VILNIAUS GEDIMINO TECHNIKOS UNIVERSITETAS

Šiuolaikinės efektyvaus lygiagretinimo su FPGA technologijos

Andrej Bugajev

2018 m. balandžio 13 d.

Programmable Spectrum

Control-flow processors

Dataflow processor
e.g. FPGA

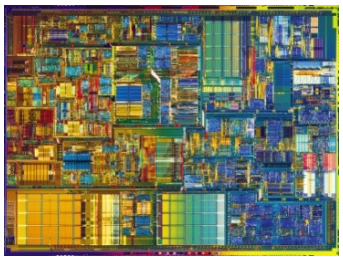

Single-Core CPU

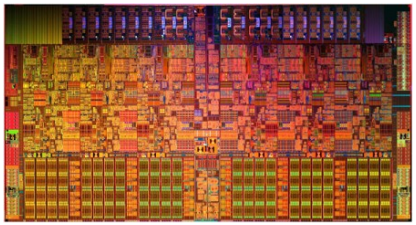

Multi-Core

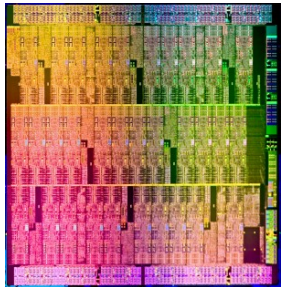

Several-Cores

GK110

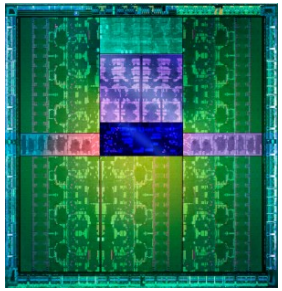

Many-Cores

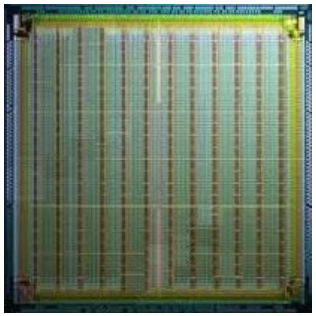

Dataflow

Increasing Parallelism (#cores)

Increasing Core Complexity

Intel, AMD

GPU (NVIDIA, AMD)
Tilera, XMOS etc...

Maxeler

Hybrid e.g. AMD Fusion, IBM Cell

Control-flow processor (CPU)

