Ultimate DataFlow SuperComputing for BigData Analytics

V. Milutinovic, University of Belgrade, SRB

O. Mencer, Imperial College, London, GBR

M.J. Flynn, Stanford University, Palo Alto, USA

Major Sources of Inspiration

A. Richard Feynman:

Impact of logic/arithmetic and memory/IO Compiler-generated execution graph

B. Ilya Prigogine:

Impact of energy, entropy, order, and optimization Compiler-generated data separation

C. Daniel Kahneman:

Impact of approximate computing on precision Compiler-controlled approx computing

D. Tim Hunt:

Impact of system latency on precision Compiler-controlled system latency

The Major Axiom of Optimal Computing

A. Whenever the Technology changes, the Fundamental Paradigm of Computer Architecture has to change, too.

aSoG (not: FPGA)

B. If several paradigms are available, the most suitable paradigm for adoption is the one most effective for modern **Applications**.

BrontoData (not: ExaBigData)

Is the von Neumann Paradigm still the most effective one?

- A. MultiCores?
- B. ManyCores?

The Holy Trinity of Generalized Computing

Applications

Architecture

-Size

-Power

Speedup

-Precisior

Technology

The von Neumann Paradigm (1940s)

$$\lim_{i \to \infty} \left(\frac{TALU(i)}{TCOMM(i)} \right) \to \infty$$

Optimal Solution: Finite Automata

The Nobel Laureate Richard Feynman Observations

$$\lim_{i \to \infty} \left(\frac{TALU(i)}{TCOMM(i)} \right) \to 0 (t \to \infty)$$

Where is the technology now?

A. Closer to 1940s?

B. Closer to $t \rightarrow \infty$?

State of the Art in Technology Today The Power Challenge The Data Movement Challenge

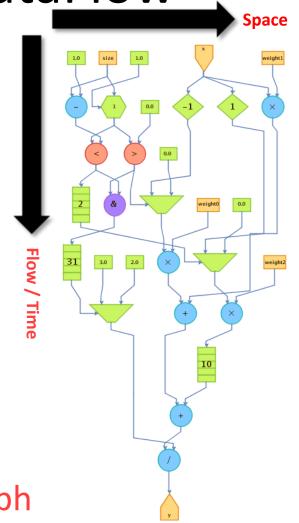
	2015	2020
Double precision FLOP	100pj	10рј

- Moving data off-chip will use 200x more energy than computing!
- Moving data in 1940s was using 1/60x ...
- Conclusion: We are getting close to the Feynman Asymptote!
- Important: Power and speed could be traded!

The Maxeler Technology Vision: MultiScale DataFlow

- ☐ Thinking in space rather than in time
- ☐ Difficult change in mindset to overcome
- ☐ Transformation of data through flow over time
- ☐ Instructions are parallelized across the available space

Optimal Solution: Execution Graph



Comparing the Two Approaches

• The Von-Neumann paradigm resembles an old wall clock



• The Feynman paradigm resembles lightning! Why?

Programming the Two Paradigms

von Neumann:

The Program Moves Data

Feynman:

The Program Configures Hardware

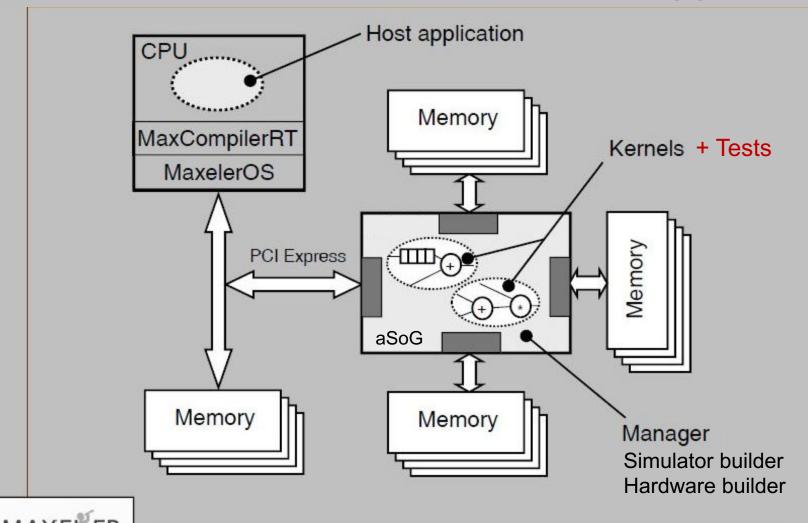
What moves data?

External sources till input.

Voltage difference through aSoG!

Voltage difference moves the important stuff!

The Maxeler Generic Architecture Application



Important: Supporting any CL and any OS!

Why The Acceleration Approach?

Nobel Laureate Ilya Prigogine: Injecting Energy to Decrease Entropy!

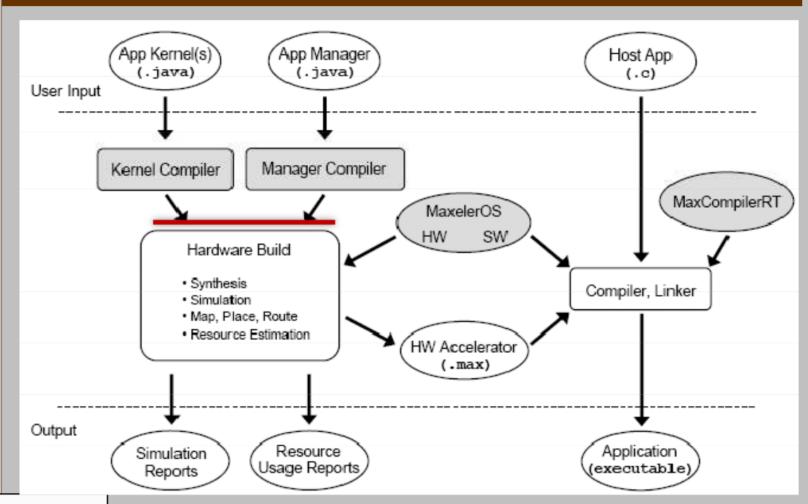
Corollary:

Burning energy to split spatial and temporal decreases the entropy of computing and enables the DataFlow compiler to create a maximally effective execution graph.

Final goal:

The execution graph with the minimal length of edges.

MaxCompiler



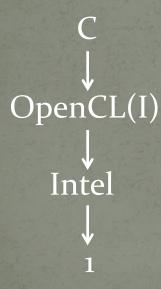


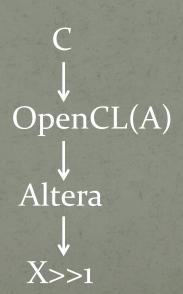
.max -> ASIC brings 30% to 50% in Speedup and Power, at the expense of no reconfigurability and no flexibility!

Alliances Being Formed

- Intel acquired Altera
- Qualcomm and IBM teaming up with Xilinx

However:







Nano Accelerators

- Invisible on the DataFlow Concept Level
- Invisible to DataFlow Programmers
- Visible to the MaxCompiler
- The MaxCompiler knows how to utilize them

Best protected by two aSoG (now FPGA) protection levels and two Vendor (e.g., Maxeler) protection levels!

Publications of Interest for NanoAcceleration

Flynn, M., Mencer, O., Milutinovic, V., Rakocevic, G., Stenstrom, P., Trobec, R., Valero, M., Moving from Petaflops (on Simple Benchmarks) to Petadata per Unit of Time and Power (On Sophisticated Benchmarks), Inspired by: Communications of the ACM (nano-acceleration), May 2013. Feynman 2. Trobec, R., Vasiljevic, R., Tomasevic, M., Milutinovic, V., Beiveide, M., Valero, M., Interconnection Networks for SuperComputing, ACM Computing Surveys (nano-acceleration), 2017. Milutinovic, V., Tomasevic, M., Markovic, B., Tremblay, M., 3. The Split Temporal/Spatial Cache: Initial Performance Analysis, Inspired by: Proceedings of the SCIzzL-5, Santa Clara, California, USA, March 26, 1996, pp 72-78. Prigogine Milutinovic, V., Tomasevic, M., Markovic, B., Tremblay, M., The Split Temporal/Spatial Cache: Initial Complexity Analysis, Proceedings of the SCIzzL-5, Santa Clara, California, USA, September, 1996. 5. Milutinovic, V., A Comparison of Suboptimal Detection Algorithms Applied to the Additive Mix of Orthogonal Sinusoidal Signals, Inspired by: IEEE Transactions on Communications, Vol. COM-36, No. 5, May 1988, pp. 538-543. Kahneman 6. Milutinovic, V., Mapping of Neural Networks on the Honeycomb Architectures, Proceedings of the IEEE, Vol. 77, No 12, December 1989, pp. 1875-1878. Helbig, W., Milutinovic, V., The RCA's DCFL E/D MESFET GaAs 32-bit Experimental RISC Machine, Inspired by: IEEE Transactions on Computers, Vol. 36, No. 2, February 1989, pp. 263-274. Hunt Jovanovic, Z., Milutinovic, V., FPGA Accelerator for Floating-Point Matrix Multiplication, IEE Computers & Digital Techniques (nano-acceleration), 2012, 6, (4), pp. 249-256. The IET 2014 Premium Award for Computing & Digital Techniques.

