HPC Accelerator Research
100X Speedup with FPGAs*

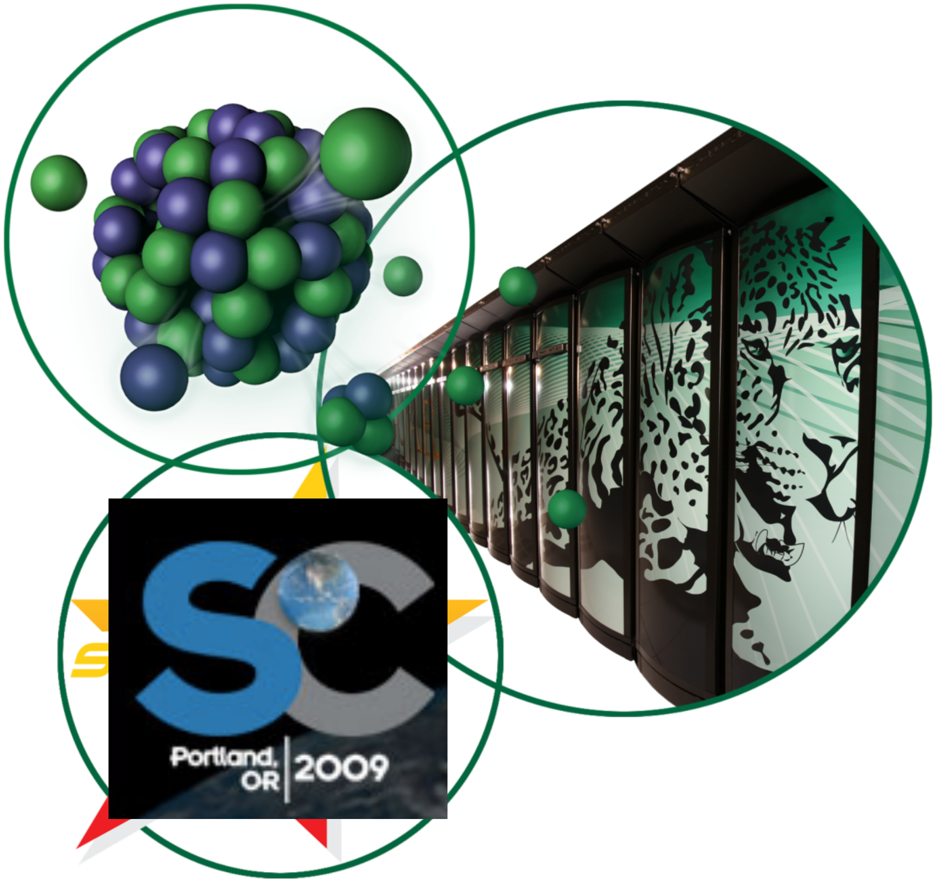

Presented by

Olaf O. Storaasli

Future Technologies Group
Computer Science and Mathematics Division

*Field-Programmable Gate Arrays

FPGA: Your “custom chip”