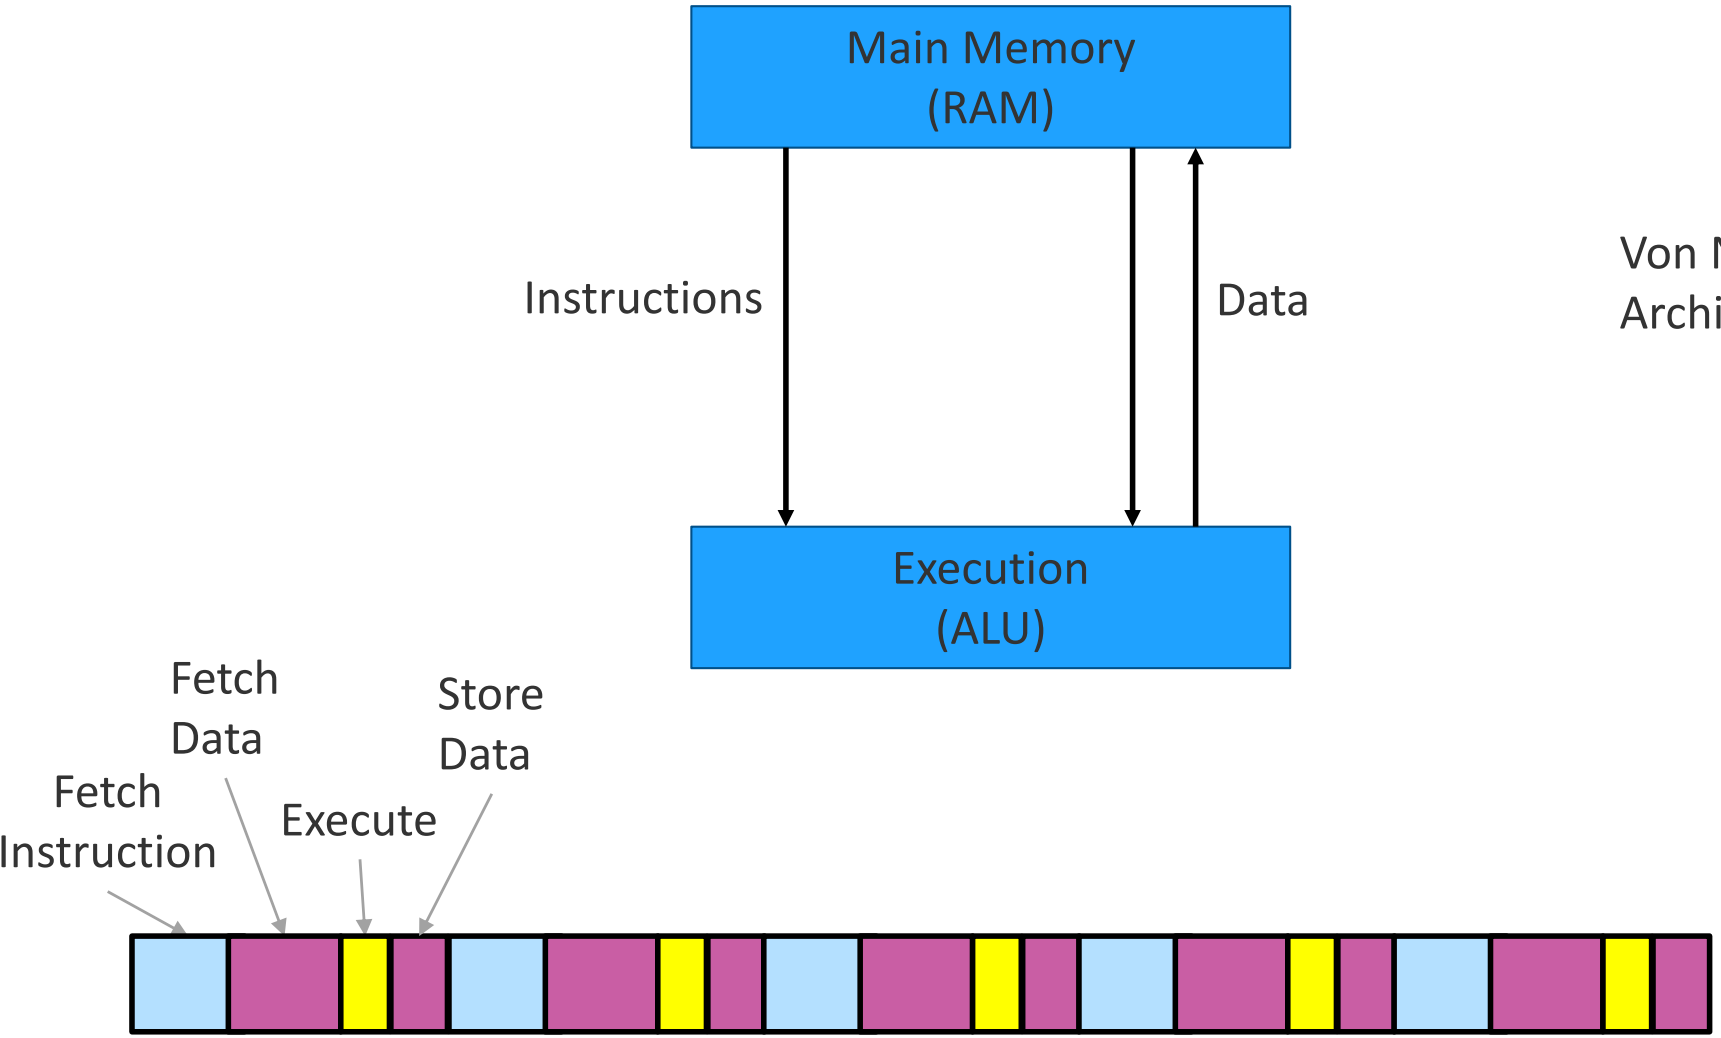

Von Neumann
Architecture

Modern Control-flow processor (CPU)