Platforms

Products

Technology

About Us



MyMaxele



Back...

Moving from Petaflops to Petadata

May 5, 2013

VIEWPOINT: Moving from Petaflops to Petadata

M. Flynn^{‡||}, O. Mencer^{‡‡}, V. Milutinovic[†], G. Rakocevic[§], P. Stenstrom⁸, R.Trobec^b, M. Valero[‡]

[‡]Maxeler Technologies, ^{II}Stanford University, [‡] Imperial College London, [†]University of Belgrade, [§]Mathematical Institute of the Serbian Academy of Sciences and Arts in Belgrade, ⁸ Chalmers University of Technology, [§]Jožef Stefan Institute, [‡]Barcelona Supercomputing Centre

Communications of the ACM, Vol. 56 No. 5 May 2013, doi: 10.1145/2447976.2447989

Special Acknowledgements to: Simon Aglionby, Georgi Gaydadjiev, Itay Greenspon, and Nemanja Trifunovic

IF(2012)=3.80

ACM Computing Surveys

From Wikipedia, the free encyclopedia

ACM Computing Surveys (CSUR) is a peer reviewed scientific journal published by the Association for Computing Machinery. The journal publishes survey articles and tutorials related to computer science and computing. It was founded in 1969; the first editor-in-chief was William S. Dorn.^[1]

In ISI Journal Citation Reports, ACM Computing Surveys has the highest impact factor among all computer science journals.^[2] In a 2008 ranking of computer science journals, ACM Computing Surveys received the highest rank "A*".^[3]

See also [edit]

· ACM Computing Reviews

References [edit]

- 1. ^ Dorn, William S. (1969). "Editor's Preview...". ACM Computing Surveys. 1 (1): 2-5. doi:10.1145/356540.356542 dd.
- 2. A "Journal Citation Reports" . ISI Web of Knowledge. Retrieved 2009-10-03. "JCR Science Edition 2008"; subject categories "COMPUTER SCIENCE, ...".
- 3. ^ "Journal Rankings" ₽. CORE: The Computing Research and Education Association of Australasia. July 2008. Archived ₱ from the original on 29 March 2010. Retrieved 2010-03-19..

External links [edit]

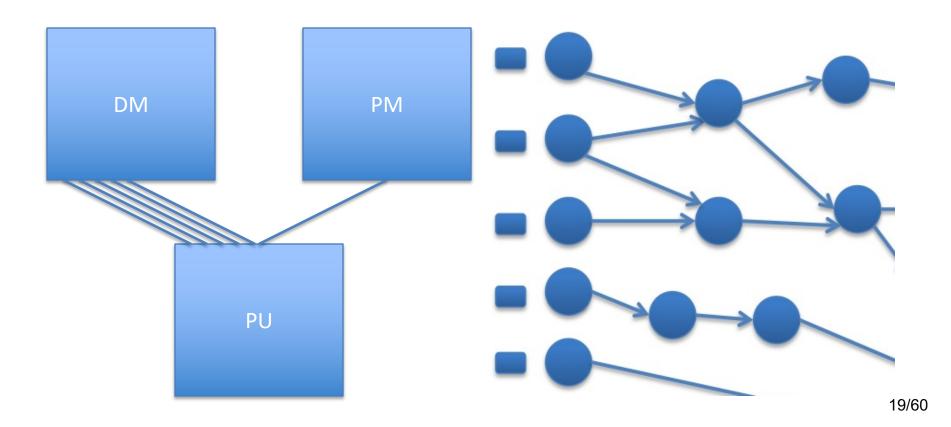
- . ACM Computing Surveys in ACM Digital Library.
- . ACM Computing Surveys do in DBLP.

ACM Computing Surveys

Abbreviated title (ISO 4)	ACM Comput. Surv.			
Discipline	Computer science			
Language	English			
Edited by	Sartaj K Sahni			
Publication details				
Publisher	ACM (United States)			
Publication history	1969-present			
Frequency	Quarterly			
Indexing				
ISSN	0360-0300 & (print)			
	1557-7341년 (web)			
Links				
Journal homepage				
Online access				
 Online archive				

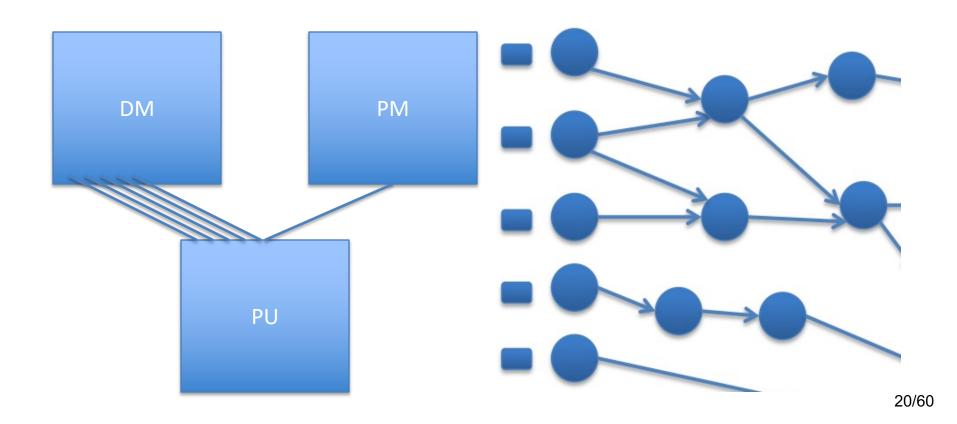
Essence: Feynman Enabled by Prigogine

- TALU possible at zero power (Arithmetic+Logic)
- TCOMM not possible at zero power (MEM+MPS)



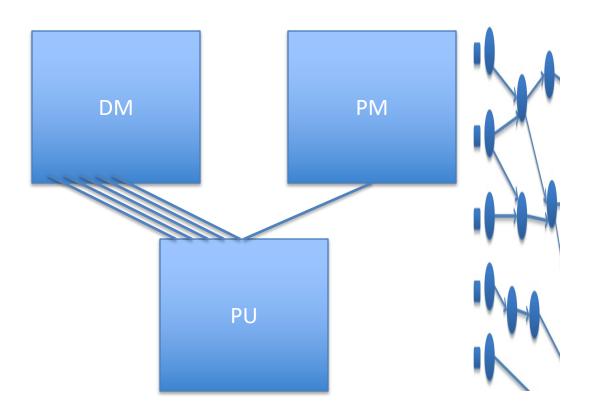
Essence: Feynman

- TALU possible at zero power (Arithmetic+Logic)
- TCOMM not possible at zero power (MEM+MPS)



Essence: Feynman

- TALU possible at zero power (Arithmetic+Logic)
- TCOMM not possible at zero power (MEM+MPS)



Programming the Maxeler Technology Generic Acceleration Architecture

MaxJ, the Maxeler Java, a DSL acting as a SuperSet of classical Java: A. A vector of built-in domain-specific classes B. Two sets of variables: SW + HW

MaxJ is a SubSet of OpenSPL, created by the Imperial-Stanford-Tokyo-Tsinghua consortium.

Possible Future Mutations of OpenSPL:
MaxPython and/or MaxR
(lower Kolmogorov complexity)
MaxHaskel and/or MaxScala
(easier extension to approximate computing).

Approximate Computing for Better Precision: Kahneman

Note: Small approximations in one domain may bring large benefits in another domain

Example: Weather forecast

A 15-bit computational precision (rather than the 64-bit precision) may decrease the forecast precision for only 2%, and at the same time, may increase the grid precision 25 times, and the forecast precision at grid intersections up to 10⁴.

Easily doable in DataFlow, difficult to do in ControlFlow.

Delayed Decision for Better Precision: Hunt

Note: Small latencies in time domain may bring large benefits in precision domains

Example: Optimal utilization of internal DataFlow pipelines

Compiler optimizations create internal pipelines that experienced DataFlow programmers know how to utilize

BigDataAnalytics

Existing Maxeler-based publications:

Ultimate aSoG-based future:

20 [Size] 20 [Power] 20, 200 [Speedup] 20 [Precision]

Applications

20-200 [Size] 20-200 [Power] 20, 200, 2000, 20000 [Speedup] 20+ [Precision]

Architecture

Technology

Maxeler Dataflow Appliance

- Software Based Solution
- Dataflow Computing in the Datacentre





The CPU

Conventional CPU cores and up to 6 DFEs with 288GB of RAM







The Dataflow Appliance

Dense compute with 8 DFEs, 768GB of RAM and dynamic allocation of DFEs to CPU servers with zero-copy RDMA access





The Networking Appliance

Intel Xeon CPUs and 4 DFEs with direct links to up to twelve 40Gbit Ethernet connections





MicroMAX.5: Good for Dew (Edge Processing of IoT Data)



The Major Application Successes

• Finances:

- Credit derivatives
- Risk assessment
- Stability of economical systems
- Evaluation of econo-political mechanisms

GeoPhysics:

- Oil&Gas
- Weather forecast
- Astronomy
- Climate changes

· Science:

- Physics
- Chemistry
- Biology
- Genomics
- Engineering: Synergy of all the Above (ML, etc...)



