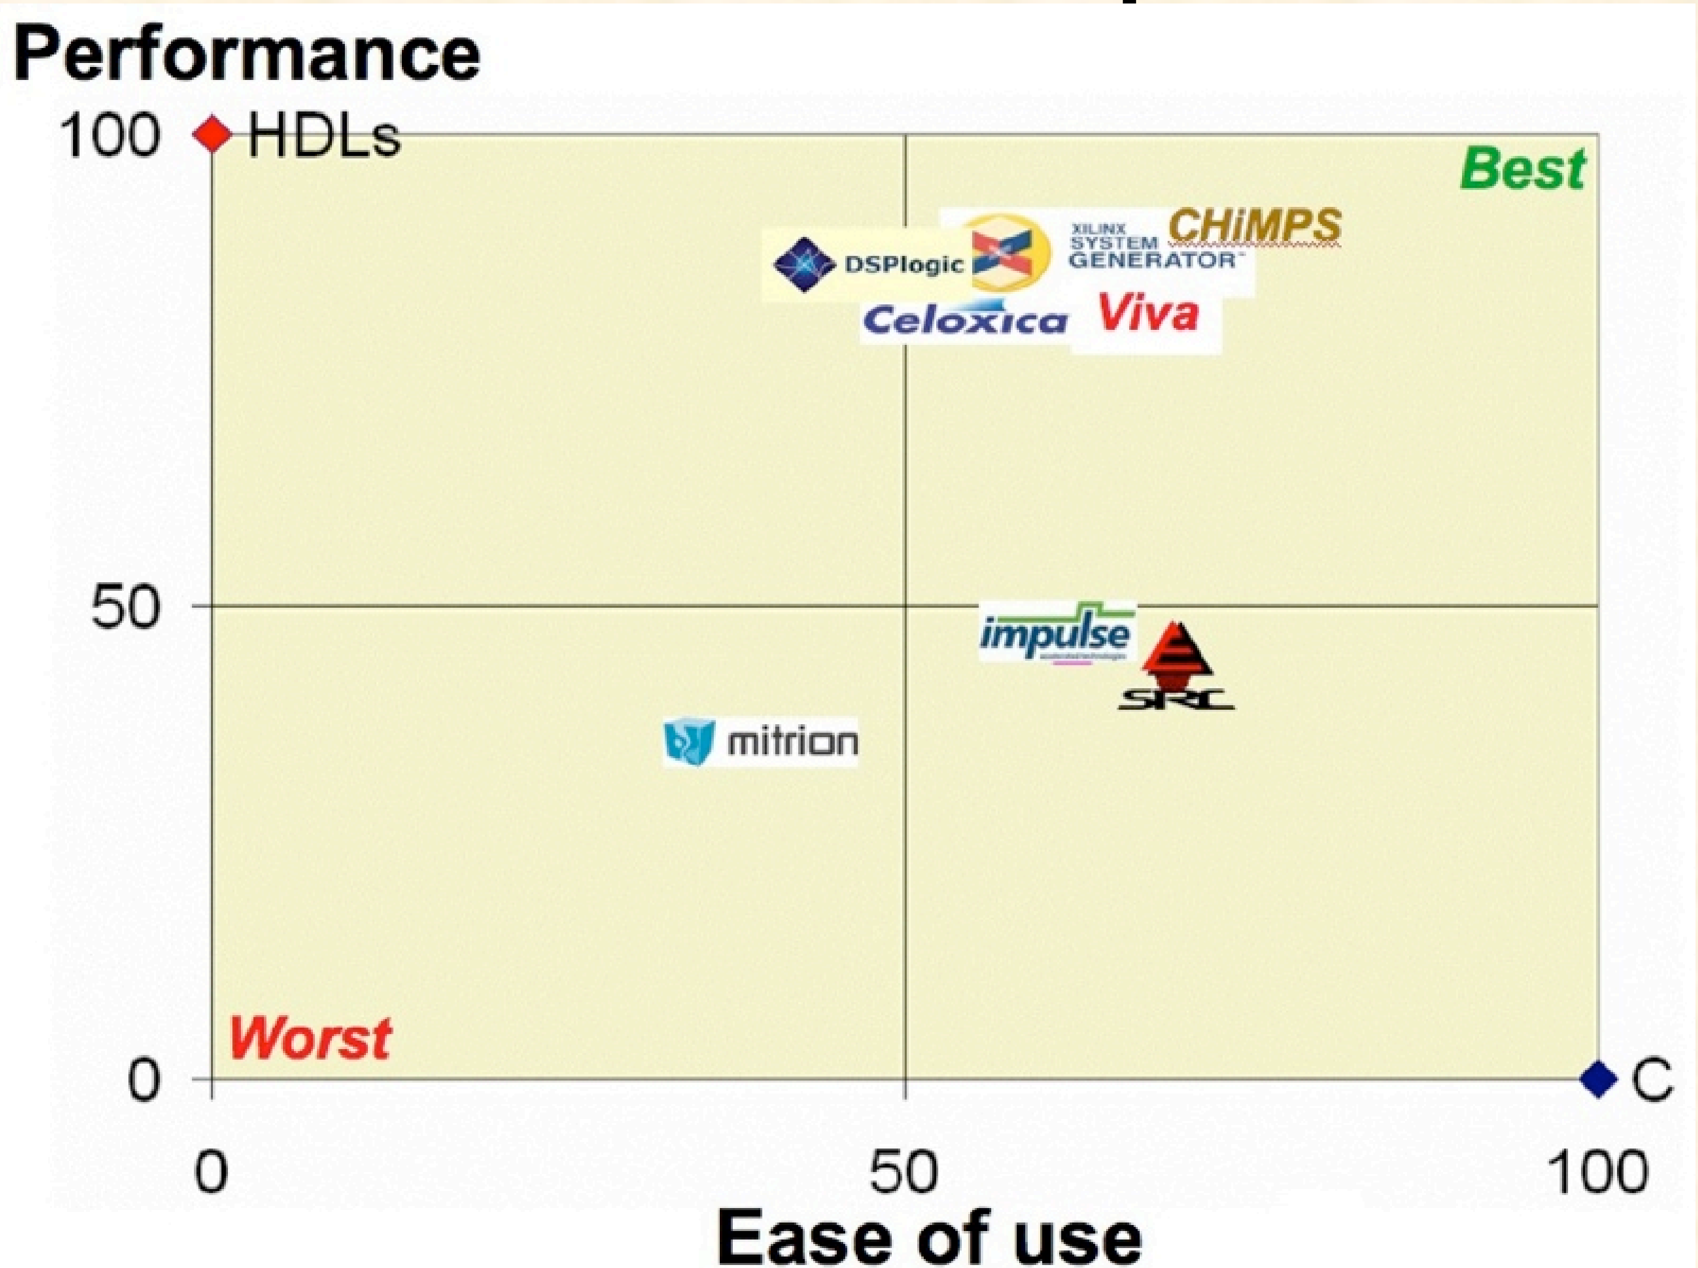

Future Technologies Group ornl

*courtesy CHREC(GWU)

Oak Ridge National Laboratory

Climate/Weather Code CHiMPS FPGA Port

Profile

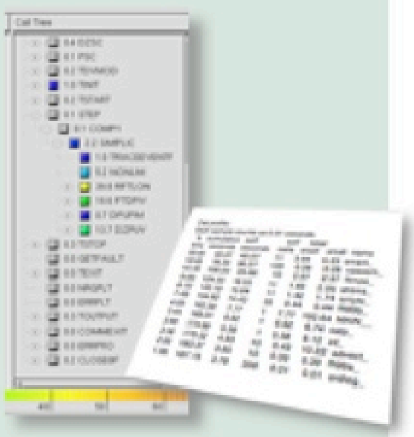

Find parallelism: 80% FFTs

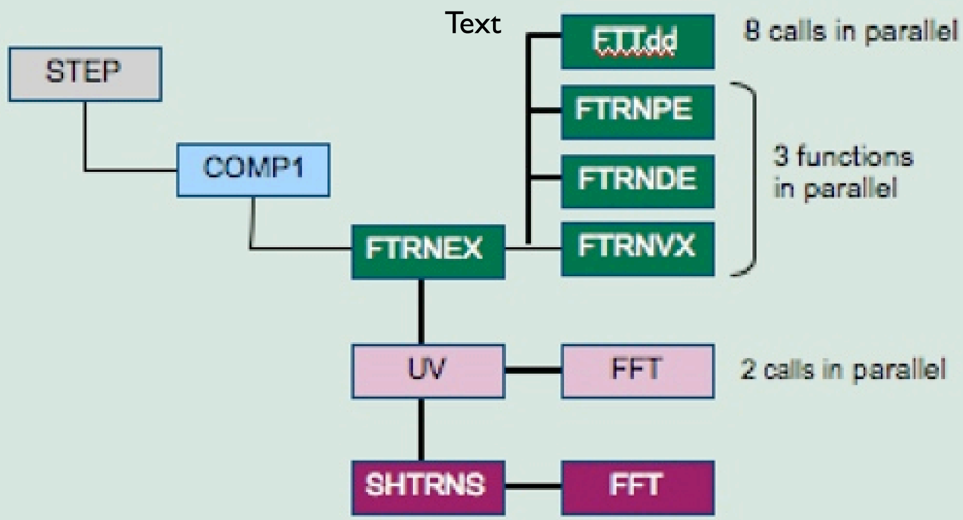

Goal

More GF/$ GF/Watt

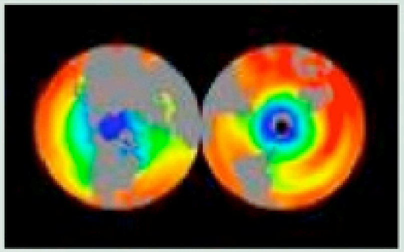

Model faster

Future Technologies Group

Oak Ridge National Laboratory

Openfpga.org Smith-Waterman Benchmark

- ***FASTA* (University of Virginia) application**

  http://fasta.bioch.virginia.edu

- **Uses search34 code & Cray SWA core**

- **Human Genome Data: 4GB compressed**

  3685 searches (MPI on ORNL Cray XD1)

Alignment of ACGAACCCCTTGC and ACGTATGC

	0	A	C	G	T	A	T	G	C
0	0	0	0	0	0	0	0	0	0
A	0	2	0	0	0	2	0	0	0
C	0	0	4	2	1	0	1	0	2
G	0	0	2	6	4	3	2	3	1
A	0	2	1	4	5	6	4	3	2
A	0	2	1	3	3	7	5	4	3
C	0	2	4	2	2	5	6	4	6
C	0	0	2	3	1	4	4	5	6
C	0	0	2	1	2	3	3	3	7
T	0	0	0	1	3	2	5	3	5
T	0	0	0	0	3	2	4	4	4
G	0	0	0	2	1	2	2	6	4
C	0	0	2	0	1	0	1	4	8