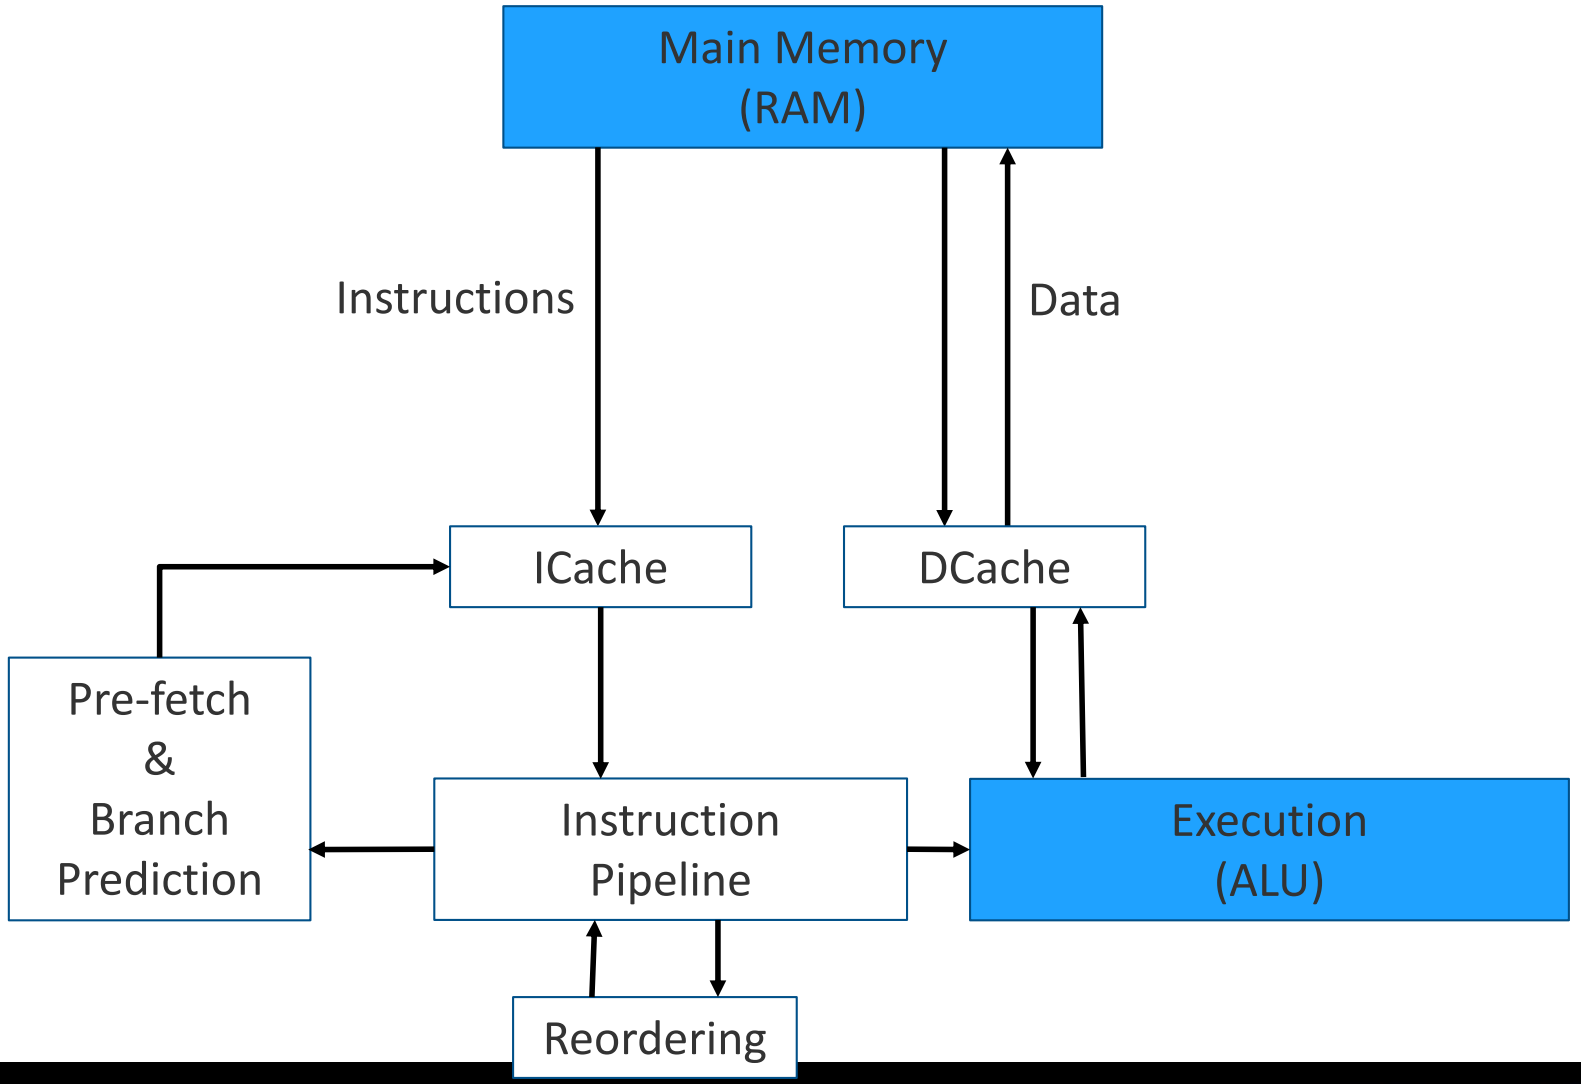

Where silicon is used?

Intel 6-Core