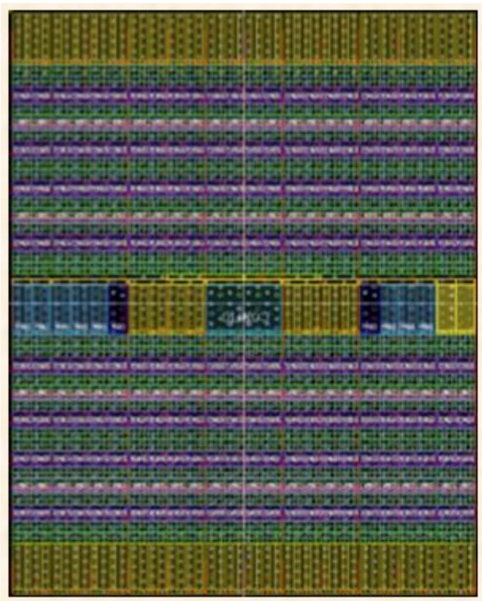

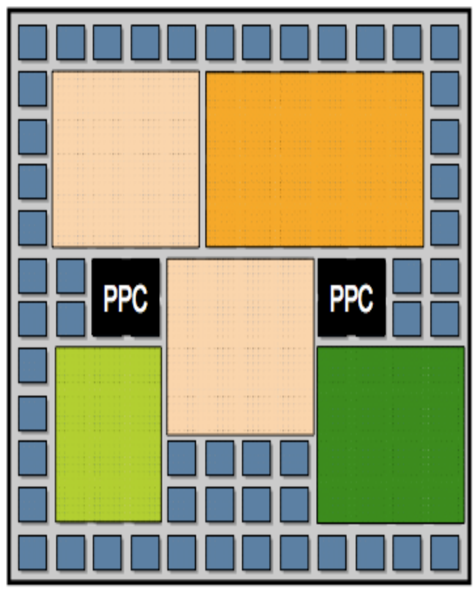

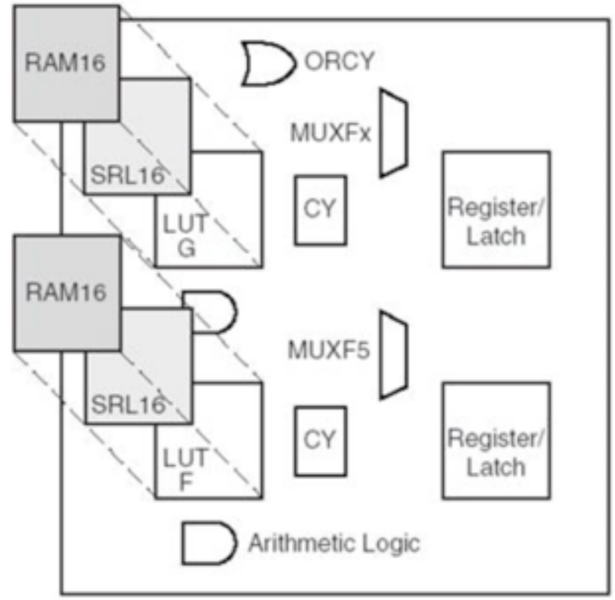

Xilinx Virtex4 FPGA: 25K Logic slices

- **Tailor Logic slices to your application**
- **On-chip RAM, multipliers and PowerPCs**
- **FastIO: Gigabit transceivers/DSP blocks**
- **100–1000 operations/clock cycle**

FPGA Logic slice
(MiniCPU)

Why FPGA accelerators?