Final alignment

A	C	G	A	A	C	C	C	T	T	G	C
|	|	|		|					|	|	|
A	C	G	T	A	-	-	-	-	T	G	C

Results

Case 1: ***Micro-RNA*** **(DNA Short Reference Sequence)**

Case 2: ***Bacillus anthracis*** **DNA comparison**

Case 3: ***Amino Acid***

Output Options (Speedup Impact)

Detailed: -Q -H -f -10 -g -3 -d 10 -b 10 -s OpenFPGA.mat -E 0.0001
Minimal: -Q -H -f -10 -g -3 -d 0 -b 10 -s OpenFPGA.mat -E 0.0001

FPGA Performance

ORNL XD1 (Virtex2): Initial Results

Case 1: *Micro-RNA*

FPGA vs Opteron Time (hrs) for FASTA

	1	2	3	4	5
CPU 2.2GHz	75	-	-	-	-
FPGA(s) 0.2GHz	7.39	3.75	2.48	1.91	1.56
FPGA Speedup vs 1 CPU	10.15	20.0	30.2	39.3	48.1

Cray XD1 FPGA Speedup vs. 2.2 GHz Opteron
Case 2: *Bacillus anthracis DNA comparison*

Virtex2 Pro 50 Speedup

	24	*25*	*26*	*27*	*28*	*29*	*30*	*31*	*32*	*33*	*34*	*35*	*36*	*37*	*38*	*39*	*40*	*41*	*Avg*	*SD*
8k	26	30	27	32	30	30	29	30	30	27	31	29	31	30	31	31	31	30	29.6	1.2
16k	22	25	26	31	30	30	28	31	28	27	30	29	29	29	32	31	32	29	28.7	2.5
8k	49	49	49	50	49	49	50	49	49	49	49	49	49	49	50	49	49	49	49.4	0.2
16k	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	49	50	50	49.9	0.3

Virtex4 LX160 Speedup

	24	*25*	*26*	*27*	*28*	*29*	*30*	*31*	*32*	*33*	*34*	*35*	*36*	*37*	*38*	*39*	*40*	*41*	*Avg*	*SD*
8k	36	43	39	47	44	45	43	45	45	39	46	42	46	44	46	46	46	43	43.5	2.9
16k	29	33	37	45	44	43	39	47	41	37	46	41	43	41	47	46	48	43	41.5	4.9
8k	98	98	98	97	97	98	98	98	98	97	98	98	98	98	98	98	98	97	97.6	0.1
16k	100	101	101	100	100	100	101	101	101	101	101	101	101	101	100	100	101	100	100.7	0.4