J.P.Morgan

Field Programmable Gate Arrays (FPGAs)

A Field Programmable Gate Array (FPGA) is a silicon chip containing a matrix of configurable logic blocks (CLBs) that are connected through programmable interconnects. By combining optimized use of available silicon with fine-grained parallelism, sustained acceleration improvements of over 300x can be achieved across a range of vanilla and complex mathematical models. The current work is the first time that FPGA technology has been employed at this scale to accelerate computational performance anywhere in the finance industry.

Power and Versatility

- Can accelerate performance by between 100 and 1,000x across a range of mathematical models, with the ability to perform a task in less than a second
- Can be reprogrammed and precisely configured to compute exact algorithm(s) at the desired level of numerical accuracy required by any given application, unlike normal microprocessors whose design is fixed by the manufacturer
- Can be deeply pipelined to achieve maximum parallelism from arithmetic, algorithms and data streaming

Key Business Challenges

- Reduce the execution time of existing applications to meet business and regulatory demands
- · Decrease cost of running existing applications and developing new ones
- Provide fast, cost-effective extra computational capacity to address problems that are currently inextricable
- Achieve a step-change improvement in price-performance and end-to-end compute time across many applications

Key Benefits (Business/Clients)

- Competitive advantage to valuation, execution, risk management and complex scenario analyses by speeding up existing applications
- Lower cost of existing applications as hardware costs can be reduced by a factor between 100 and 1,000
- Ability to perform previously difficult calculations, such as complex trading strategies or risk evaluations of global portfolio simulations.

Technology Overview

- Low clock speed chips
- Maximal usage of available silicon resources
- Acceleration through use of fine-grained parallelism
- Reconfigurable hardware
- Silicon configurable to fit algorithm

LOB/Function(s) Impacted

- Credit & interest rates
- Equities & commodities
- Loan & mortgage modeling
- Finance & accounting
- High frequency trading
- Risk management & VaR

Industry/External Recognition

- Used by Cisco in all routers
- Simulation of real and theoretical systems
- Geophysics for oil and gas exploration
- Astrophysics & hydrodynamics
- Defense for cryptography
- Video games
- Genotyping

Functionality Overview

Double precision floating point-capable FPGAs became commercially available in 2002, but it was the arrival of the Virtex 5 and 6 series chips from market leader Xilinx that really provided the scale required for the development of production-grade accelerated solutions. Using FPGAs in high performance compute solutions provides distinct advantages over conventional CPU clusters.

Operational Advantages

- Significantly increases performance for two main types of applications: those based around highly complex mathematical models and those using simpler algorithms that can be massively parallelized
- Enables a dramatic increase in compute density per cubic meter by using FPGAs as computational accelerators
- · Consumes around 1% of the power of a single CPU core

Performance Improvements

- Performance improvements in the range 200-300x faster than the existing CPU cores used on the Compute BackBone (CBB) have been achieved in credit and interest rates hybrids businesses
- In equities, direct market access can run risk and loan stock at wire speed
 (3.5 micro secs) using a low-latency FPGA solution
- Benchmarked average throughput for J.P. Morgan's existing 40-node hybrid FPGA machine of 984MFlops/watt/cubic meter
- Potential standing at the top of the Green-500 ecological global supercomputer performance table

Development/Delivery

Timeline

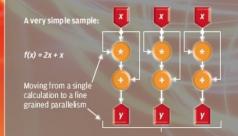
- Initial porting of an algorithm can vary from one to three months depending on complexity.
- Production capabilities then depend on the scale of the application and the scope and intensity of the testing and reconciliation cycle

Partners

- London-based Applied Analytics group: includes three technology and business specialists with extensive experience in developing and delivering high performance solutions across a range of asset classes, models and lines of business
- Maxeler Technologies: external consultants trained in Imperial College, Stanford and MIT research labs

FPGAs at Work

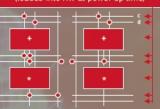
- An algorithm is implemented as a special configuration of a general purpose electric circuit
- Connections between prefabricated wires are programmable
- Function of calculating elements is itself programmable
- FPGAs are two dimensional matrix-structures of configurable logic blocks (CLBs) surrounded by input/output blocks that enable communication with the rest of the environment



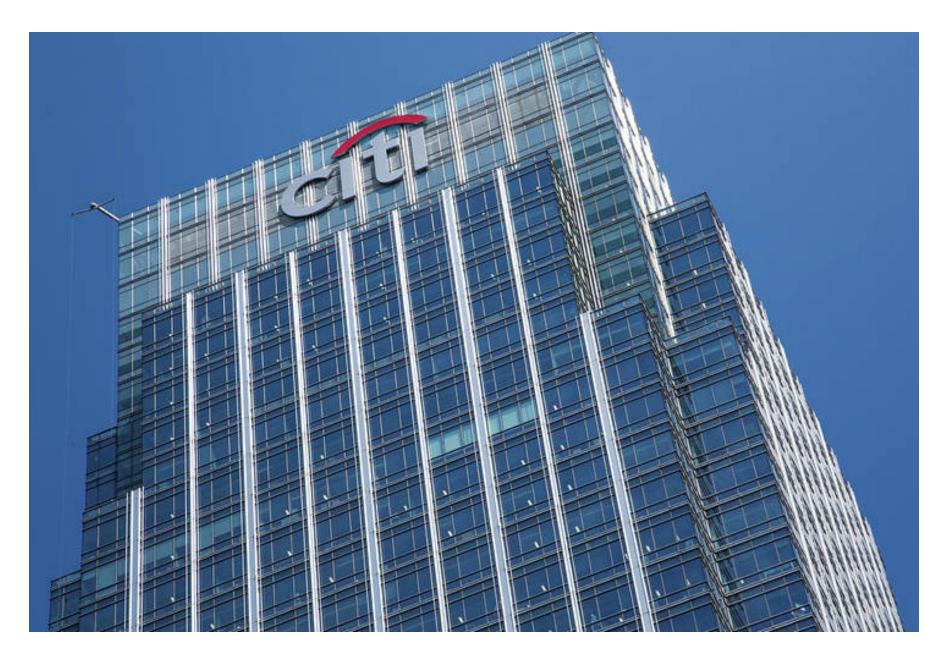
A slightly more complex example:

e = (a+b)*(c+d)

Configuration Memory (loaded into HW at power up time)



Migrating algorithms from C++ to FPGAs involves doing a Fourier Transform from time domain execution to spatial domain execution in order to maximize computational throughput. It's a paradigm shift to stream computing that provides acceleration of up to 1,000x compared to an Intel CPU.



The know-how needed for deep security!



www.cmegroup.com/trading/interest-rates/dsf-analytics.html

Designed for educational use only using Maxeler Technologies' curve construction methodology. This tool uses delayed data and displayed results are indicative representations only.

Please hover your mouse pointer over column titles and links for further information.

CME Tieles	Bloomberg	DSF Pricing					- !
	Ticker	Price	Coupon	PV01	NPV	Implied Rate	Timestamp
T1UM4 2Y	CTPM4	100'057	0.750%	\$19.97	\$179.69	0.6600%	4:00:03 PM CT 4/4/2014
F1UM4 5Y	CFPM4	100'115	2.000%	\$48.49	\$359.38	1.9259%	4:00:03 PM CT 4/4/2014
N1UM4 10Y	CNPM4	100'225	3.000%	\$90.16	\$703.12	2.9220%	4:00:03 PM CT 4/4/2014
B1UM4 30Y	CBPM4	102'270	3.750%	\$195.07	\$2,843.75	3.6042%	4:00:03 PM CT 4/4/2014
T1UU4 2Y	CTPU4	100'085	1.000%	\$19.93	\$265.62	0.8668%	4:00:03 PM CT 4/4/2014
F1UU4 5Y	CFPU4	100'110	2.250%	\$48.27	\$343.75	2.1788%	4:00:03 PM CT 4/4/2014
N1UU4 10Y	CNPU4	101'125	3.250%	\$89.55	\$1,390.62	3.0948%	4:00:03 PM CT 4/4/2014
B1UU4 30Y	CBPU4	106'020	4.000%	\$193.47	\$6,062.50	3.6868%	4:00:03 PM CT 4/4/2014

Quotes and analytics are updated every 15 minutes.

