# 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 communications/memory

- B. Ilya Prigogine: Impact of energy, entropy, order, and optimization
- C. Daniel Kahneman: Impact of approximate computing on precision
- D. Andre Geim: Impact of latency on precision

### 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. BigData

Is the von Neumann Paradigm still the most effective one?

A. MultiCores?B. ManyCores?

### The Holy Trinity of Generalized Computing



-Size -Power -Speedup -Precision

### Architecture

# 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  $t-\infty$ ?

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

	Today	2018-20
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 traided!

# 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



9/60

• 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



2n+3

### 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 Maxeler compiler to create a maximally effective execution graph.

Final goal: The execution graph with the minimal length of edges.

### MaxCompiler



### 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

Protected by two aSoG protection levels and two Maxeler protection levels!

### **Publications of Interest for NanoAcceleration**

- 1. Milutinovic, V., Mapping of Neural Networks on the Honeycomb Architectures, Proceedings of the IEEE, Vol. 77, No 12, December 1989, pp. 1875-1878.
- 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., The Split Temporal/Spatial Cache: Initial Performance Analysis, Proceedings of the SCIzzL-5, Santa Clara, California, USA, March 26, 1996, pp 72-78.
- 4. 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, IEEE Transactions on Communications, Vol. COM-36, No. 5, May 1988, pp. 538-543.
- 6. 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), Communications of the ACM (nano-acceleration), May 2013.

- Helbig, W., Milutinovic, V., The RCA's DCFL E/D MESFET GaAs 32-bit Experimental RISC Machine, IEEE Transactions on Computers, Vol. 36, No. 2, February 1989, pp. 263-274.
- 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



MyMaxeler

Back ....

### Moving from Petaflops to Petadata

May 5, 2013

VIEWPOINT: Moving from Petaflops to Petadata

M. Flynn<sup>‡II</sup>, O. Mencer<sup>‡</sup><sup>‡</sup>, V. Milutinovic<sup>†</sup>, G. Rakocevic<sup>§</sup>, P. Stenstrom<sup>8</sup>, R. Trobec<sup>b</sup>, M. Valero<sup>F</sup>

<sup>‡</sup>Maxeler Technologies, <sup>II</sup>Stanford University, <sup>‡</sup> Imperial College London, <sup>†</sup>University of Belgrade, <sup>§</sup>Mathematical Institute of the Serbian Academy of Sciences and Arts in Belgrade, <sup>8</sup> Chalmers University of Technology, <sup>§</sup>Jožef Stefan Institute, <sup>F</sup>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

Search

### 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.<sup>[1]</sup>

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

#### 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 2.
- 2. ^ "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 home page ₽.
- ACM Computing Surveys
   in ACM Digital Library.
- ACM Computing Surveys ₽ in DBLP.



#### ACM Computing Surveys

Q

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) 1557-7341 & (web)		
Links			
• Journal homepage 🗗			
<ul> <li>Online access 🗗</li> </ul>			
<ul> <li>Online archive G</li> </ul>			

Article Talk

### 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 and precision 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.

Easily doable in DataFlow, difficult to do in ControlFlow.

### Delayed Decision for Better Precision: Geim

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

### Maxeler @ BigDataAnalytics



20 [Size] 20 [Power] 20, 200, 2000 [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 DEEs to CPU servers with zero-copy RDMA access



The Networking Appliance Intel Xeon CPUs and 4 DFFs with direct links to up to twelve 40Gbit **Ethernet** connections







### **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



### J.P.Morgan

### Innovation in Investment Banking Technology 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)

28/60

- 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 emised excellellem
- 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

Technology ignites our business. Be the spark.

#### **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

#### 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 very simple sample:

f(x) = 2x + x

Moving from a single calculation to a fine grained parallelism

#### 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.

### J.P.Morgan

#### 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

0 0

#### 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.

CME Ticker Bloomberg Ticker	DSF Pricing				-		
	Price	Coupon	PV01	NPV	Implied Rate	Timestamp	
T1UM4 2Y	CTPM4	100'057	0.750%	<mark>\$</mark> 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 <mark>%</mark>	4:00:03 PM CT 4/4/2014
N1UM4 10Y	CNPM4	100'225	3.000%	\$90.16	\$703.12	2.9220 <mark>%</mark>	4:00:03 PM CT 4/4/2014
B1UM4 30Y	CBPM4	<mark>102'270</mark>	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%	<mark>\$</mark> 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

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

Quotes and analytics are updated every 15 minutes.

#### (K) 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



### 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<sup>†</sup>, J. Stefani<sup>†</sup>, W. Liu<sup>†</sup>, R. Dimond<sup>‡</sup>, O. Pell<sup>‡</sup>, R.Ergas<sup>§</sup> <sup>†</sup>Chevron, <sup>‡</sup>Maxeler, <sup>§</sup>Formerly Chevron, SEG 2008

### **Global Weather Simulation**



global atmospheric equations through mixed-precision data flow engine, FPL2013]

London

Imperial College

### Weather Model – Performance Gain

Platform	Performance	Speedup
	()	
6-core CPU	4.66K	1
Tianhe-1A node	110.38K	23x
MaxWorkstation	468.1K	<b>100x</b>
MaxNode	1.54M	330x

Meshsize:  $1024 \times 1024 \times 6$ 14xMaxNode speedup over Tianhe node: 14 times





Imperial College



### Weather Model -- Power Efficiency

Platform	Efficiency	Speedup
6-core CPU	20.71	1
Tianhe-1A node	306.6	14.8x
MaxWorkstation	2.52K	<b>121.6</b> x
MaxNode	3К	144.9x