24 => Sequence AE017024

XD1 Virtex2 Speedup vs. 2.2 GHz Opteron

Case 2: *Bacillus anthracis DNA comparison*

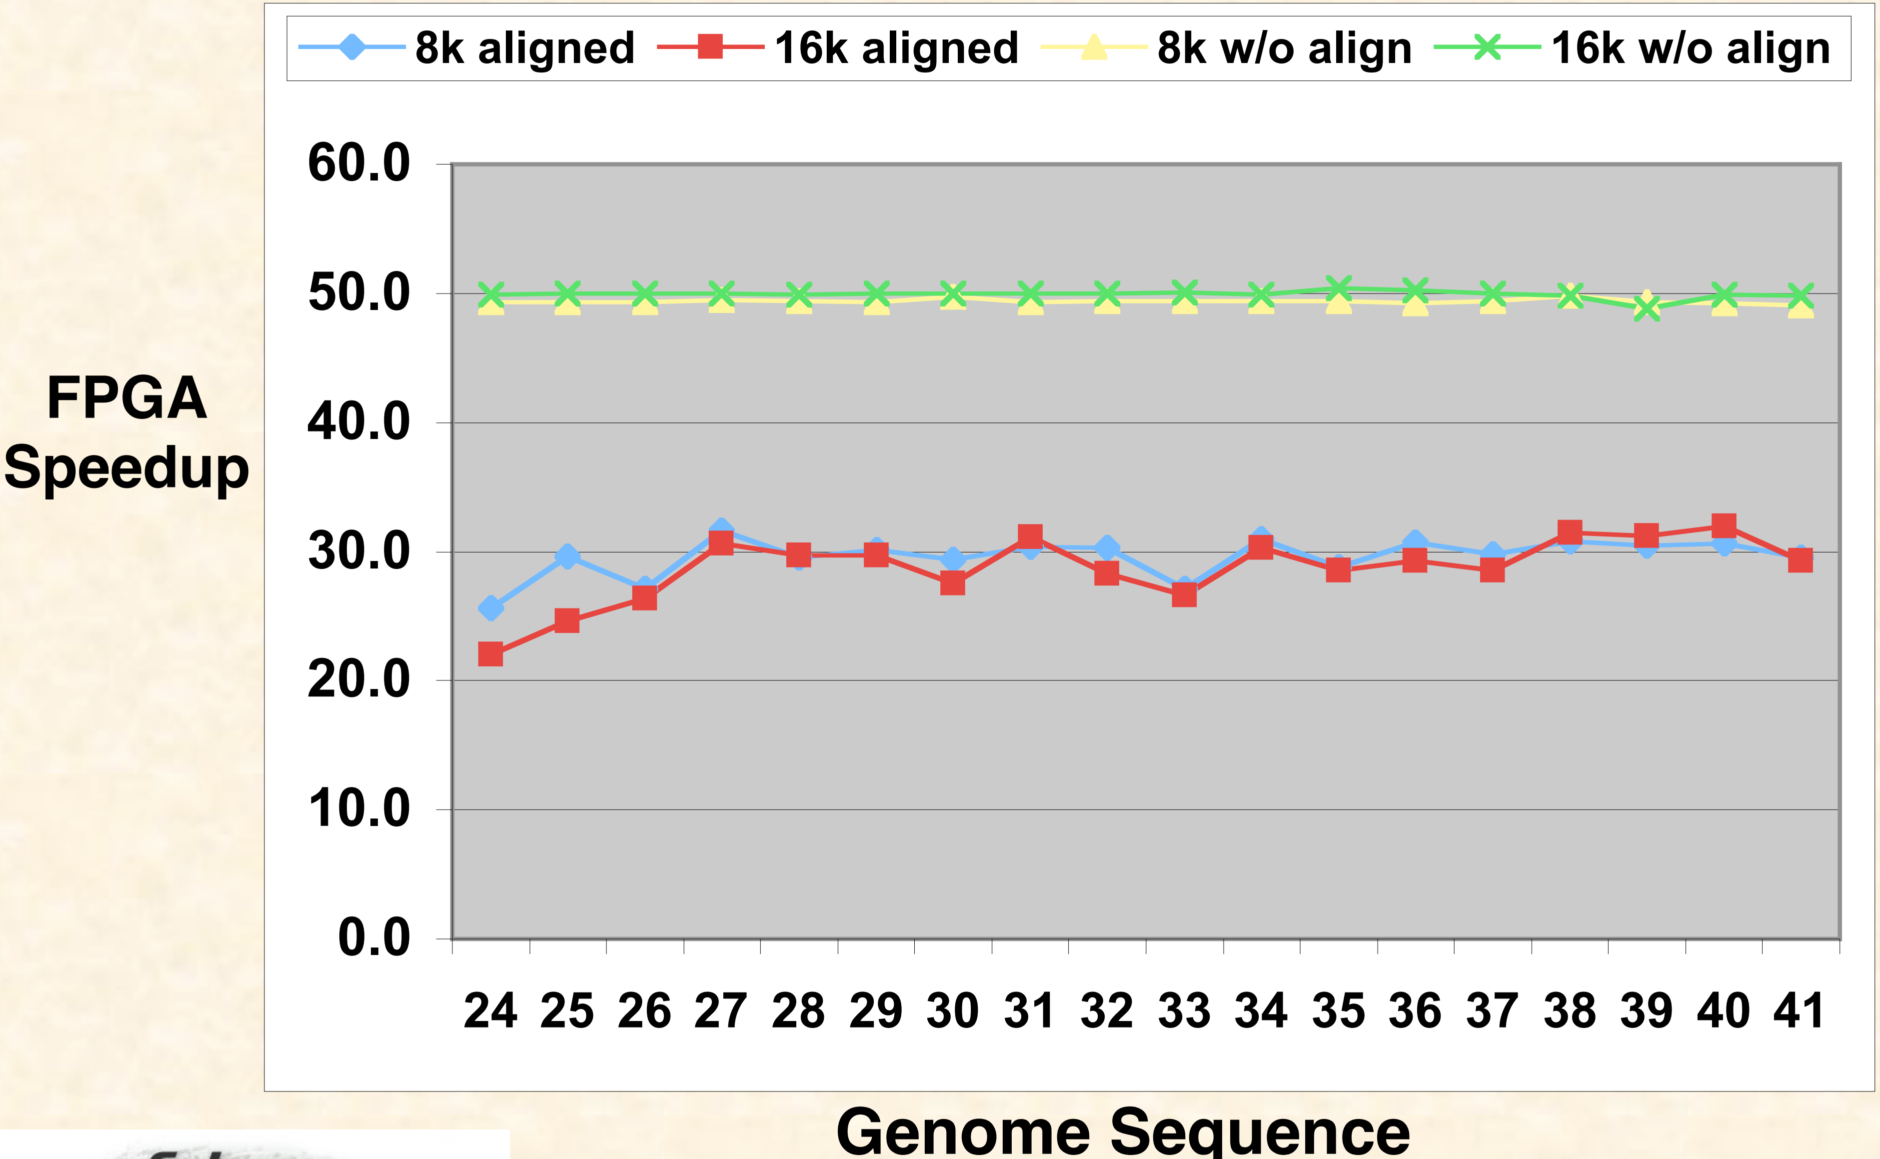

Cray XD1 Virtex-4 100X Speedup*

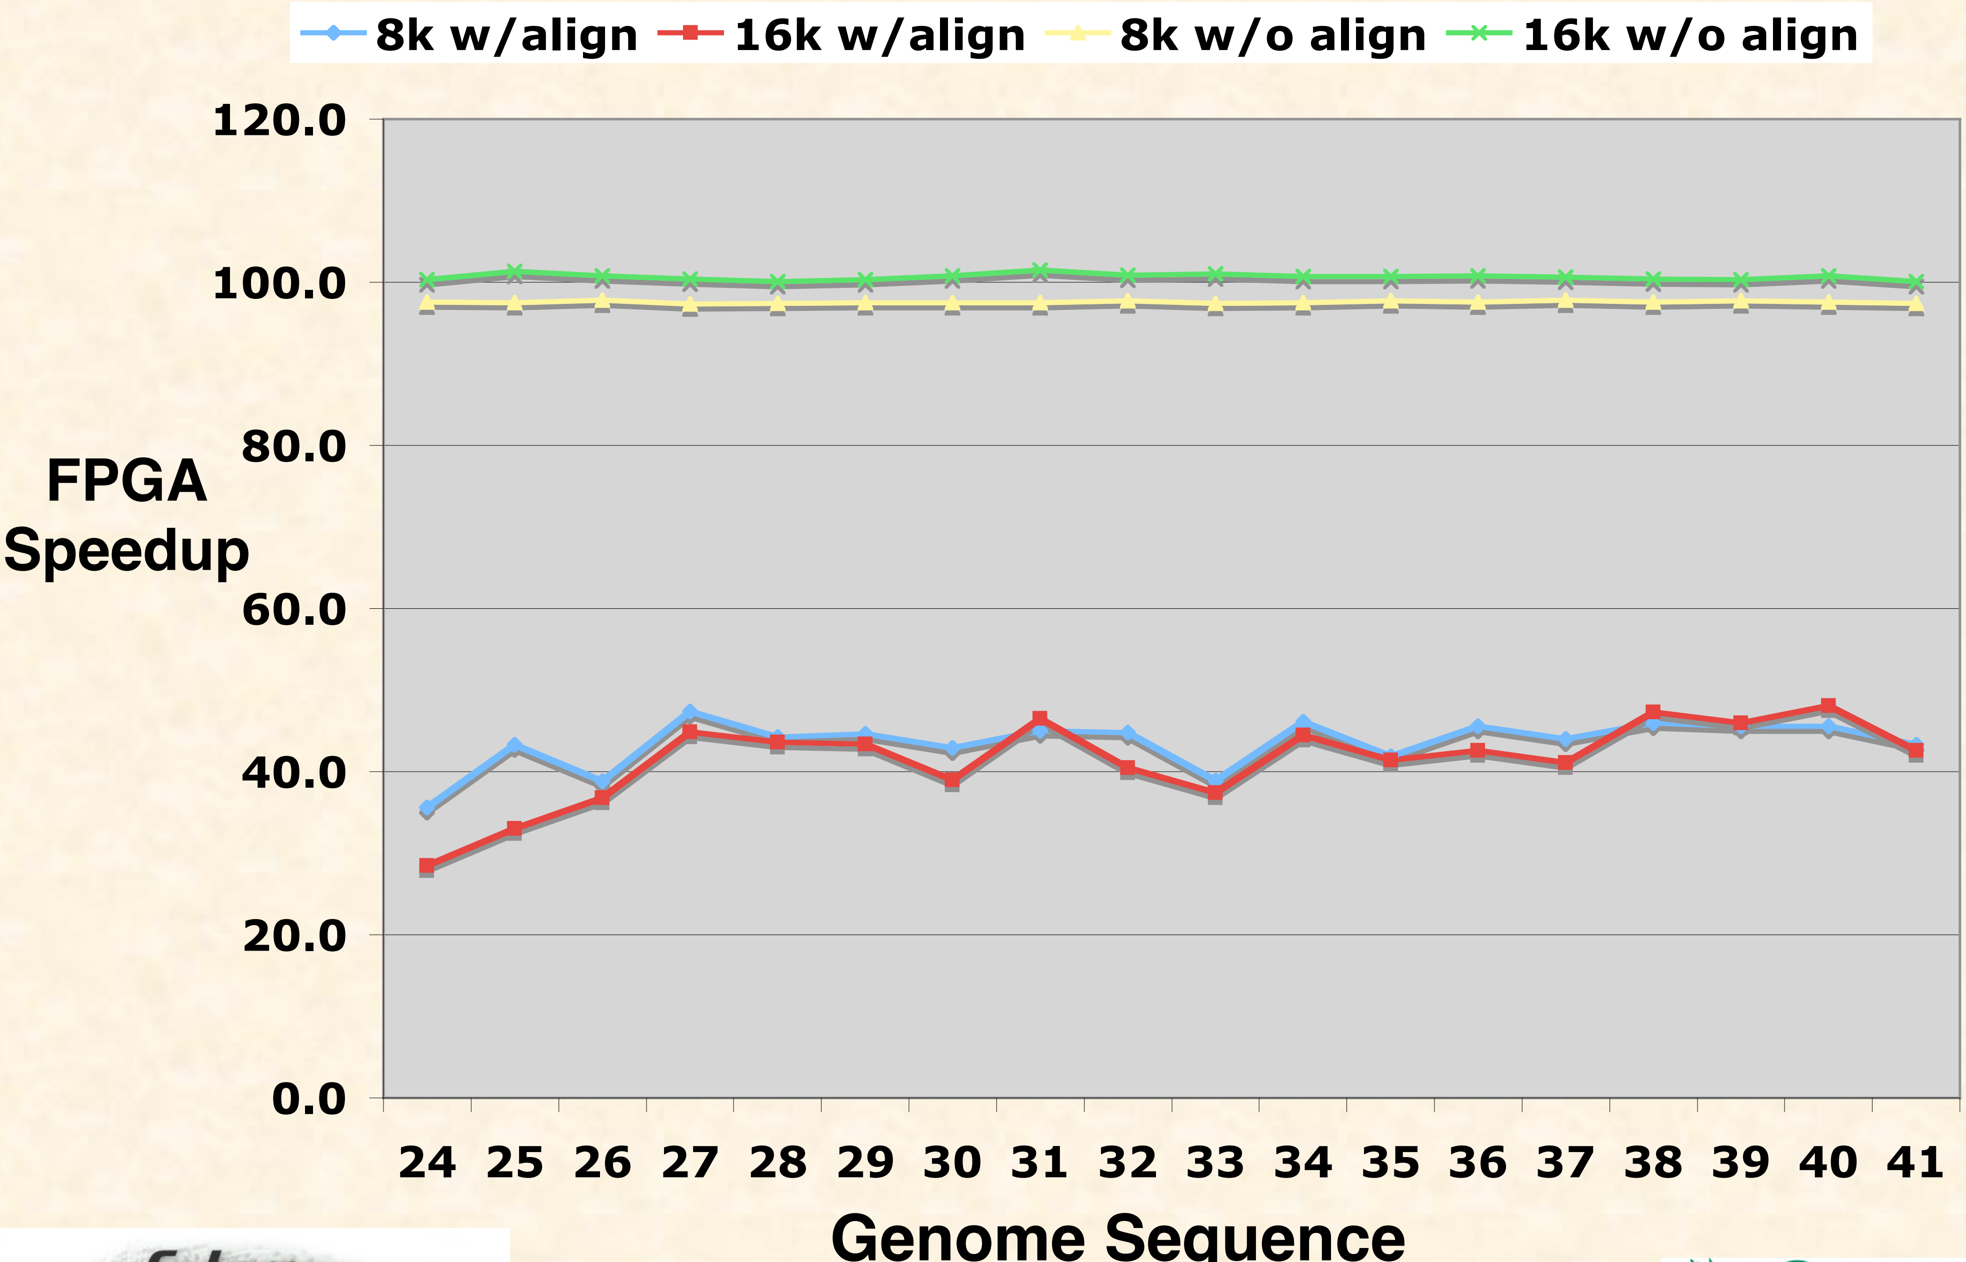