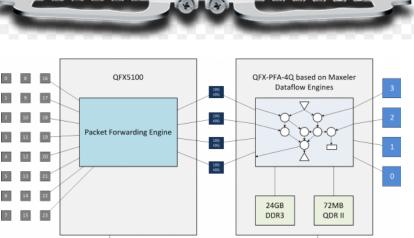
(1) Analytics powered by Maxeler Technologies®

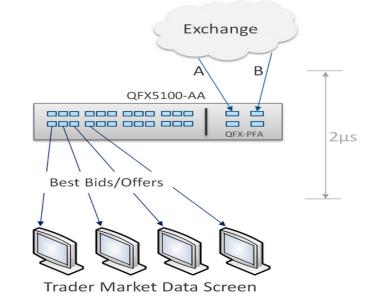
Instrument	CPU 1U-Node	Max 1U-Node	Comparison	
European Swaptions	848,000	35,544,000	42x	
American Options	38,400,000	720,000,000	19x	
European Options	32,000,000	7,080,000,000	221x	
Bermudan Swaptions	296	6,666	23x	
Vanilla Swaps	176,000	32,800,000	186x	
CDS	432,000	13,904,000	32x	
CDS Bootstrap	14,000	872,000	62x	

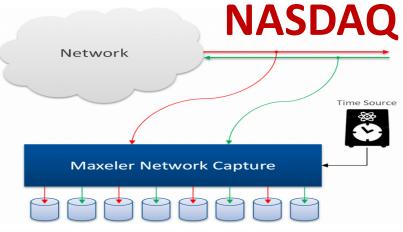
MAXELEER Technologies

Juniper for Online Trading





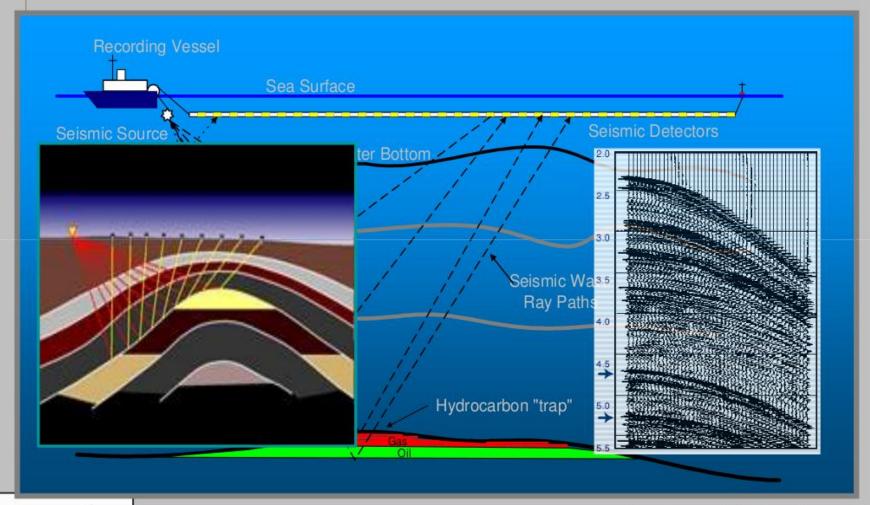








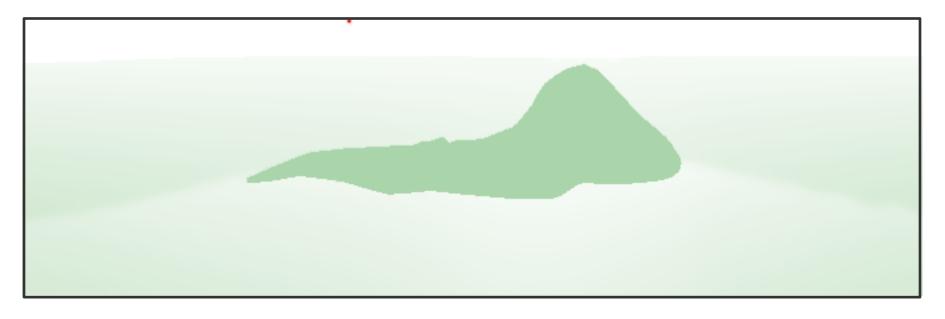
Seismic Data Acquisition





Courtesy of Schlumberger

Seismic Imaging



Running on MaxNode servers

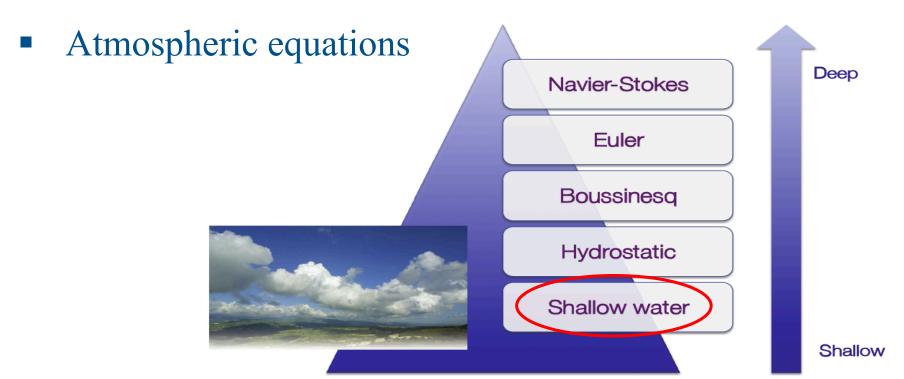
- 8 parallel compute pipelines per chip
- 10x less power: 150MHz vs 1.5GHz
- 30x faster than microprocessors

An Implementation of the Acoustic Wave Equation on FPGAs

T. Nemeth[†], J. Stefani[†], W. Liu[†], R. Dimond[‡], O. Pell[‡], R.Ergas[§]

[†]Chevron, [‡]Maxeler, [§]Formerly Chevron, SEG 2008

Global Weather Simulation: Size is Relevant



Equations: Shallow Water Equations (SWEs)

$$\frac{\partial Q}{\partial t} + \frac{1}{\Lambda} \frac{\partial (\Lambda F^1)}{\partial x^1} + \frac{1}{\Lambda} \frac{\partial (\Lambda F^1)}{\partial x^2} + S = 0$$

[L. Gan, H. Fu, W. Luk, C. Yang, W. Xue, X. Huang, Y. Zhang, and G. Yang, Accelerating solvers for global atmospheric equations through mixed-precision data flow engine, FPL2013] Tsinghua 34/60













Weather Model – Performance Gain

Platform	Performance	Speedup
6-core CPU	4.66K	1
Tianhe-1A node	110.38K	23x
MaxWorkstation	468.1K	100x
MaxNode	1.54M	330x

Meshsize: $1024 \times 1024 \times 6$

14x

MaxNode speedup over Tianhe node: 14 times













Weather Model -- Power Efficiency

Platform	Power Efficiency	Speedup
	()	
6-core CPU	20.71	1
Tianhe-1A node	306.6	14.8x
MaxWorkstation	2.52K	121.6x
MaxNode	3K	144.9x

Meshsize: $1024 \times 1024 \times 6$

9 x

MaxNode is 9 times more power efficient





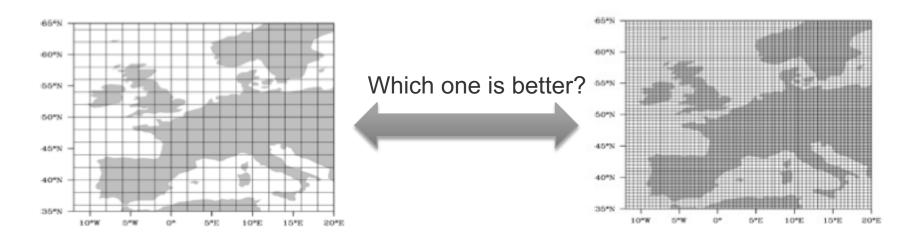








Weather and Climate Models: Precision



Finer grid and higher precision are obviously preferred but the computational requirements will increase → Power usage → \$\$

What about using reduced precision? (15 bits instead of 64 double precision FP)



We use only 15 bits for 98% of the computation:



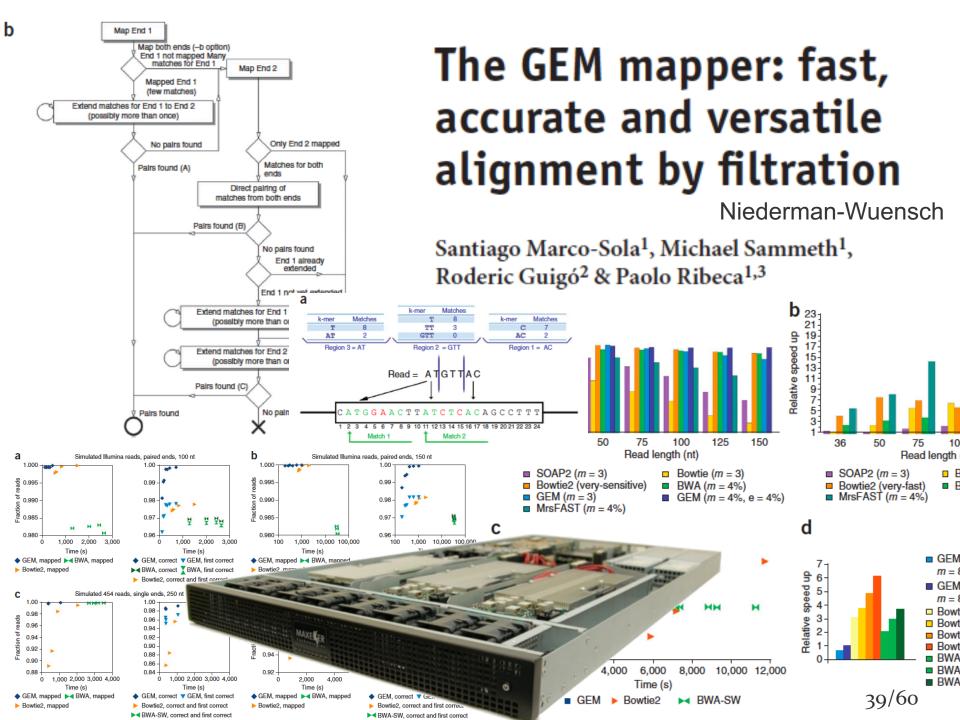
37/60



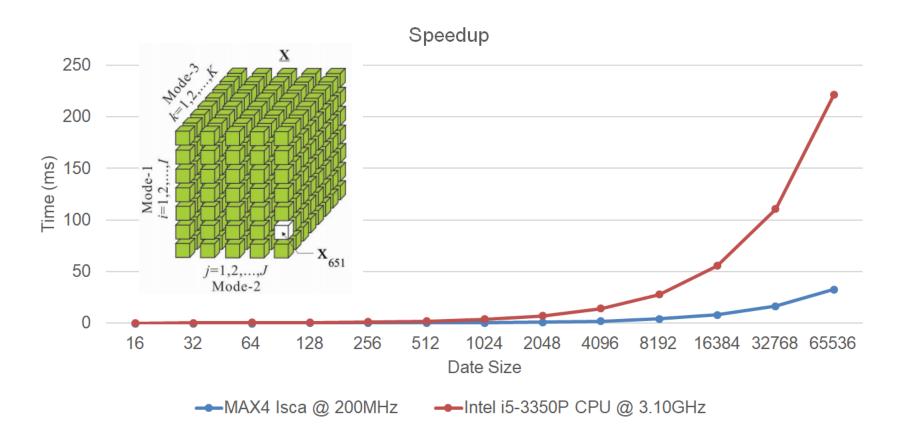
Maxeler Running Smith Waterman







Analysis of the Tensor Calculus Operations on DataFlow (PhD Thesis by Miloš Kotlar, on DataFlow-based Machine Learning)

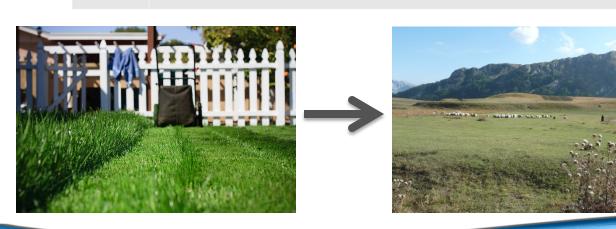


The speedup of **6.75x** achieved as early as for KiloData (Perceptron), with **10x** less on-chip transistors and the power savings of **4.6x**



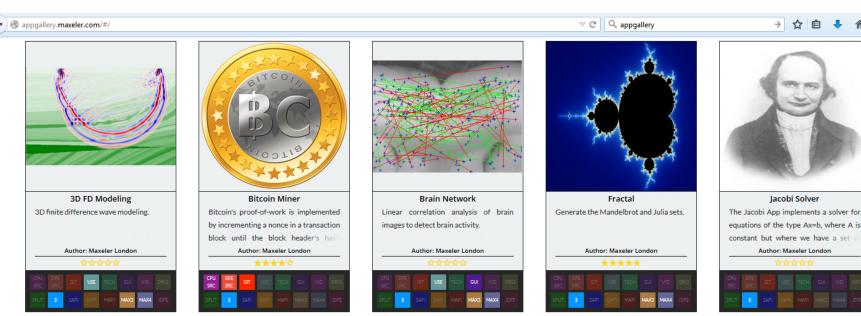
Conditions for the Y-Chart-Based "Kernelization" of Loops @ML (PhD Thesis by Nenad Korolija, on the Mapping of Algorithms onto DataFlow)

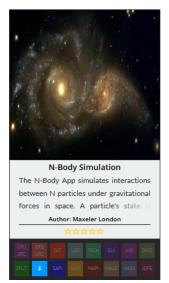
1.	BigData (RAM vs. STREAM)	$O(n^2)$
2.	Code reusability (WORO vs. WORM)	+
3.	Overall application tolerance to latency	+
4.	Over 95% of run time in loops	++
5.	Reusability of the data in loops	++
6.	Potential for utilization of pipes	O(n)

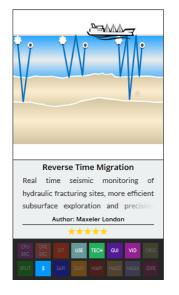


Essentials for speedup: algorithmic modifications, pipeline utilization, data choreography, decision making on precision

