



HPRC Applications R&D

@ CHREC-UF




BOF ON HPRC
@SC12

UF UNIVERSITY of
FLORIDA

Virginia
Tech
1872
VIRGINIA POLYTECHNIC INSTITUTE
AND STATE UNIVERSITY

BYU
BRIGHAM YOUNG
UNIVERSITY


THE GEORGE
WASHINGTON
UNIVERSITY
WASHINGTON DC

Herman Lam

Associate Director, CHREC
Associate Prof. of ECE
University of Florida

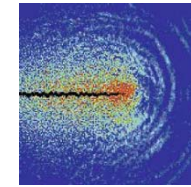
HPRC Applications

Opportune Convergence of:

- ❑ **Needs:** Escalating demands of HPC domains
 - Beyond conventional HPC w.r.t. **performance** and **energy sustainability**
- ❑ Emergence of:
 - **High-performance RC devices** and **productivity tools**
- ❑ Expanding availability of **HPRC systems**
- ❑ Promising R&D and commercial **HPRC apps**

Boston University

- ◆ Molecular dynamics



Convey

- ◆ Bioinformatics
- ◆ Graph 500 (big data analytics)



IBM (Nallatech)

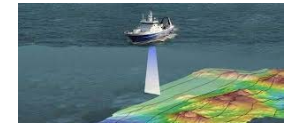
- ◆ Altera OpenCL – Barrier options pricing

Maxeler

- ◆ J. P. Morgan – Financial analytics
- ◆ Oil and Gas

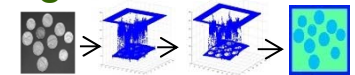
Novo-G (GiDEL)

- ◆ Monsanto – BLAST Toolset
- ◆ Veritomyx – Isotope Pattern Calculator
- ◆ UBS – Multi-asset barrier options pricing
- ◆ Scalable arch for image segmentation



Pico Computing

- ◆ Burrows-Wheeler Aligner (BWA)



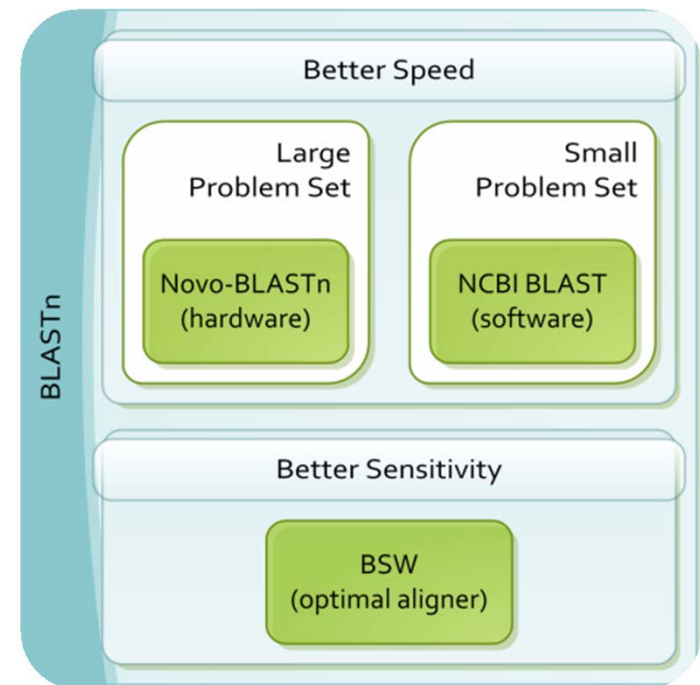
BioRC: CHREC BLAST Toolset

BLAST: Basic Local Alignment Search Tool

- ❑ Like Smith-Waterman (SW)
 - Computes local alignment of two biological sequences
- ❑ Unlike Smith-Waterman
 - **Much faster**, based on heuristics
 - Generates **non-optimal** alignments
- ❑ Currently the **most commonly used** general sequence alignment tool in bioinformatics

CHREC BLAST Toolset

- ❑ **Novo-BLAST**
 - Better speed*
 - Same Sensitivity*
- ❑ **BLAST-wrapped SW (BSW)**
 - Better Sensitivity*
 - Same Speed*