X5680 “Westmere”

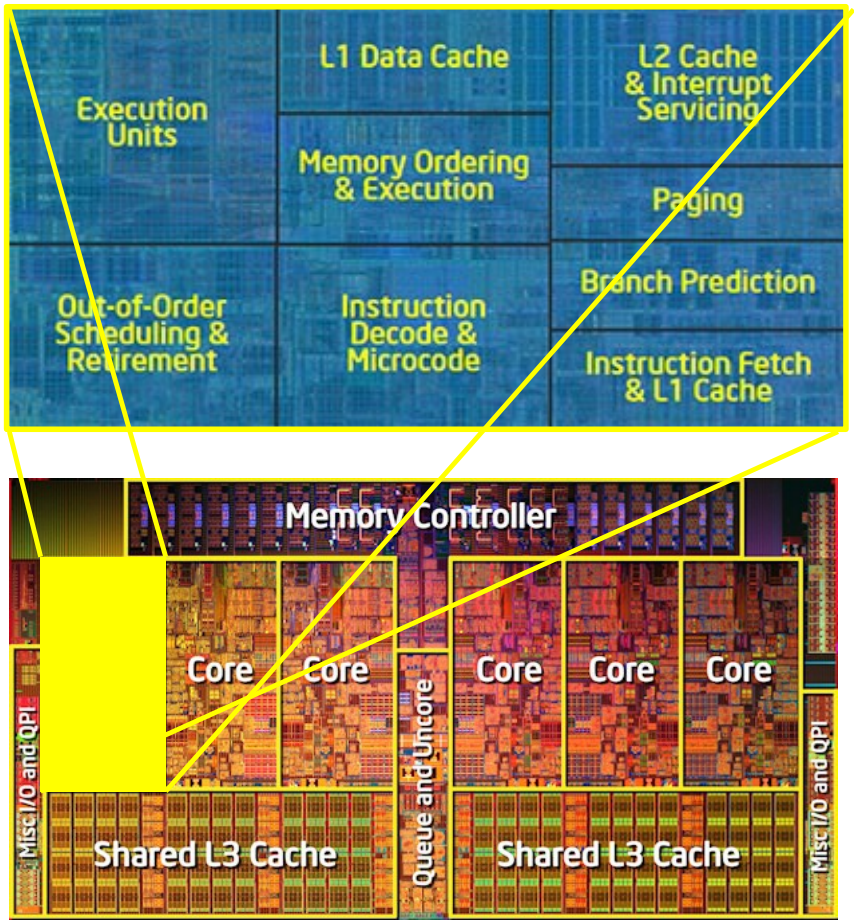

Where silicon is used?