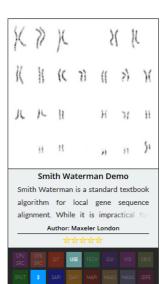


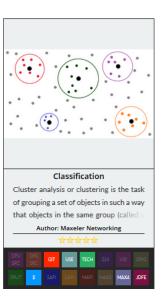




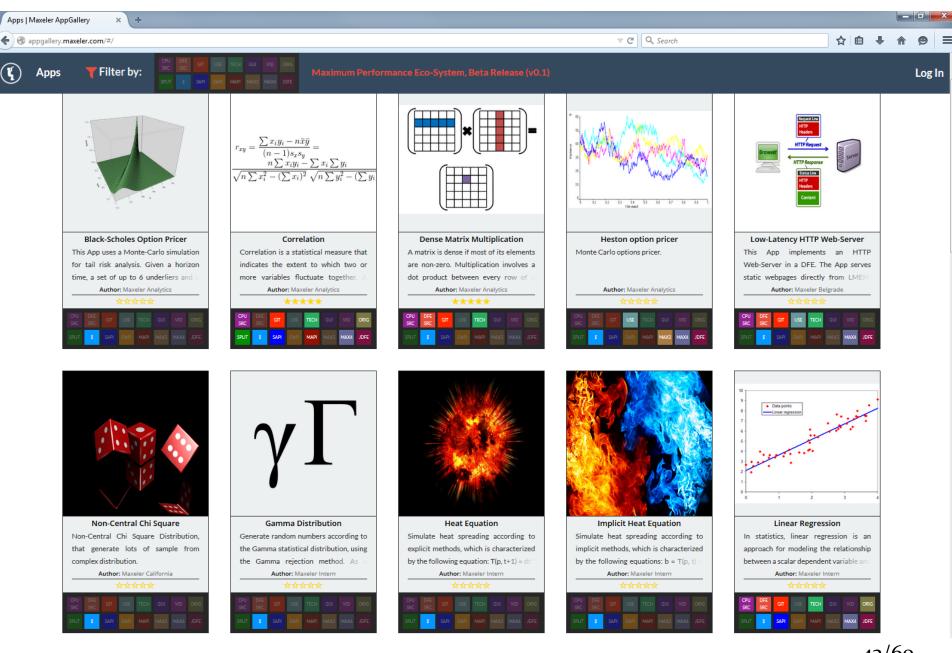


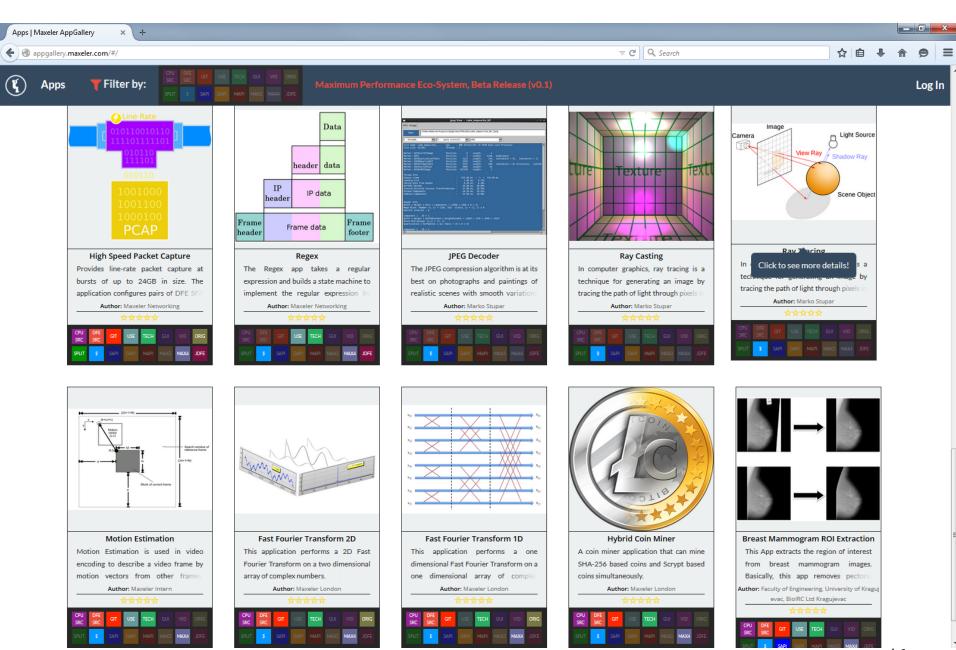


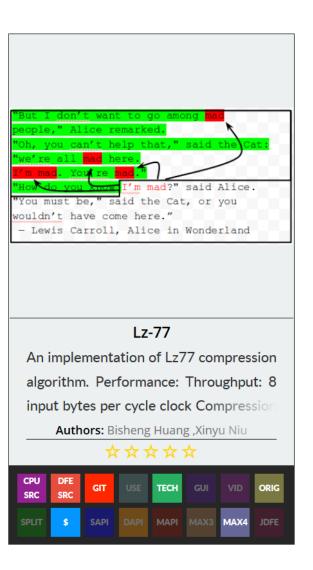


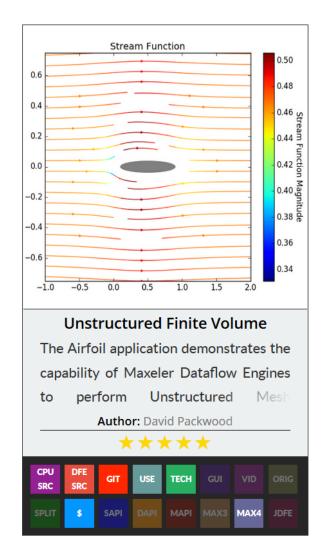


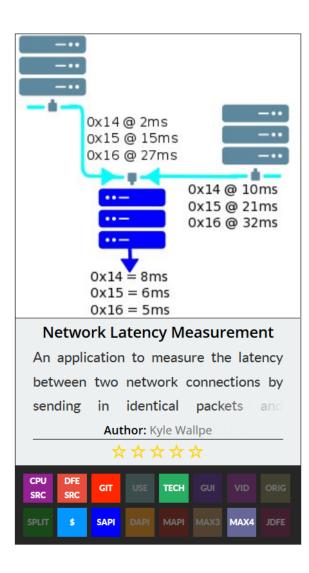
appgallery.maxeler.com





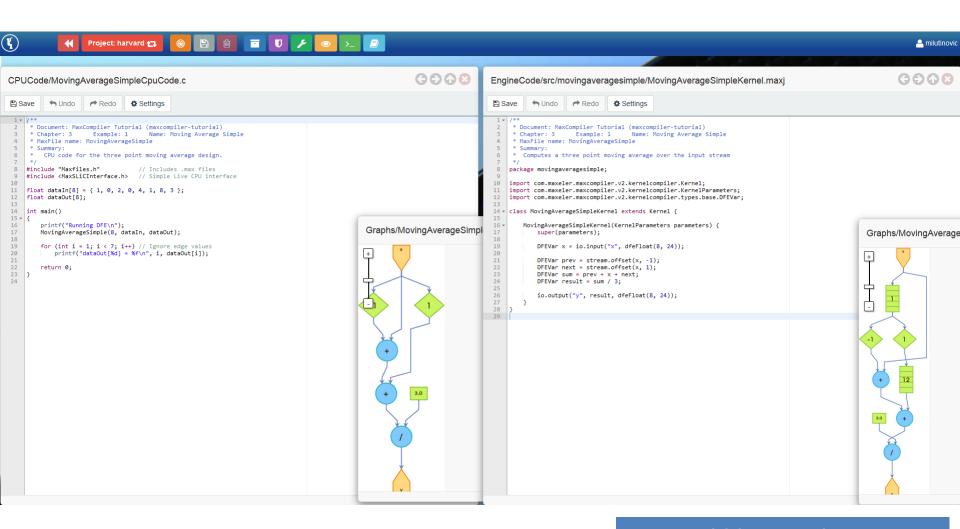






appgallery.maxeler.com





webide.maxeler.com https://maxeler.mi.sanu.ac.rs





MultiCore



Which way are the horses going?





ManyCore

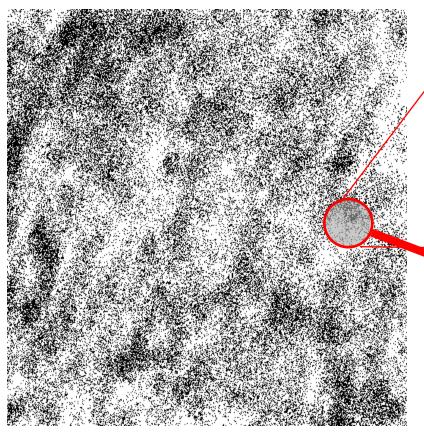
• Is it possible to use 2000 chicken instead of two horses?

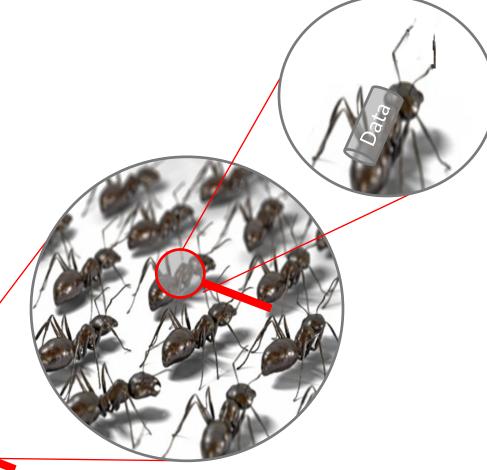






What is better, real and anecdotic?

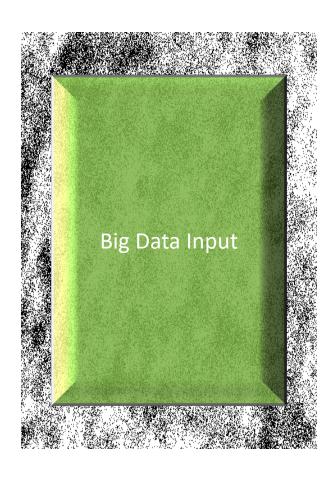


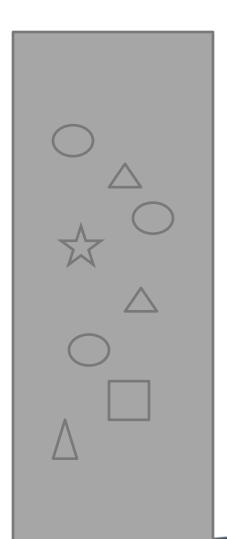


How about 2 000 000 ants?

52/60

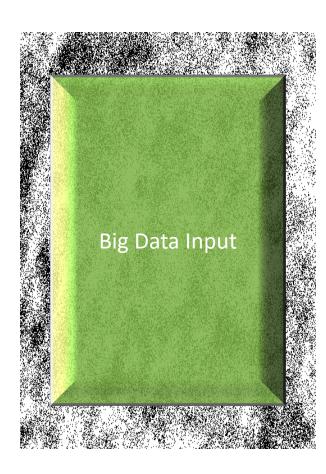




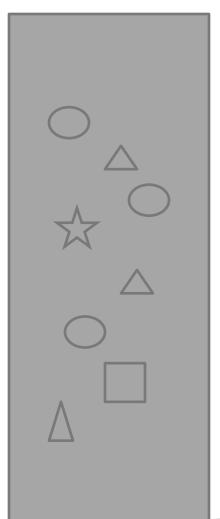


53/60





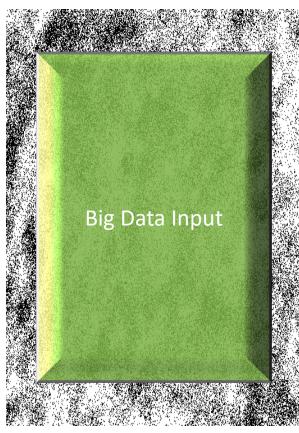
aSoG



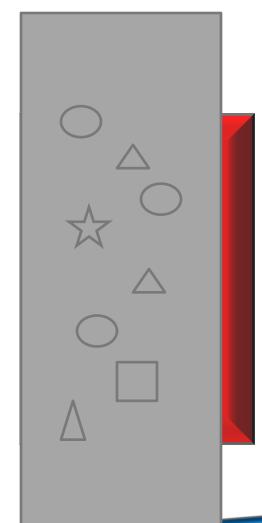




Bronto



aSoG





Marmelade

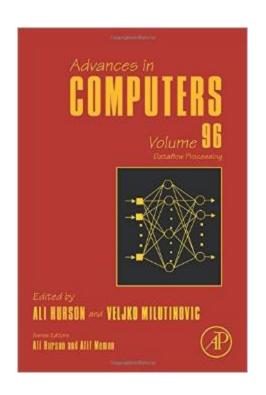
55/60



An Edited Book Covering the Applications

- http://www.amazon.com/Dataflow-Processing-Volume-Advances-Computers/dp/0128021349
- http://www.elsevier.com/books/dataflow-processing/milutinovic/978-0-12-802134-7

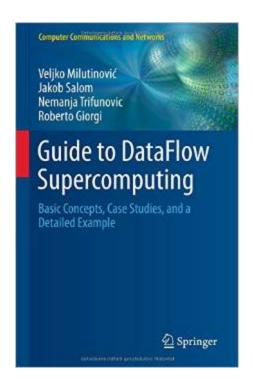
Indexed by: WoS (SCI)



Contributions welcome for the follow-ups: Vol. 102 + Vol. 104 + etc...

An Original Book Covering the Essence

- http://www.amazon.com/Guide-DataFlow-Supercomputing-Concepts-Communications/dp/3319162284
- <u>http://www.springer.com/gp/book/9783319162287</u>



The first source to use the term the Feynman Paradigm in contrast with the Von Neumann Paradigm

InformationWeek CONNECTING THE BUSINESS TECHNOLOGY COMMUNITY

Search InformationWeek Q

News & Commentary INTEROP Follow IW: Home Authors Slideshows Video Reports White Papers Events University STRATEGIC CIO SOFTWARE **BIG DATA** INFRASTRUCTURE DEVELOPER SECURITY MOBILE INDUSTRIES IT LIFE

CLOUD // SOFTWARE AS A SERVICE



Google I/O: Hello Dataflow, Goodbye MapReduce

Google introduces Dataflow to handle streams and batches of big data, replacing MapReduce and challenging other public cloud services.