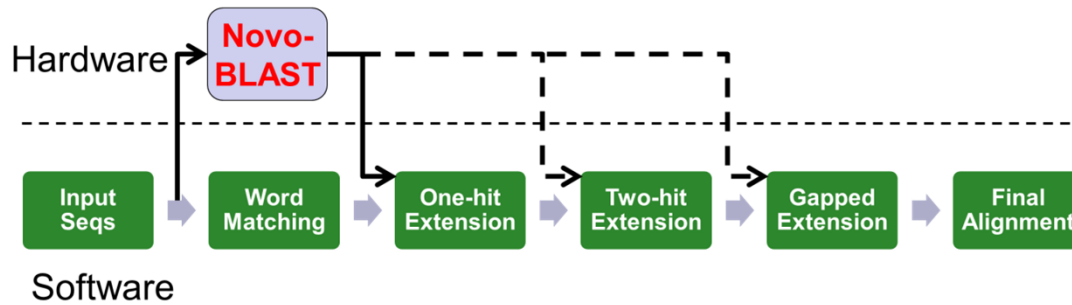


BioRC: Novo-BLAST

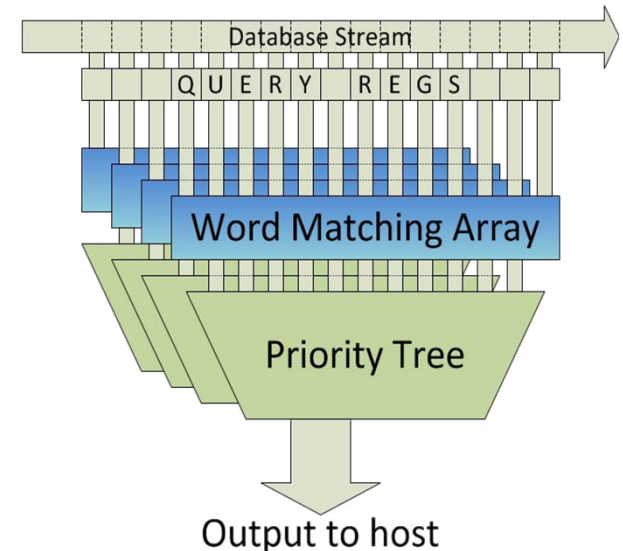
Goal: Accelerate NCBI BLAST with 100% compatible results

Approach:

- Host software integrated with BLAST
 - Processes queries/extensions in parallel
- Parts of BLAST performed in hardware
 - Word matching & parts of ungapped extension



Result: Measured speedup* of up to 265 v/s pre-compiled blastn from NCBI-BLAST+ toolset (v2.2.24)



Reference: Glycine max, chr 1
Query Set: 5000 from Gmax_cds
Software Version: NCBI BLAST 2.2.24+
Software Baseline: Intel Xeon E5520
Hardware Baseline: 1 Altera Stratix III E260
Config (Novo-BLAST): 4000-length query engine

BLAST Word Size	NCBI BLAST Runtime	Novo-BLAST Speedup
11 default	125	1.81
10	252	3.65
9	585	8.42
8	1358	19.13
7	2246	16.04
6	8782	60.99
5	41384	265.28

BioRC: BLAST-Wrapped SW (BSW)

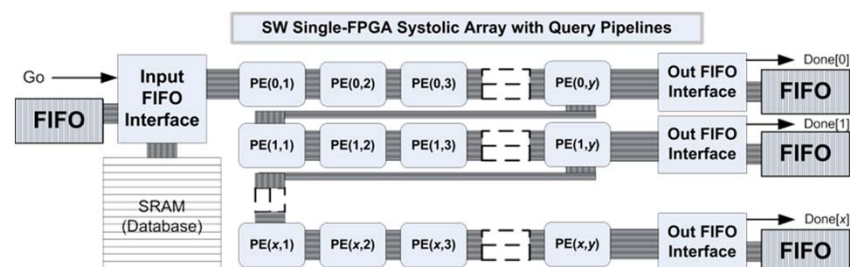
Goal: Achieve *Optimal* sensitivity with comparable speed to **NCBI BLAST** & superior speed to **FASTA SSEARCH**

Approach:

- Smith-Waterman hardware core with software alignment traceback
 - Aligns multiple queries concurrently
- Interface mimics BLAST for ease of deployment and higher familiarity

Result:

- Speedup vs. NCBI BLAST: .28 (word size 11); 1.03 (performance parity at word size 9); 244 (word size 5)
- Speedup of up to 141 vs. FASTA SSEARCH



Reference: Glycine max, chromosome 1

Query Set: 32-5000 of length 240 from Gmax_cds

Software Version: SSEARCH 36.3.5a

Software Baseline: Intel Xeon E5620 (one core)

Hardware Baseline: Altera **Stratix IV E530** (95%)

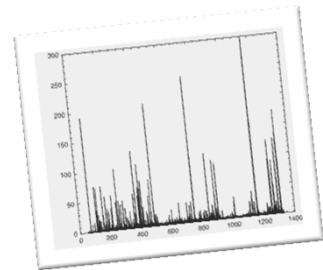
Config (BSW): 7x2 pipelines of 240 PEs, 140 MHz

Number of Queries	SSEARCH Runtime	BSW Speedup
32	00:04:51	58.1
64	00:09:16	84.1
128	00:18:00	101.8
256	00:35:29	119.7
512	01:10:25	131.3
5000	11:22:33	141.9

BioRC: Isotope Pattern Calculator (IPC)

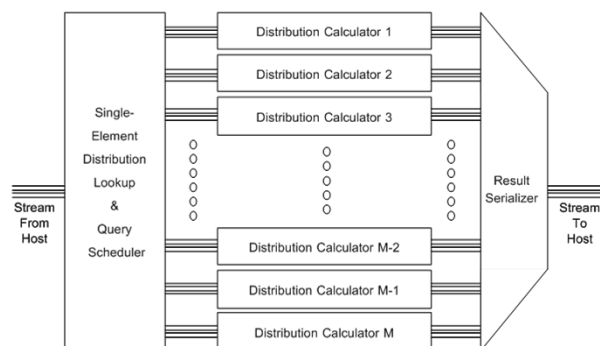
Goal: *Increase accuracy* of Protein Identification Algorithms (PIAs)

- ❑ Essential for pharmaceutical research & cancer diagnostics
- ❑ Existing PIAs has potential to revolutionize accuracy
 - Prohibitive execution times
 - Must accelerate for feasible use



Approach:

- ❑ Accelerate *Isotope Pattern Calculator (IPC)*, a dominant subroutine common in de novo PIAs
- ❑ Provide *customizable design* for general use
- ❑ Reconfigurable computing at scale to achieve *sustainable supercomputing performance*

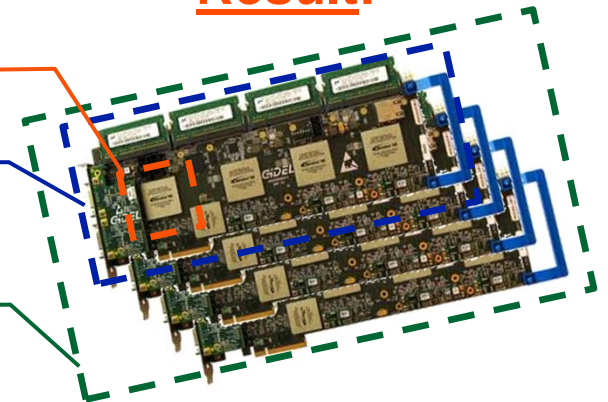


Between 72 and 566 speedup[†] on a single FPGA

Up to 1259 speedup[†] on a single board (4 FPGAs)

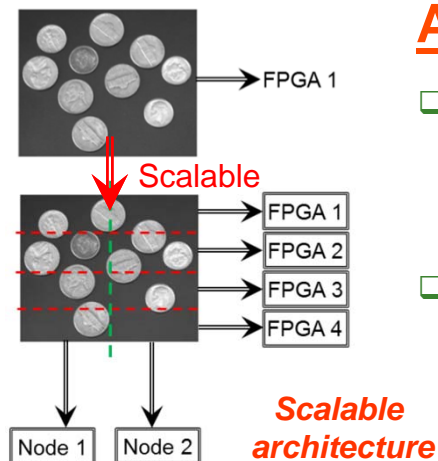
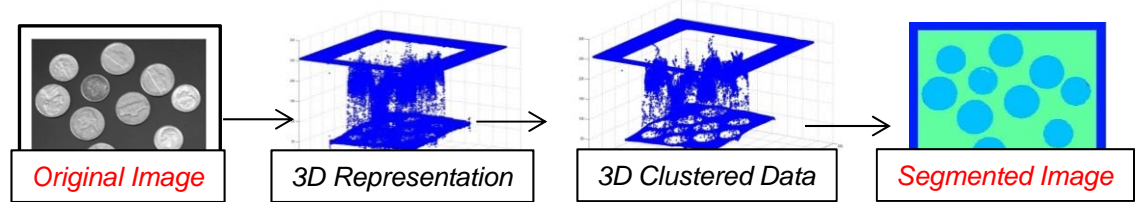
Up to 3340 speedup[†] on a single node (16 FPGAs)