Intel 6-Core X5680 “Westmere”

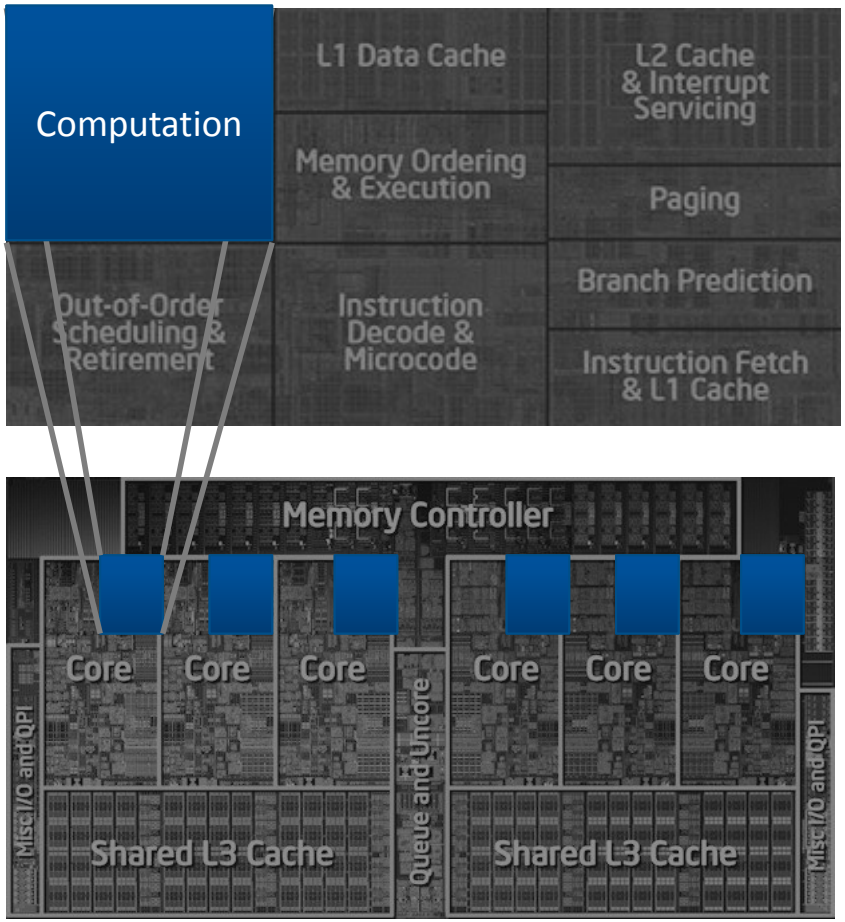

Dataflow Processor
e.g. FPGA

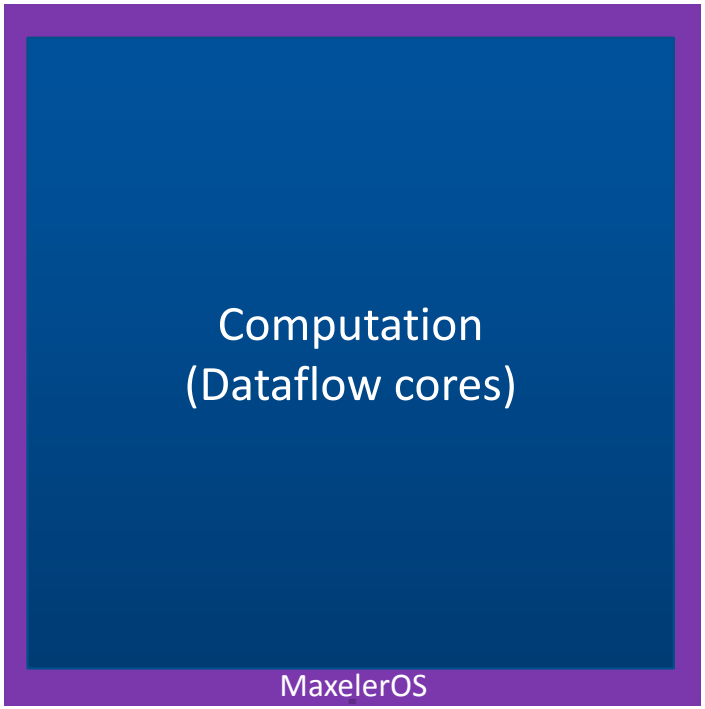

On chip resources

- Each application has a