*vs. 2.2 GHz Opteron

XD1 Virtex2 Speedup vs. 2.2 GHz Opteron

Case 3: *Amino Acid*

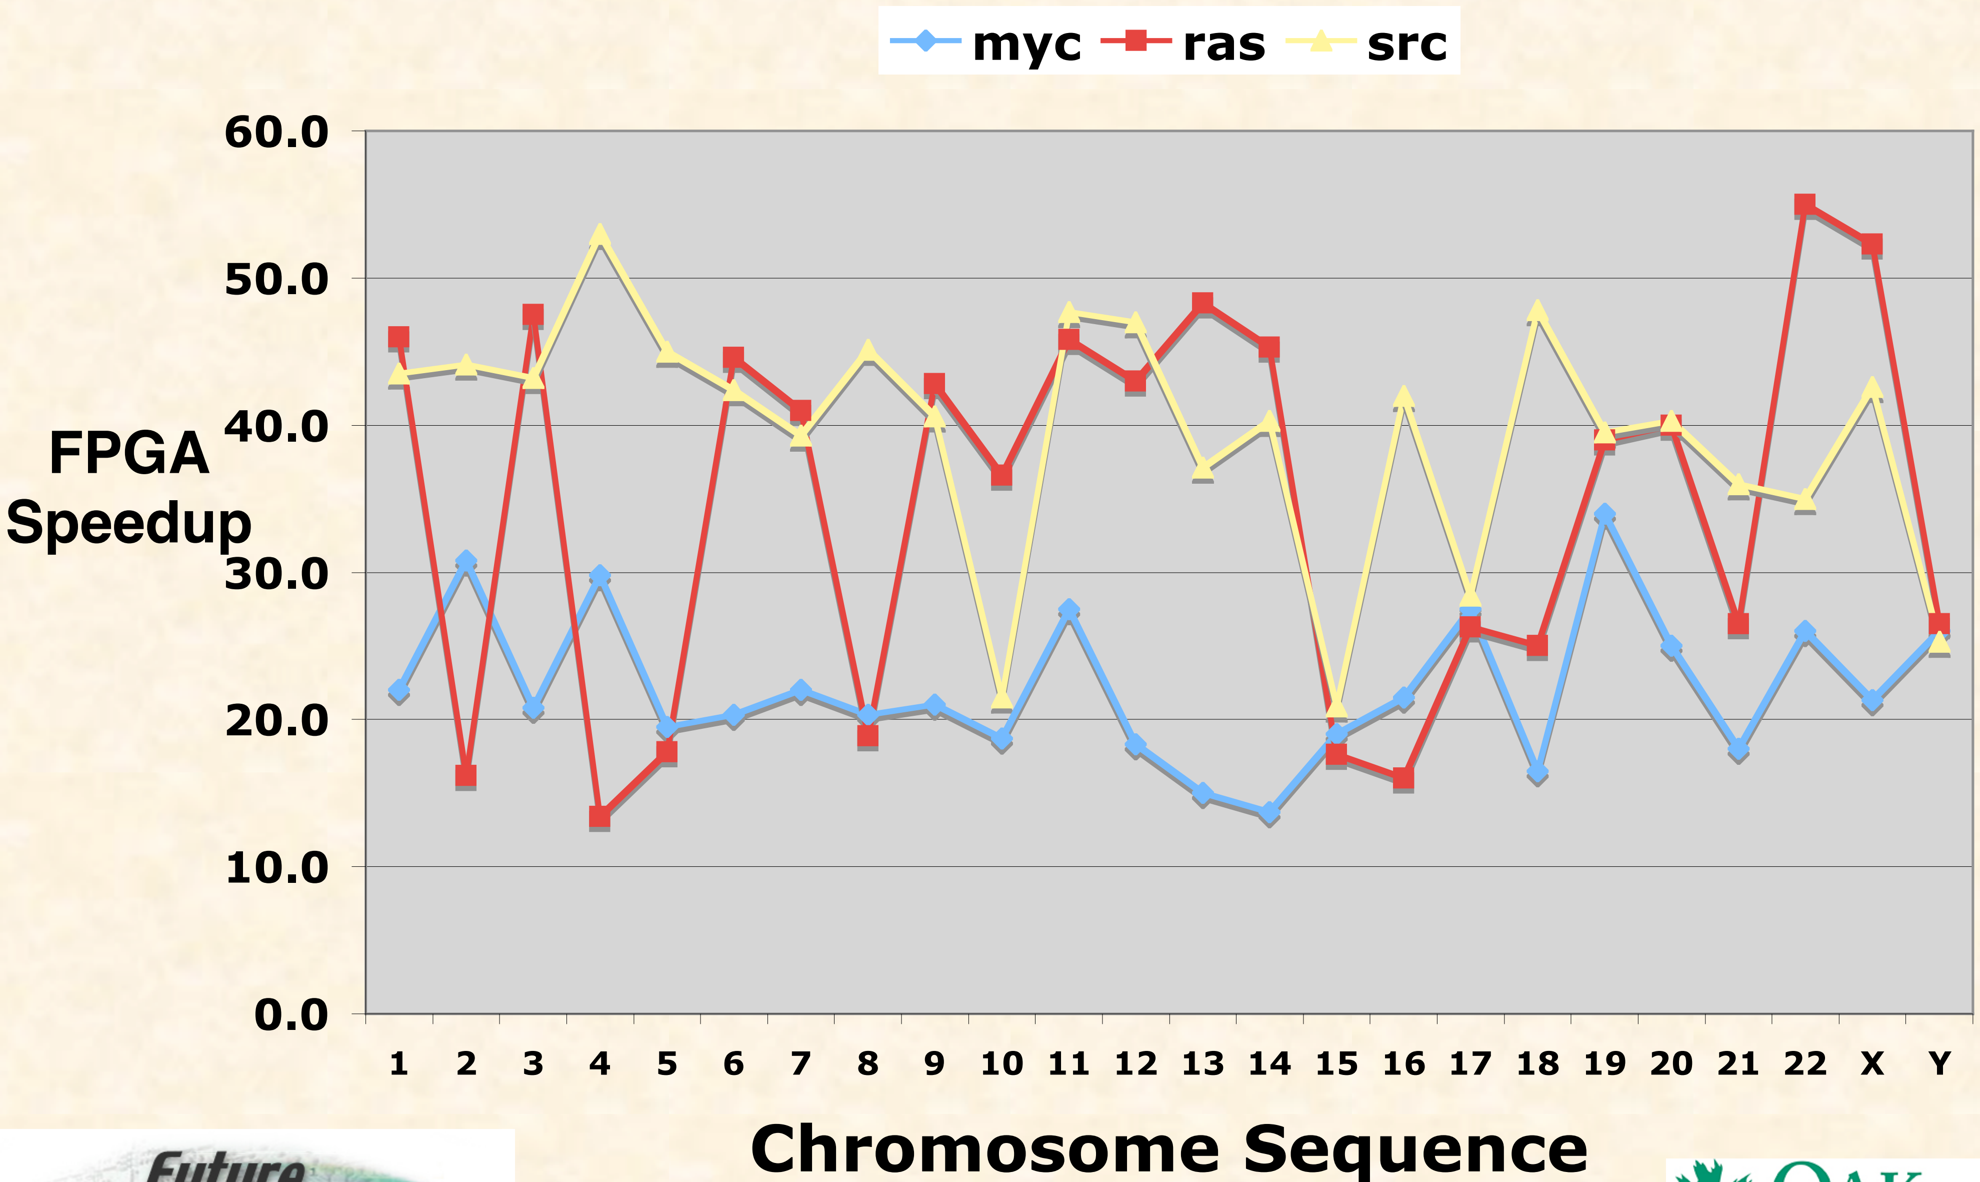

FPGA Speedup grows with Query Size

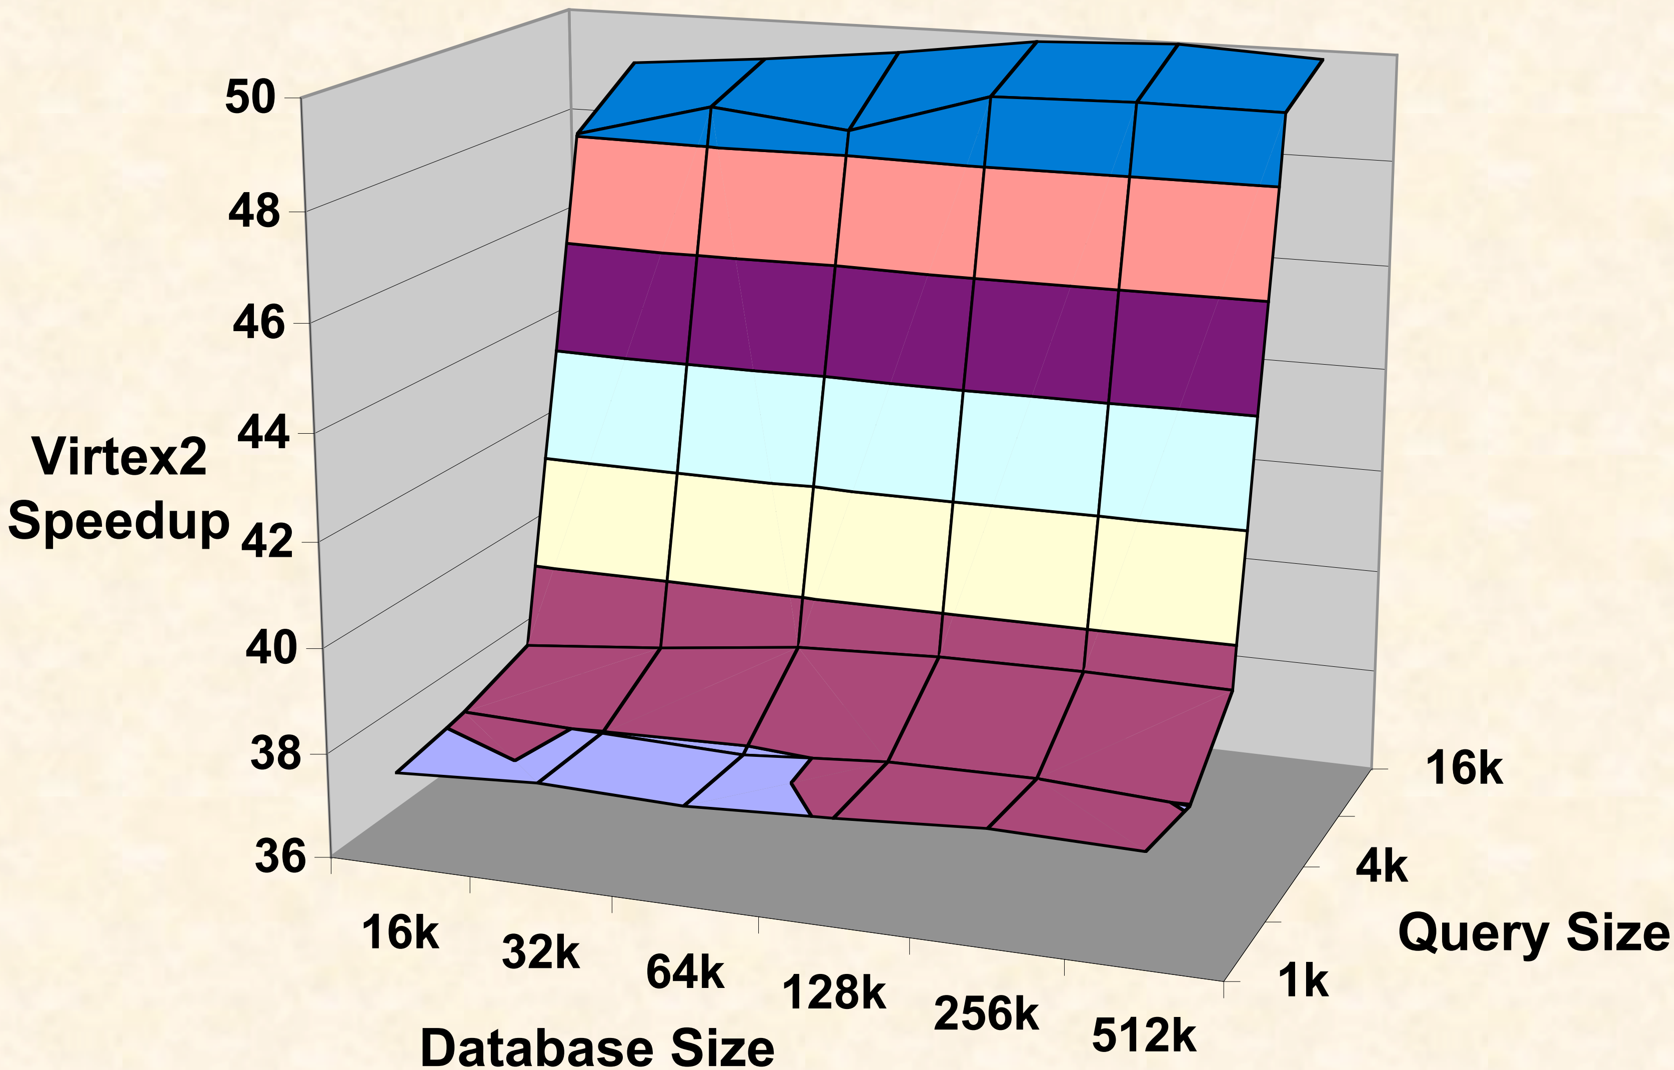