Result:



DspRC: Unsupervised Image Segmentation

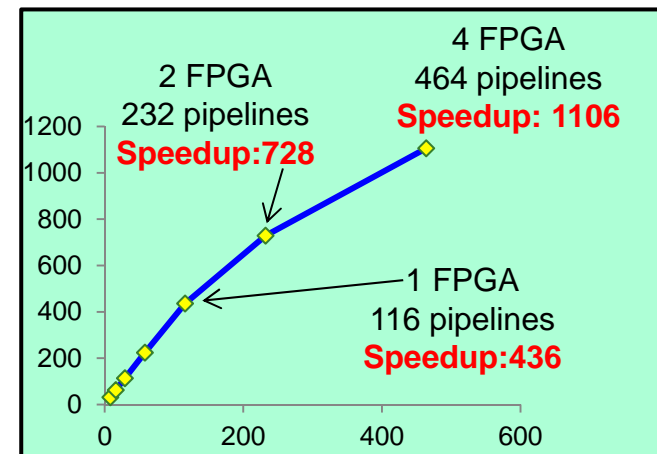
Goal: Achieve **real-time segmentation** for HD (1920x1080p) images



Approach:

- Dedicate one pipeline per pixel:
 - Pixels cluster independently
 - Replicate pipelines to efficiently utilize FPGA resources
- Scalable architecture
 - Divide input image among multiple boards

Result: Measured **speedup of 1106*** on PROCStar III board (4 Stratix III E260 FPGAs @150 MHz)



FinRC: Multi-asset Barrier Options Simulation

Goal: Use of reconfigurable supercomputing to enable **business-relevant** financial apps

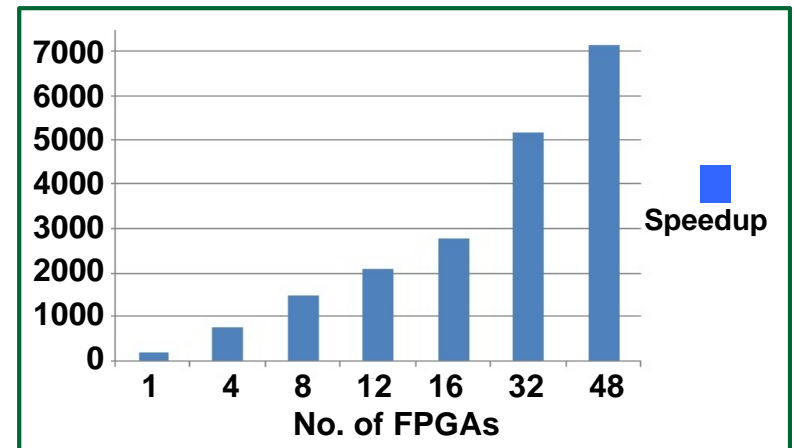
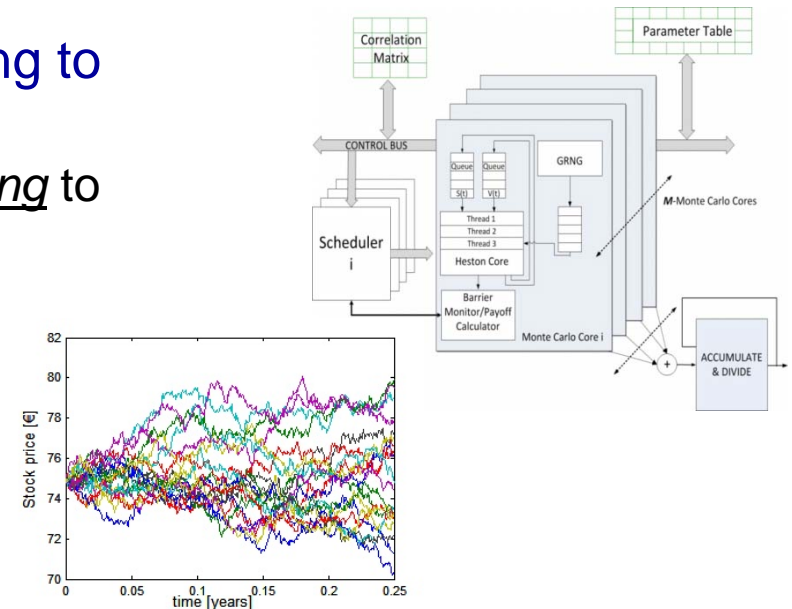
- Accelerate multi-asset barrier options pricing to meet **rigorous time constraint**