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

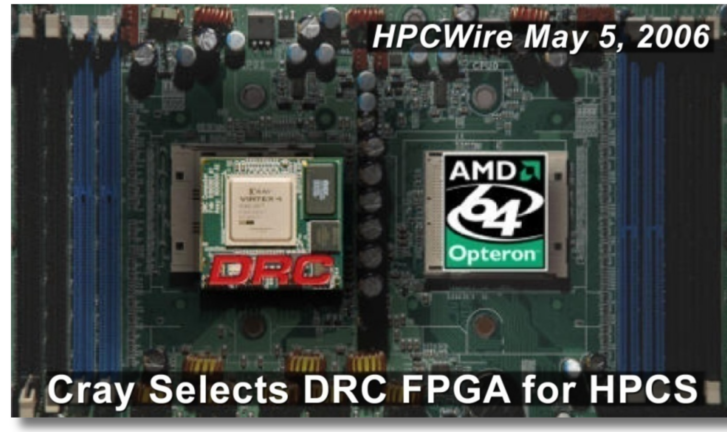

Cray XT5 FPGA accelerator

Porting climate code to FPGAs

ORNL-Xilinx Collaboration

Profile

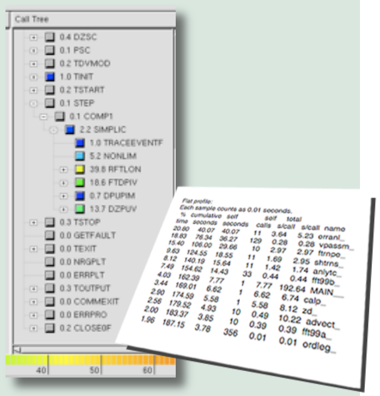

Find parallelism: 80% FFTs

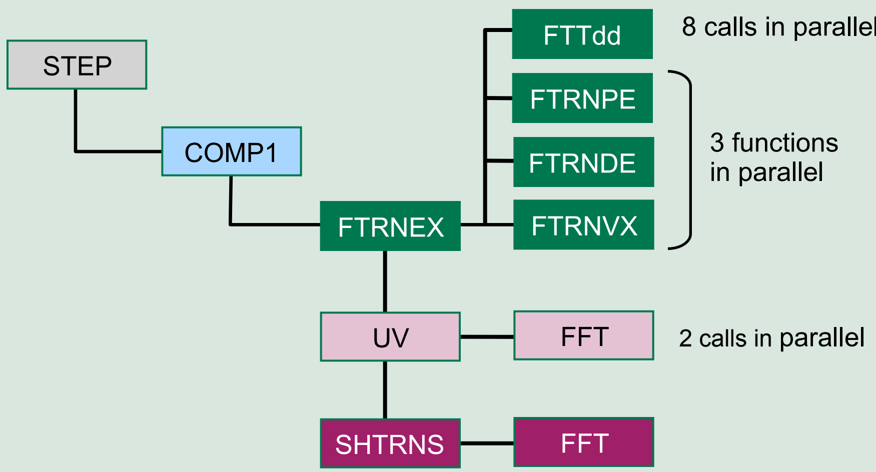

Goal

More GF/$ GF/Watt

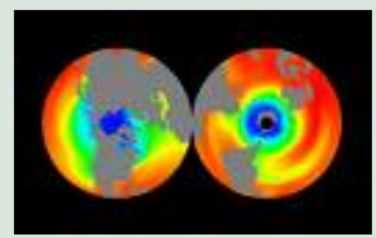

Model faster

FPGA coding options

Gauss matrix solver

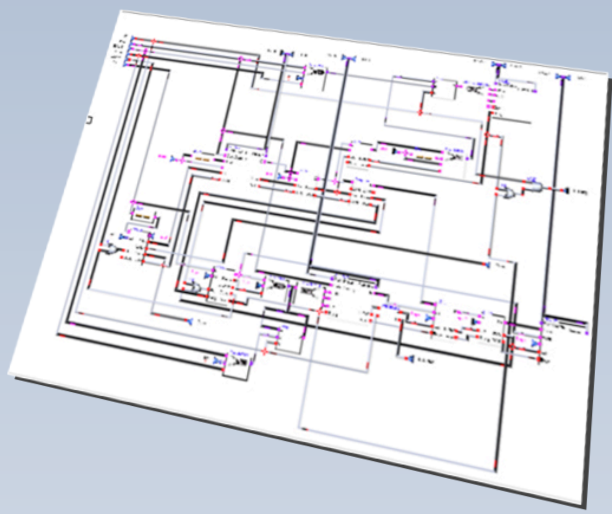

Graphical: 3D via icons (Viva)

Compile-simulate-debug

Text: 1D flow (Mitrion C)

Others: Carte, CHiMPS-VHDL, DSPlogic

37×* LU Matrix Factor Speedup
10× Matrix solver Speedup

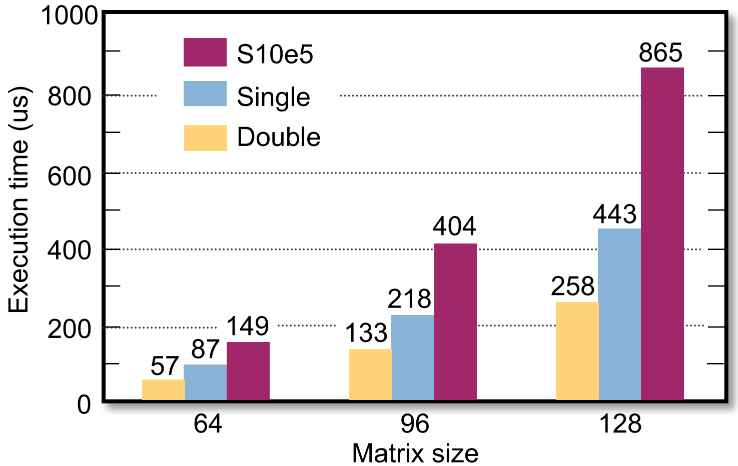

Design	Double FP	Single FP	S10e5
PE amount	8	16	32
Max size	128	256	256
Achievable frequency	120 MHz	150 MHz	150 MHz
Slices	27,005 (57%)	14,792 (59%)	14,730 (62%)
BRAMs	68 (29%)	129 (55%)	65 (28%)
MULT18X18	128 (55%)	64 (27%)	32 (13%)

Benefits:

High-performance: SP arithmetic

High-precision: DP accuracy refine

Speedup grows with matrix size

as LU dominates calculations

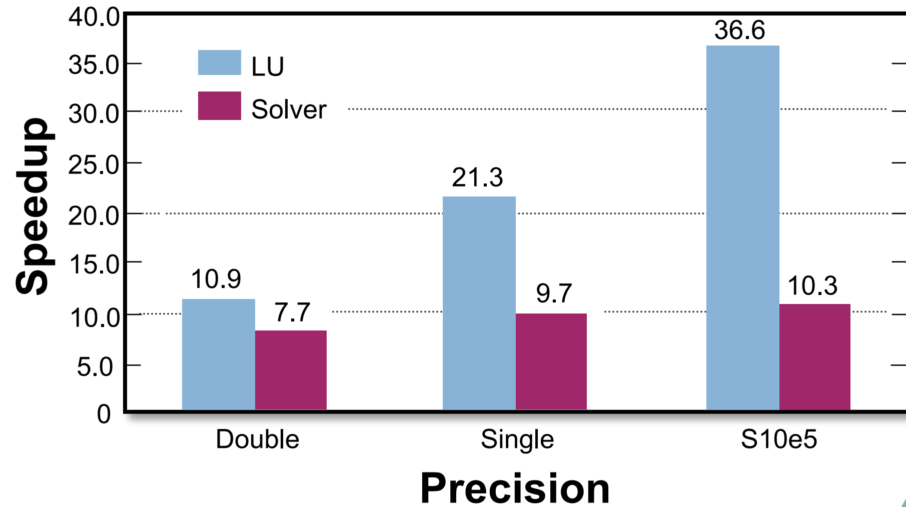

1st mixed-precision LU and solver for FPGAs

*2.2 GHz Opteron

100× speedup*: human DNA sequencing

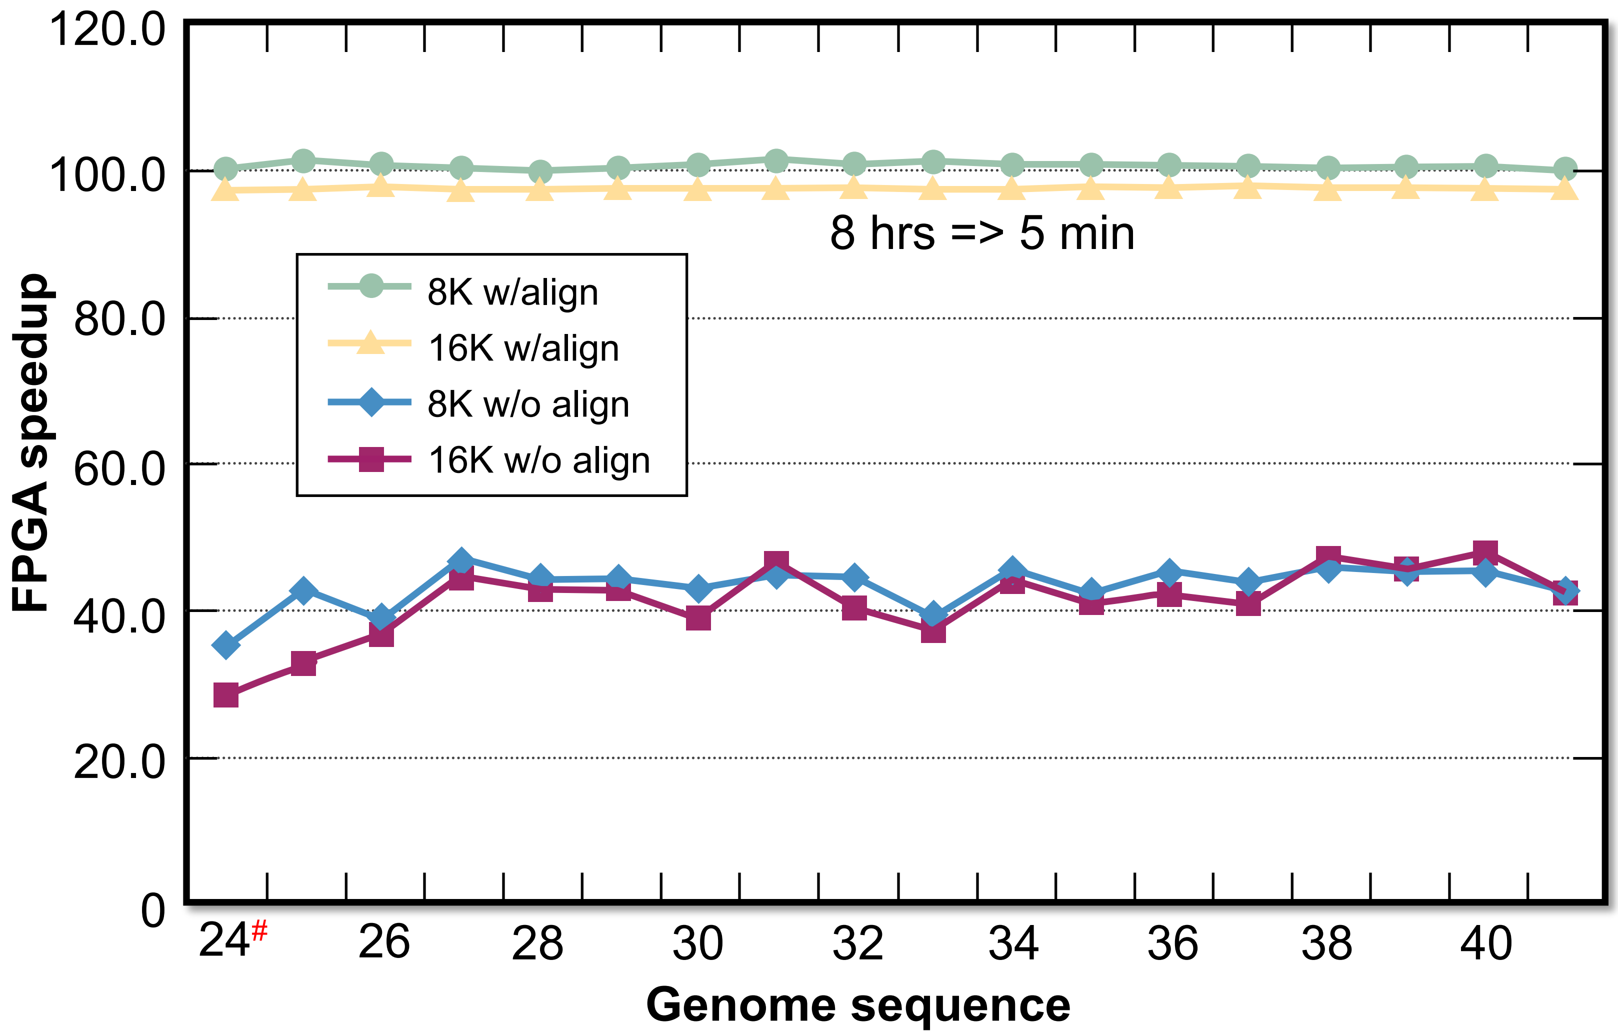

*Virtex-4 FPGA vs 2.2 GHz Opteron on Cray XD1

24= Sequence AE17024

Faster DNA sequencing* using 150 FPGAs

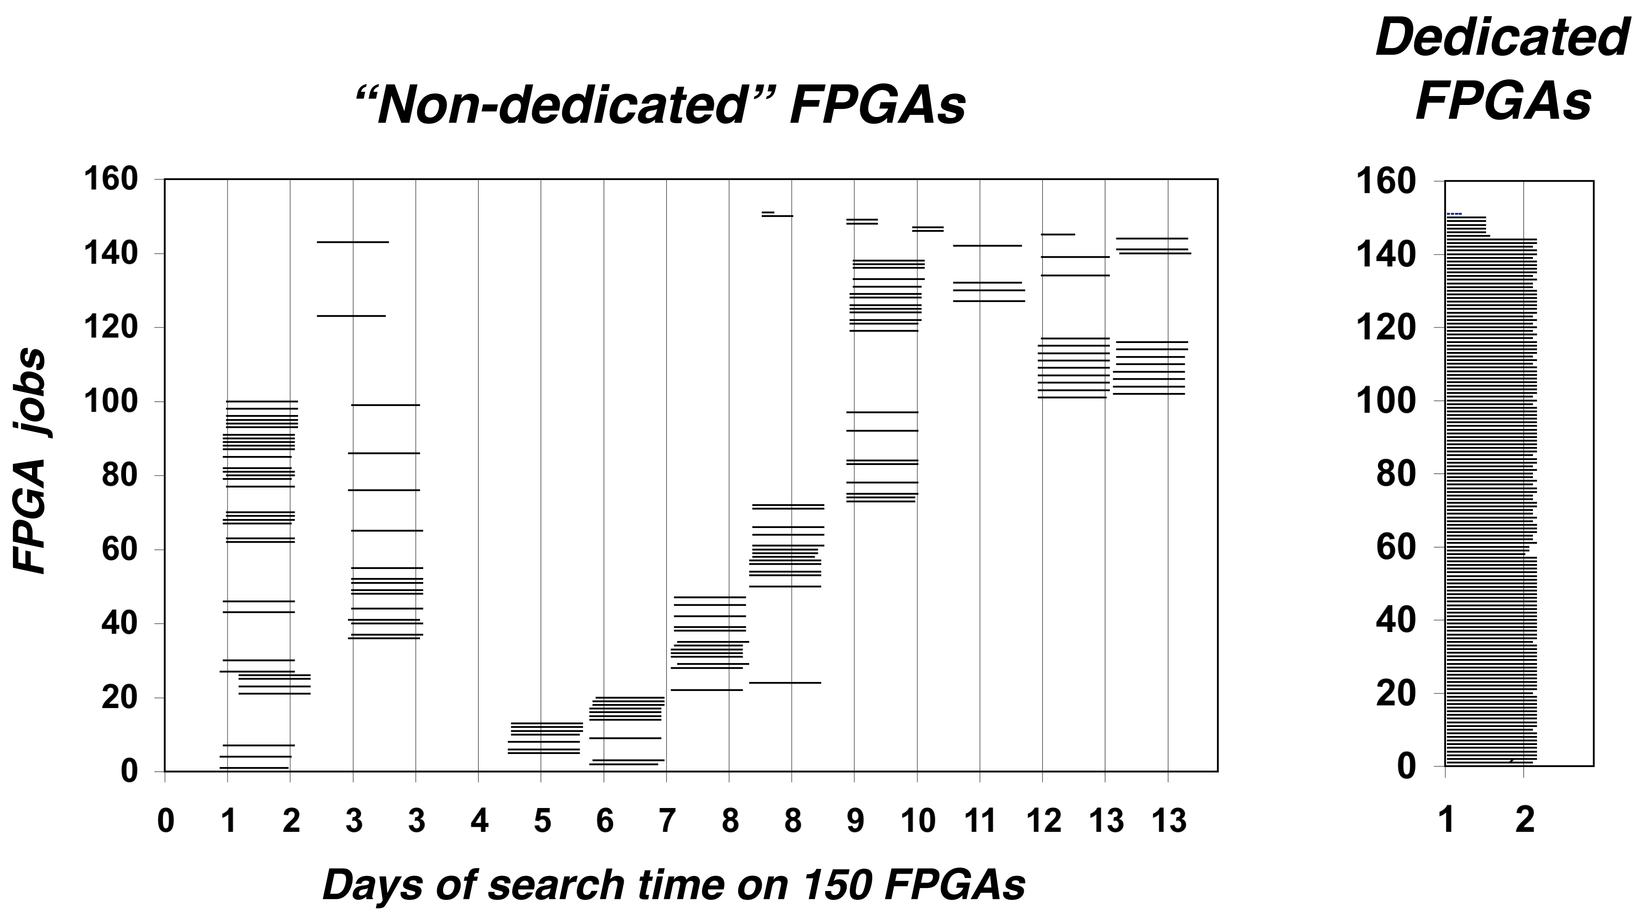

*Human-mouse DNA compare (FASTA)

DNA Sequence speed*on 150 FPGAs

