# Novo-G#: A Community Resource for Exploring Large Scale Reconfigurable Computing with Direct and Programmable Interconnects

Martin Herbordt Alan George Herman Lam

BOSTON UNIVERSITY

# Motivation

Problem: Applications are increasingly communication bound

→ one motivating example: MD is limited by global communication

High-level solutions

-- reduce latency

-- compute in the network

Explore compute/communicate heterogeneity

Use FPGAs – Combine MGTs w/ computation

# Novo-G Reconfigurable Supercomputer

- Developed and deployed at CHREC
  - Most powerful reconfigurable computer in research community (2009-present)
    - **448 (soon 512) high-end Altera FPGAs with 3.5TB (soon 4.5TB) of FPGA-attached SDRAM**
  - 2012 Alexander Schwarzkopf Prize for Technology Innovation @ NSF
- App acceleration
  - Computer vision, finance, bioinformatics, molecular dynamics, crypto, et al.
- Hardware emulation
  - Behavioral emulation of future apps and systems, up to Exascale
- Latest upgrade
  - 64 Altera Stratix-V D8 FPGAs
    - On Gidel ProcE-V PCIe boards
    - Additional 64 in development
  - 3D torus interconnect (4x4x4)
    - 6 links per Stratix-V (40 Gb/s per link)

**2009:** 96 Altera Stratix-III E260 FPGAs, each with 4.25GB SDRAM, on Gidel ProcStar-III cards
**2010:** 96 Altera Stratix-III E260 FPGAs, each with 4.25GB SDRAM, on Gidel ProcStar-III cards
**2011:** 96 Altera Stratix-IV E530 FPGAs, each with 8.50GB SDRAM, on Gidel ProcStar-IV cards
**2012:** 96 Altera Stratix-IV E530 FPGAs, each with 8.50GB SDRAM, on Gidel ProcStar-IV cards
**2014:** 32 Altera Stratix-V D8 FPGAs, each with 16GB SDRAM and 3D torus i/f, on Gidel ProcE-V card
**2015:** 32 Altera Stratix-V D8 FPGAs, each with 16GB SDRAM and 3D torus i/f, on Gidel ProcE-V card
**2016:** 64 Altera Stratix-V D8 FPGAs (coming soon)

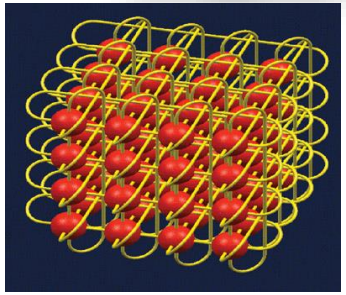


BU Department of Electrical & Computer Engineering

Made possible by support from UF, Altera, Gidel, NSF, and DOE

# Novo-G#

- NSF CRI funded project with major donations from Altera and Gidel
- An enhancement to Novo-G (appended to Novo-G)
- BU / U. Florida collaboration
- Assembled in Florida, prototype & router at BU
- We are developing public infrastructure and you are invited to use it and/or collaborate. See Novo-G# Forum for more info
- Specifications
  - 128 nodes, each with 1 top-of-the-line Altera FPGA
  - Boards by Gidel
  - 7 MGTs (40Gb/s ports) for direct connection to other FPGAs
    - Augments Novo-G with low-latency, high-throughput, and reconfigurable communications capabilities among FPGA accelerators

# Sample Results – Distributed 3D FFT

| fft size | Fixed(us) | Variable(us) | Reference(us) |
|---|---|---|---|
| $16^3$ | 1.63 | 1.64 | 3.98 |
| $32^3$ | 2.24 | 2.48 | 5.46 |
| $64^3$ | 6.72 | 6.92 | 16.76 |
| $128^3$ | 35.13 | 41.14 | 101.11 |

# Novo-G# Forum

BOSTON UNIVERSITY

UF UNIVERSITY of FLORIDA

UCLA

Imperial College London

THE UNIVERSITY of TENNESSEE KNOXVILLE

Northeastern UNIVERSITY

UNIVERSITY OF SOUTH CAROLINA

CLEMSON UNIVERSITY

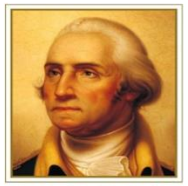

THE GEORGE WASHINGTON UNIVERSITY WASHINGTON DC

GiDEL

Washington University in St. Louis

Federal University of Pernambuco (Brazil) UFPE

UNIVERSITY of GLASGOW

ALTERA

BU Department of Electrical & Computer Engineering

# Reconfigurable Supercomputing BoF
# Tuesday 5:30 – 7:00

## Introduction (15 minutes)

- Overview of Reconfigurable Computing – Alan George
- Featured Talk on Catapult @ Microsoft – Derek Chiou

## Highlight Briefings (5 minutes each)

- Altera – Mike Strickland
- Boston University – Martin Herbordt
- CHREC/Florida – Herman Lam
- DRC – Roy Graham
- Gidel – Reuven Weintraub
- IBM – Jonathan DeMent or Mike Paolini
- Micron – Steve Palowski
- Nallatech – Allan Cantle
- Xilinx – Michael Leventhal

## Q&A Panel (all speakers above) with Audience (30 min)