Future Opportunities

- Speedup XD1 code another 2X => 200X (LX160)
- LX200 speedup (89k/68k = 1.3) => 262X (LX200)
- 144 XD1 FPGAs => 144X50+ => 7,200X
- DRC LX200 module => Cray XT4
- ORNL-Xilinx-Cray-DRC CHiMPS collaboration
  - widen range of applications (climate, MD, solvers,...)

Future Technologies Group ornl

Oak Ridge National Laboratory

Summary

- ***FPGA, Genome matching background***

  - FASTA, search34, Smith-Waterman

- **Results: 3 openfpga.org cases**
  XD1 Speedups: 50X (Virtex2), 100X (Virtex4)
  (*promise of 200X or more*)

- ***Future: 144 FPGAs,* DRC LX200 => XT4**

THANK YOU

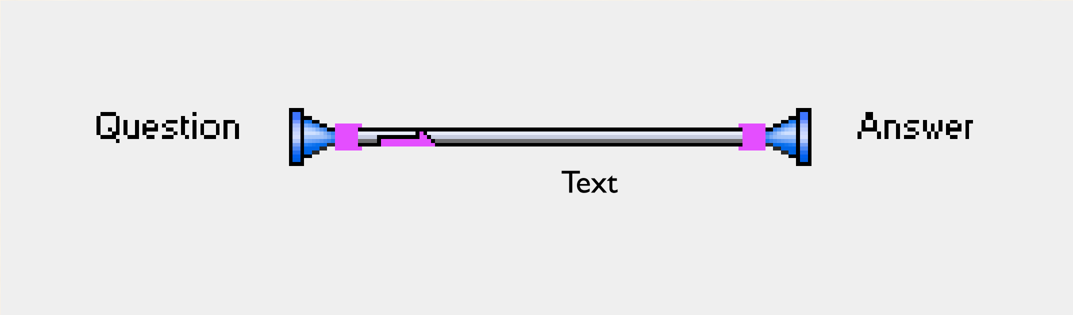

Google: Olaf ORNL