Approach:

- Simulate multi-asset barrier options under **Heston volatility dynamics** with a Monte Carlo (MC) process
- Architecture consists of:
 - **Parallel MC cores**, each capable of simulating multiple MC paths
 - **Customizable payoff kernel**, flexibility to price different types of contracts

Result: Measured **speedup of 350** on one FPGA and **7134** on 48 FPGA

- w.r.t. SSE2 optimized code using one E5-2687 core



RC Middleware

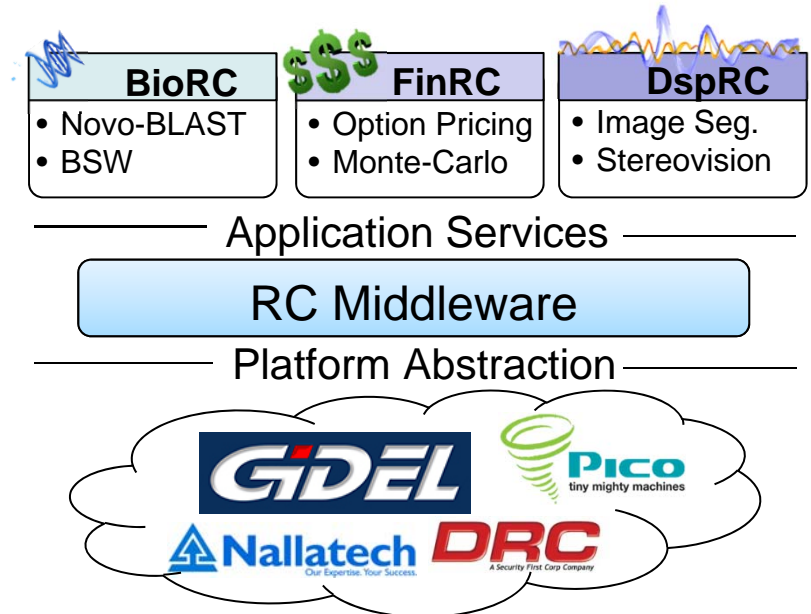
Motivation: Lack of standards between FPGA platforms limits app & tool portability, and productivity

- Major factor in limited acceptance of FPGAs in HPC community

Approach: *RC Middleware* (RCMW) provides uniform, standardized interfaces and programming model across heterogeneous platforms

Portability

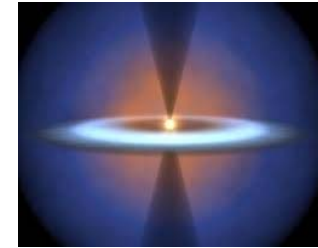
- ❑ Abstracts away *platform-specific interfaces* (hardware & software)
- ❑ Enables app & tool portability across *6 platforms* from *four vendors*
- ❑ Modest overhead
 - <1% area, <10% performance



Productivity

- ❑ RCMW toolset handles application resource *mapping* and *translation*
- ❑ Productivity improvement:
 - Hardware & software interfaces simplify HPC app development

Conclusions



Opportune Convergence of:

- ❑ *Escalating demands* of key HPC domains
- ❑ Emergence of *high-performance RC devices* and productivity *tools*
- ❑ Expanding *availability* of *HPRC systems*
- ❑ *Promising* R&D and commercial *HPRC apps*

And an opportune time for
High-Performance Reconfigurable Computing

For more details on apps and middleware,
Come see us at the **CHREC booth #2405**

Time for