different configuration of dataflow cores
- Dataflow cores are built out of basic operating resources on-chip

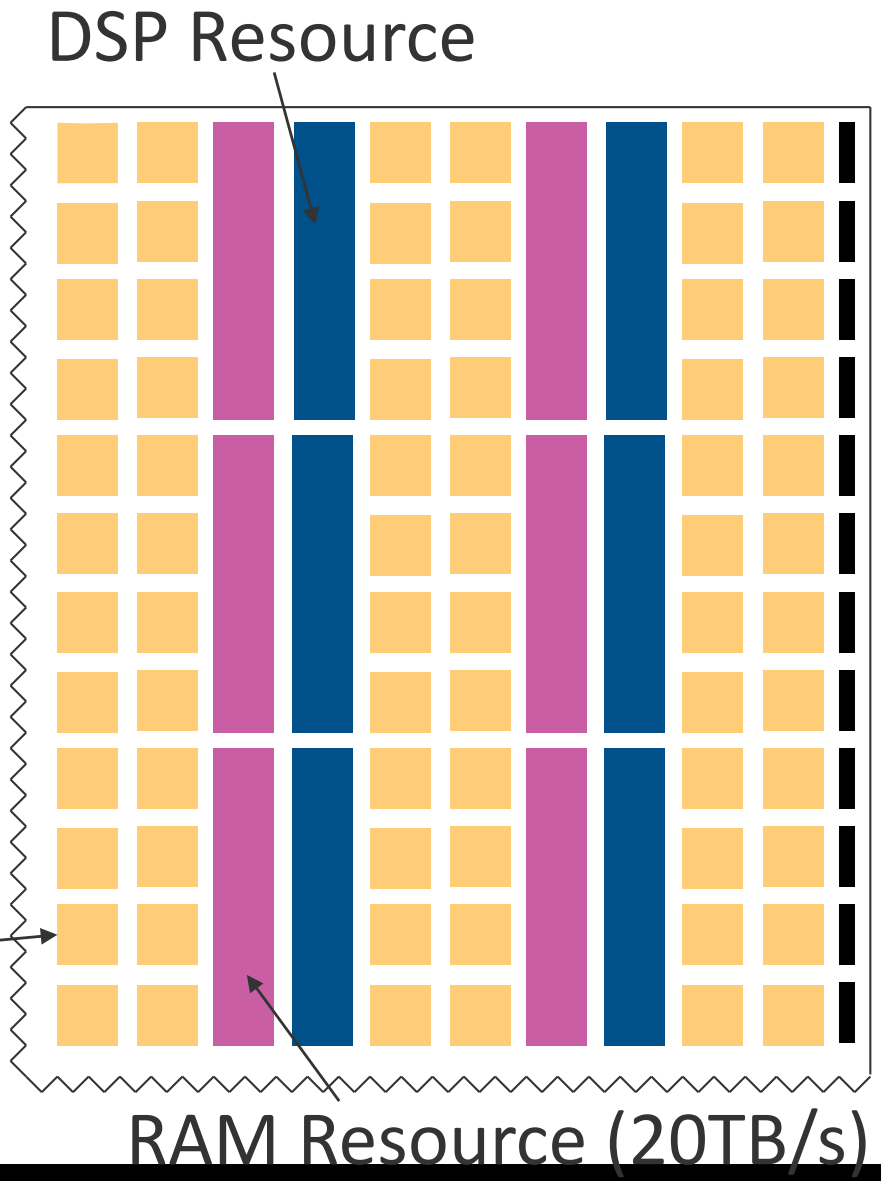

Dataflow Engine (DFE) - ‘Spatial Computing’