Meshsize:  $1024 \times 1024 \times 6$ MaxNode is 9 times more power efficient

Steph Performance, Computing



Imperial College London

scas

34/60

**9** x

### Weather and Climate Models



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

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





### Maxeler Running Smith Waterman

Ŭ,	Smith Waterman Demo - Maxeler Technologies
GCA AGA GAT AA Ala Arg Asp Asp	Query : UniRef50_F2T2I7 Histone-lysine N-methyl transferase n=8 Tax=E(1280)Best scores : spl01DR06 SET1_COCIM Histone-lysine N-methyl transferase, H3(1271)4077spl02UMH3 SET1_ASPOR Histone-lysine N-methyl transferase, H3(1229)3849spl04WNH8 SET1_ASPFU Histone-lysine N-methyl transferase, H3(1220)3683spl04WNH8 SET1_MENL Histone-lysine N-methyl transferase, H3(1220)3683spl04WNH8 SET1_GIBZE Histone-lysine N-methyl transferase, H3(1220)3683spl04WNH8 SET1_MEUCR Histone-lysine N-methyl transferase, H3(1252)2150spl04UFSR3 SET1_GIBZE Histone-lysine N-methyl transferase, H3(1076)2089spl04UFSR3 SET1_DEBHA Histone-lysine N-methyl transferase, H3(1076)2089spl06GKL7 SET1_DEBHA Histone-lysine N-methyl transferase, H3(1076)2089spl06CEK8 SET1_VARLI Histone-lysine N-methyl transferase, H3(1070)938spl05ABG1 SET1_CANAL Histone-lysine N-methyl transferase, H3(1040)893
1     2     3     4       uniprot_sprot_fasta       Number of sequences : 532224       Number of residues : 188726448	5 Best alignment : MSRASAGFADFFPTAPSVLQKKRSSKAAQDRPKGKLKHDDDPQSSNPAPTAATAAVTVTGVGVPGAEEGGASDNNTNSDV MSRAPAGFADFFPTAPSVLQKKRS-KAAQDR-HAANTPKAADPLPNLGLSS-TPDIK-GGVGTSAD-
Scoring matrix :	HNNINSNNNNKNNSSSHTNINSNTOFDESAGAVARGDVNITPGDANGVGSSSSTSTGSS-VFSASILPQPGLTTSNGITH -NPVRAVGE-RSAE-TT-LALGDTNG-ATSSSSLSTGSSGFFSASA-P-PGVAKPNGISS PHALTPLTNTDSSPSCKIASPSGQKS-IA-ATGEIVPTSRFVDDIK-ATITPLQTPPTPRIQARPAGNAPKGYKLTYDPD C-ALTPLTNTDSSPPCKIESPLGSKSGSTDAAPQLAPTCEAHGGPEPVTITPLHTPPTPRVQARPANSEVKGHKITYDPD
BLOSUM62	LERK-PLTKEKRRKPQYEVFDTTED-EAPPADPRIAIANYTRGAGCKQKTKYRPAPYILRPWPYDPATSVGPGPPTQIVV LDRKFP-SKARRRKPQYETFGVDDEKDPPPCDPRMAIANYTRGAACKQKTKYRPTPYILRPWAYDPTTSVGPGPPTQIVV TGYDPLTPLAPISALFSSFGDIAEIKNRTDPNTGRFLGVCSIRYKDSRMFRGGGPLLAAQAARRAYLECKKEQRIGVRRI TGFDPLTPIAAISALFSSFGDIGEINNRTDPMTGRFLGVCSIKYKDSRAFRGGISLSASQAVRRAYLECKKEQRIGTRRI
Open Gap Penalty	-5 QVSLDRDGVVSDRLVARIIGSQR-QQEP-PPLVME-E-KMKSE-EQDNLPPPTAPKGPS-RXPNM RVELDRNGVVSGRMVAKLITAQKAEFPSLEESRKESVGDNDNRLPIGDGAKKDNEQSKDNLPPSTAPKGPSGRSSLHPSL LIPEGPRATMMKPPAPSLIEETPILDQIKRDPYIFIAHCYVPVLSTTIPHLERRLKLFNWKSVRCDKTGYYIIFDNSRRG LAPDGPRA-VLKSPVPSRIEETPILQQIKRDPYIFIAHCYVPVLSTTVPHLERRLKLYDWKAVRCDKTGYYIIFENSRRG
Stop         Comp           Stopping computation - please         V	vait Performance: 812.0759 GCUPS







b

#### → ☆ 自 🕹 兪 😕



## appgallery.maxeler.com

38/60

http://www.mi.sanu.ac.rs/~appgallery.maxeler/















44/60

MAXELER Technologies

### ManyCore

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



What is better, real and anecdotic?





### How about 2 000 000 ants?





















Marmelade



### 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



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

## An Original Book Covering the Essence

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

□ <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



#### 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



4 COMMENTS COMMENT NOW 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.)

CLOUD **OPERATIONS** PLATFORM™ FOR AWS

Powering automated

### newelectronics

### Intel says logic is faster than GPUs



### Intel's Programmable Systems Group takes its first step towards an 28 June 2016 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.



## 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, 20-200b) monthly electricity bills, reduced 20 timesc) size, 20 times smaller

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



DataFlow



# Why are Electricity Bills so Small?

• Factor: 20

MultiCore/ManyCore

DataFlow

# MMM



 $\mathsf{P} = \mathsf{k} \mathsf{f} \mathsf{U}^2$ 

# Why is the Cubic Foot so Small?

• Factor: 20

MultiCore/ManyCore

Data Processing	
Process Control	Data Processing
	Process Control

DataFlow

### **New Challenges of 2017**

# Amazon AWS

# Hitachi