* *State-of-the-art:* **Giga Cell Updates Per Second (GCUPS)**

- **DNA characters: Human = 155 million, mouse = 165 million**

Total compares = $155M \times 165M \times 1062 \times 2 = 51 \times 10^{15}$ **cell updates**

- **Sequential FPGAs take 11,923,200 s** (138 days)
  ==> $51 \times 10^{15}/11{,}923{,}200$ **= 4.3 TCUPS (*Tera CUPS*)**
- **Parallel (actual) = 1,114,560 s** (12.9 days) **==> 46 TCUPS**
- **Parallel (dedicated) = 86,400 s** (1 day) **==> 605 TCUPS**

Summary

Speedup* on 1 FPGA:

10X - *general matrix equation solution*
100X - *DNA sequencing*

Speedup on 150 FPGAs - *DNA Sequencing*

1 Opteron ==> 20 years   150 Opterons ==> 6 weeks
1 FPGAv2 ==> 5 months   150 FPGAs ==> 1 day ***49X speedup***

==> 7,350X speedup over one Opteron (VirtexIIs)
==> 14,700X speedup (Virtex4s)

More petaflops at reduced power

*Compared with one 2.2 GHz Opteron

Contact

Olaf Storaasli

Future Technologies Group
Computer Science and Mathematics Division
Olaf@ornl.gov
Google **Olaf ORNL**

Acknowledgment:
Thanks are extended to the Naval Research Lab
for use of its Cray XD1 with 150 FPGAs