Charles Babcock News

Connect Directly







Google I/O this year was overwhelmingly dominated by consumer technology, the end user interface, and extension of the Android universe into a new class of mobile devices, the computer you wear on your wrist.

At the same time, there were one or two enterprise-scale data handling and cloud

computing gems scattered among all the end user announcements.



Hadoop Jobs: 9 Ways To Get Hired

(Click image for larger view and

slideshow.)

OPERATIONS
PLATFORM
OF AWS
Powering automated

CLOUD







Comment





Intel says logic is faster than GPUs





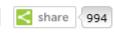


f Share











Intel's Programmable Systems Group takes its first step towards FPGA based system in package portfolio

Speaking in 2012, Danny Biran – then Altera's senior VP for corporate strategy – said he saw a time when the company would be offering 'standard products' devices featuring an FPGA, with different dice integrated in the package. "It's also possible these devices may integrate customer specific circuits if the business case is good enough," he noted.

There was a lot going on behind the scenes then; already, Altera was talking with Intel about using its foundry service to build 'Generation 10' devices, eventually being acquired by Intel in 2015.



Jordan Inkeles, Altera's director of product marketing for high end **FPGAs**

Now the first fruit of that work has appeared in the form of

Stratix 10 MX. Designed to meet the needs of those developing high end communications systems, the device integrates stacked memory dice alongside an FPGA die, providing users with a memory bandwidth of up to 1Tbyte/s. 28 June 2016

QoL

Maxeler is one of the Top 10 HPC projects to impact QoL in the World:)

Scientific Computing

[www.scientificcomputing.com/articles/2014/11]

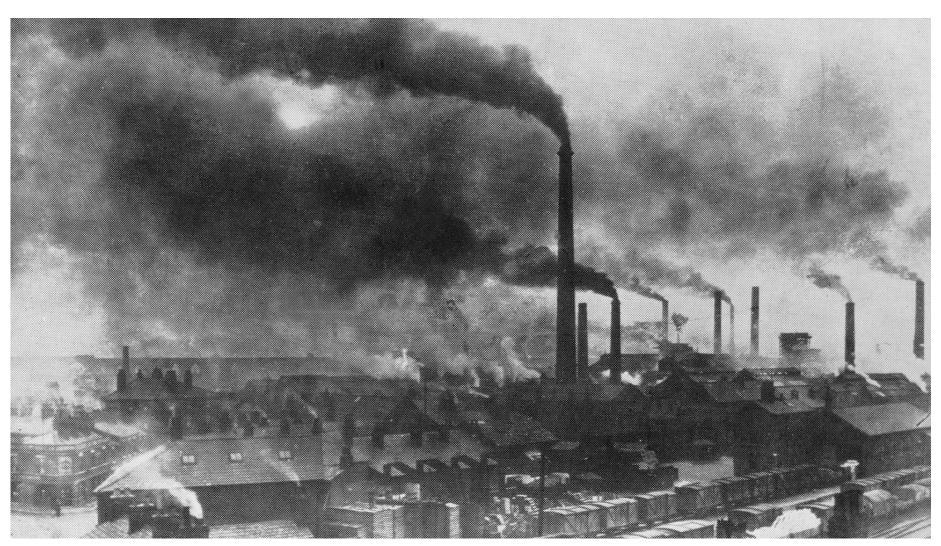
by

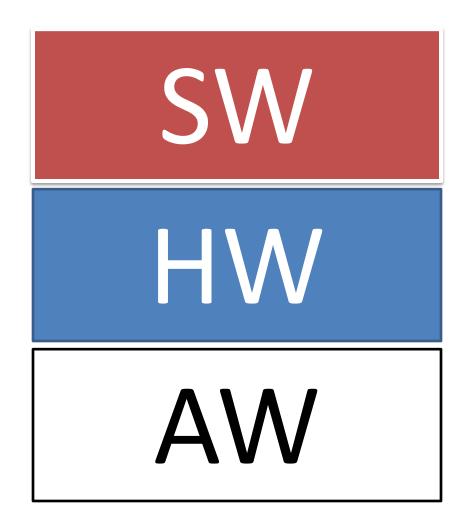
Don Johnson

of

Lawrence Livermore National Labs [editor@ScientificComputing.com]

How About QoL?





Essence of the Paradigm:

For Big Data algorithms and for the same hardware price as before, achieving:

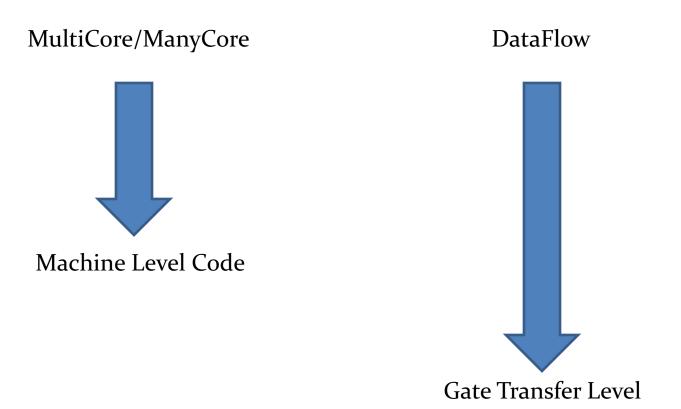
- a) speed-up, 20-200
- b) monthly electricity bills, reduced 20 times
- c) size, 20 times smaller
- d) precision, X times better

The major issues of engineering are: design cost and design complexity.

Remember, economy has its own rules: production count and market demand!

Why is DataFlow so Much Faster?

• Factor: 20 to 200



Why are Electricity Bills so Small?

Factor: 20

MultiCore/ManyCore

DataFlow





$$P = \mathbb{R}^2$$

Why is the Cubic Foot so Small?

Factor: 20

MultiCore/ManyCore

DataFlow

Data Processing

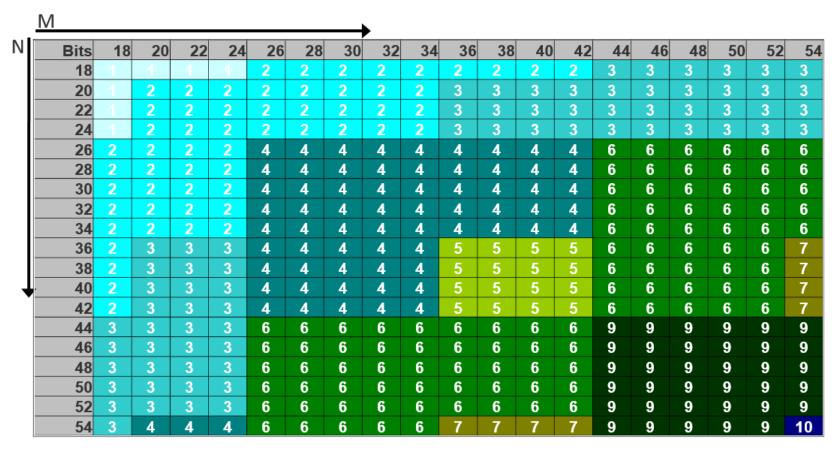
Process Control

Process Control

Process Control

Why is the Precision Better?

Factor: X



Successes of 2018 and 2019

Hitachi Cloud Amazon AWS BQCD Quark Endorsed by Jerome Friedman





(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2018/0189063 A1 FLEMING et al.

- Jul. 5, 2018 (43) **Pub. Date:**
- (54) PROCESSORS, METHODS, AND SYSTEMS WITH A CONFIGURABLE SPATIAL ACCELERATOR
- (71) Applicant: Intel Corporation, Santa Clara, CA (US)
- (72) Inventors: KERMIN FLEMING, Hudson, MA (US); KENT D. GLOSSOP. Merrimack, NH (US); SIMON C. STEELY, Jr., Hudson, NH (US)
- (21) Appl. No.: 15/396,395
- (22) Filed: Dec. 30, 2016

Publication Classification

(51) Int. Cl. G06F 9/30 (2006.01)G06F 13/42 (2006.01) (52) U.S. Cl. CPC G06F 9/3016 (2013.01); G06F 13/4221 (2013.01)

(57)ABSTRACT

Systems, methods, and apparatuses relating to a configurable spatial accelerator are described. In one embodiment, a processor includes a core with a decoder to decode an instruction into a decoded instruction and an execution unit to execute the decoded instruction to perform a first operation; a plurality of processing elements; and an interconnect network between the plurality of processing elements to receive an input of a dataflow graph comprising a plurality of nodes, wherein the dataflow graph is to be overlaid into the interconnect network and the plurality of processing elements with each node represented as a dataflow operator in the plurality of processing elements, and the plurality of processing elements are to perform a second operation by a respective, incoming operand set arriving at each of the dataflow operators of the plurality of processing elements.