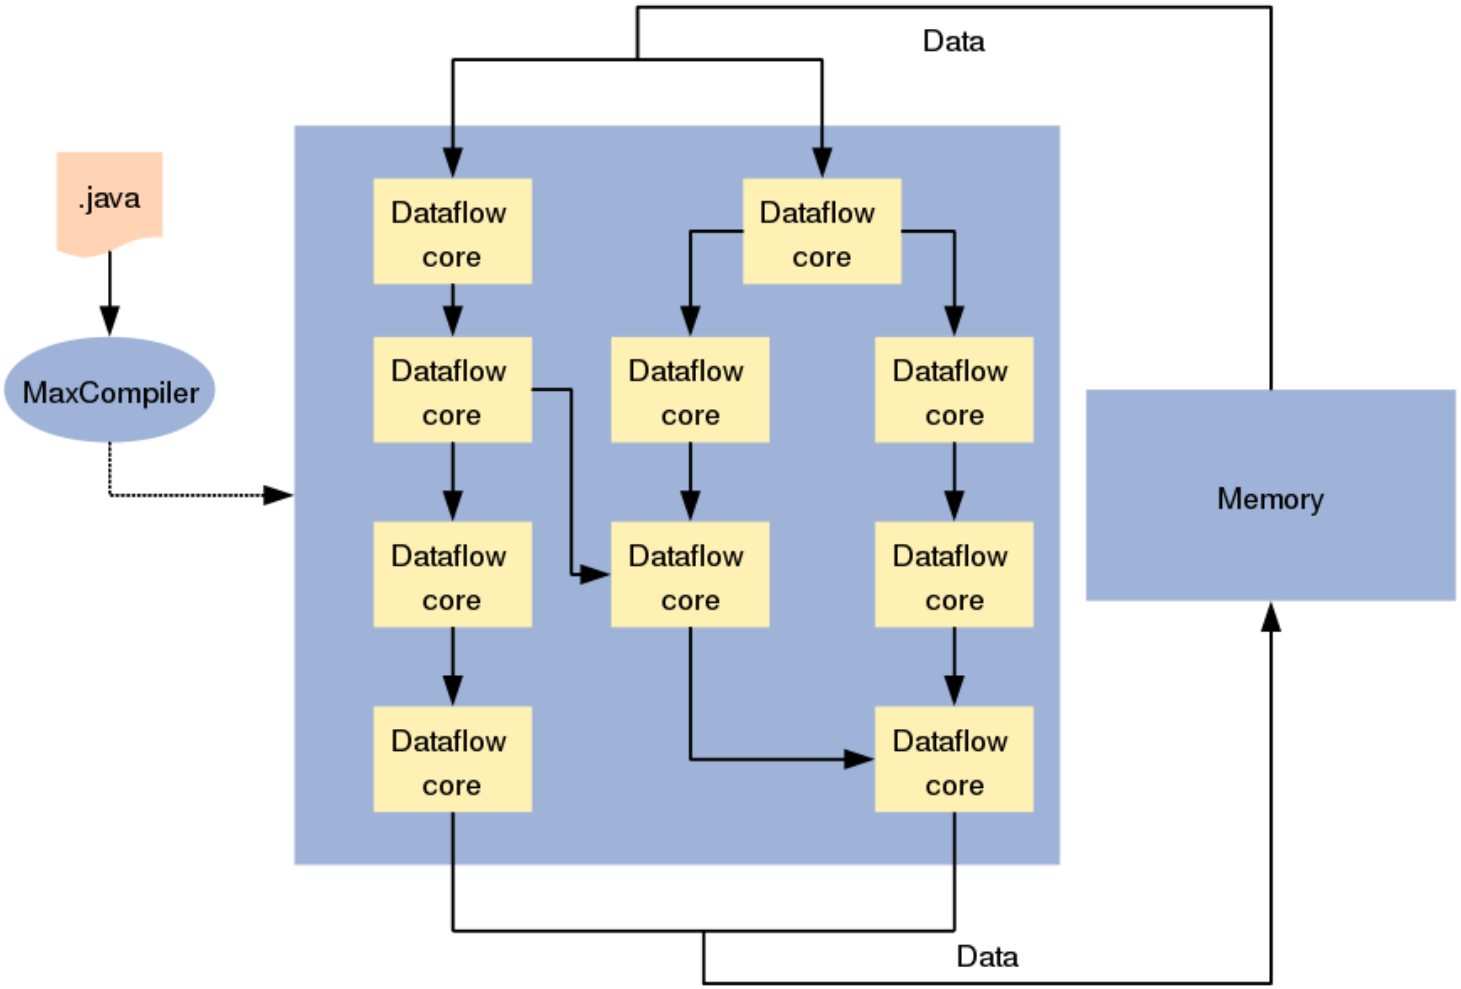

Explaining Control Flow versus Data Flow

Analogy 1: The Ford Production Line

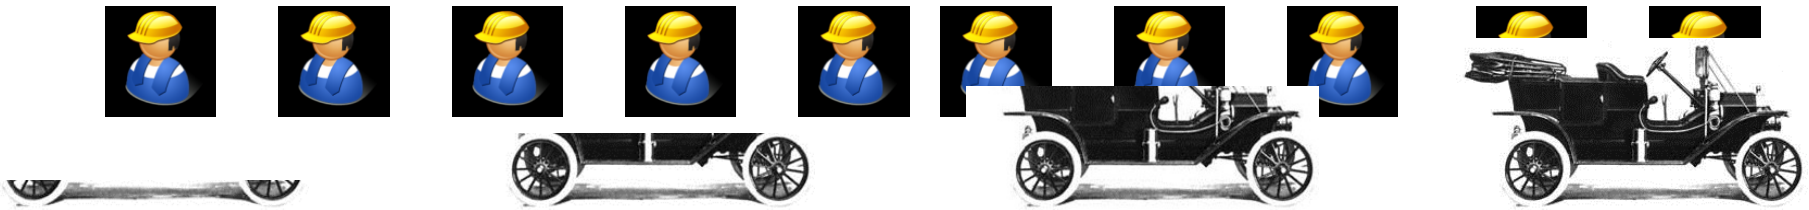

- Experts are expensive and slow (control flow)
- Many specialized workers are more efficient (data flow)

A Dataflow Kernel

```
public class MyKernel extends Kernel {