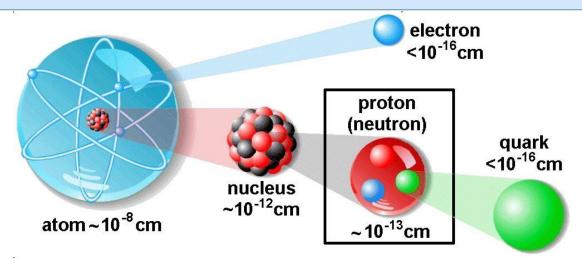
BQCD on a Maxeler Dataflow Computer





BQCD on a Dataflow Computer

Porting BQCD from BlueGene to a Maxeler Dataflow Computer



- Quantum Chromodynamics (QCD): models interactions of subatomic particles
- ◆ Lattice QCD (LQCD): its discretisation, suitable for numerical computation
- Berlin QCD (BQCD): most popular implementation of the LQCD algorithm
- Conjugate Gradient (CG): Majority of the compute time (benchmark: 68%)
 - ◆ CG iteratively solves linear algebra problem of form *Mx* = *b*
 - Operator M contains Wilson-dslash and Clover operators
- ◆ PROJECT TARGET 40x speedup of CG part of BQCD, followed by speedup of the entire application by 20x comparing same size boxes Dataflow vs BlueGene/Q



Maxeler QCD - Deployment

Maxeler QCD solution is deployed at Jülich Supercomputing Center, running on a Maxeler Dataflow system.

	2 racks of Jülich BlueGene/Q machine	On-premise Maxeler Dataflow system: scale to 1PF equivalent	Factor
Volume	6.75 m ³	0.87 m ³	7.76
Overall Time to Solution	1576.60 s	689 s	2.29
Overall Energy to Solution	169.6 kWh	4.42 kWh	38.4
Volume x TTS	10,642.05 m ³ s	599.43 m³s	17.8

• Evaluate 64x64x64x64 problem, 5 MC steps, 200 HMC steps Volume x TTS x ETS = 682







Maxeler QCD on Amazon EC2 F1

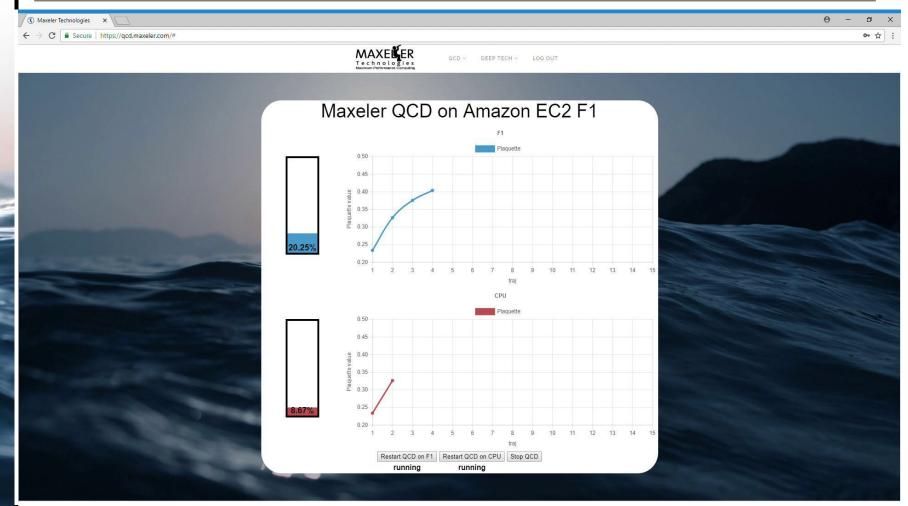
The Maxeler QCD solution is now running on Amazon EC2 F1: Software portable from Maxeler Dataflow system to Amazon Cloud

- Elastic computing: expand from on-premise to Cloud
- Scale up computation as workload grows
- Accelerated HPC as a service





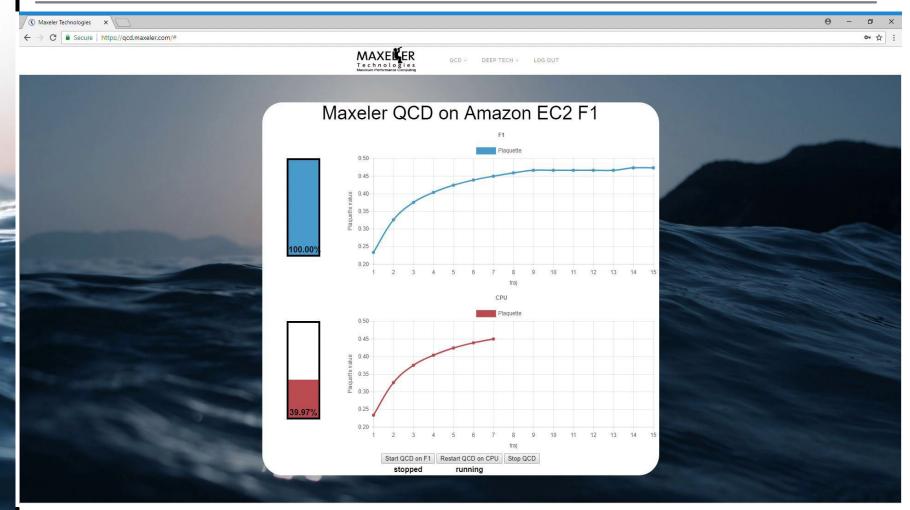
QCD Demo







QCD Demo

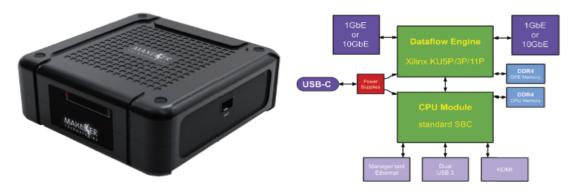






miniMAX5 Edge of IoT Platform





Dimensions	5.7 in (144mm) Wide x 5.7 in Deep (144mm) x 2.2 in High (57mm), excluding power supply				
Form factor	Desktop enclosure, fanless design. Wall or rail mount options available				
Weight	34 oz (950g), excluding power supply				
Power Supply	Separate wall plug unit, providing 60W of USB-PD power from 100-240V, 50-60Hz mains				
	Ethernet	1GbE or 10GbE (copper or fibre)	SFP+ Cage		
Input and Output (Standard Ports)	USB-C	Power input over USB-PD (Min 15V 3A supply required) USB-3 SuperSpeed II I/O on same connector supporting DisplayPort Alternate mode			
Instruction of October 14	Management LAN	1GbE	RJ45		
Input and Output (Optional Ports)	USB	Dual USB-3 Type A ports			
(Optional Ports)	Video Output	HDMI Type A			
	СРИ	AMD 3rd Generation R- or G-Series - choose from - Quad Core Merlin Falcon RX-416GD @1.6GHz - Dual Core Brown Falcon GX-217 @17GHz	Other SBC options available on request		
Controlflow	Memory	2x 4GB DDR4-2133 SODIMM, total 8GB	Higher or lower		
Engine	Storage	64GB Solid State Memory	capacities as required		
	Operating System	Linux - CentOS 7	Other OS options available on request		
Dataflow	FPGA	Xilinx Kintex Ultrascale Plus series - choose from - KU5P (217K LUTs, 544 BRAMs, 1,824 DSPs) - KU3P (163K LUTs, 408 BRAMs, 1,368 DSPs) - KU3P (299K LUTs, 680 BRAMs, 2,928 DSPs)	KU5P fitted as standard		
Engine	Memory	1x 16GB DDR4-2400 SODIMM	Or 8GB or 32GB		

Purdue, IU, MIT, Harvard, Boston, NEU, Dartmouth, U of Massachusetts at Amherst, USC, UCLA, Columbia, NYU, Princeton, NJIT, CMU, Temple U, UIUC, Michigan, Wisconsin, Minnesota, FAU, FIU, Miami, Central Florida, U of Alabama, U of Kentucky, GeorgiaTech, Ohio State, Imperial, King's, Manchester, Huddersfield, Cambridge, Oxford, Dublin, Cork, Cardiff, Edinburgh, EPFL, ETH, TUWIEN, UNIWIE, Graz, Linz, Karlsruhe, Stuttgart, Bonn, Frankfurt, Heidelberg, Aachen, Darmstadt, Dortmund, KTH, Uppsala, Karlskrona, Karlstad, Napoli, Salerno, Siena, Pisa, Barcelona, Madrid, Valencia, Oviedo, Ankara, Bogazici, Koc, Istanbul, Technion, Haifa, BerSheba, Eilat, Belgrade, Podgorica, Koper, Ljubljana, Maribor, Nova Gorica, etc. Also at the World Bank in Washington DC, IMF, the Telenor Bank of Norway, the Raiffeisen Bank of Austria, Brookhaven National Laboratory, Lawrence Livermore National Laboratory, IBM TJ Watson, HP Encore Labs, Intel Oregon, Qualcomm VP, NCR, RCA, Fairchild, Honeywell, Yahoo NY, Google CA, Microsoft, Finsoft, ABB Zurich, Oracle Zurich, and many other industrial labs, as well as at Tsinghua University, Shandong, NIS of Singapore, NTU of Singapore, Tokyo, Sendai, Seoul, Pusan, Sydney University of Technology, University of Sydney, Hobart,

Auckland, Toronto, Montreal, Durango, MontereyTech, Cuernavaca, UNAM...

MAXELER
Technologies
MAXIMUM PERFORMANCE COMPUTING



MECO 2019

8th Mediterranean Conference on Embedded Computing

Including:

7th EUROMICRO/IEEE Workshop on Embedded and Cyber-Physical Systems (ECYPS 2019)

Montenegro, Budva, 10th-14th June 2019