    public MyKernel (KernelParameters parameters) {
        super(parameters);

        DFEVar x = io.input("x", dfeInt(32));

        DFEVar result = x * x + 30;

        io.output("y", result, dfeInt(32));
    }
}
```

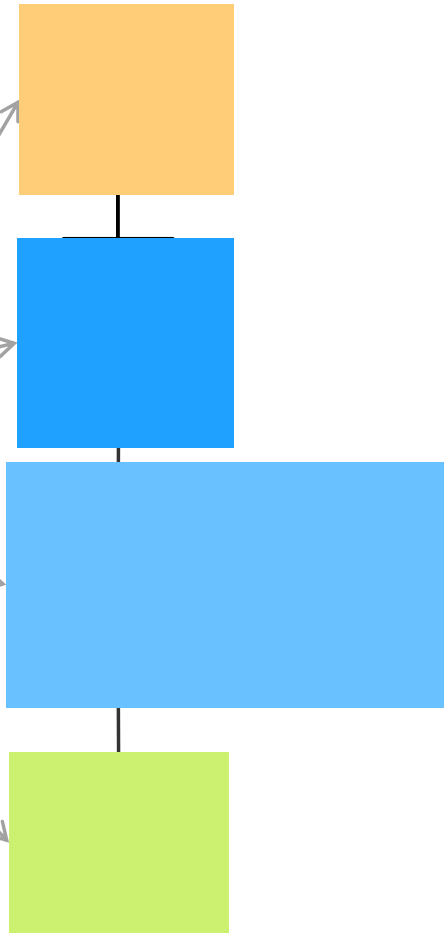

Streaming Data through the Kernel

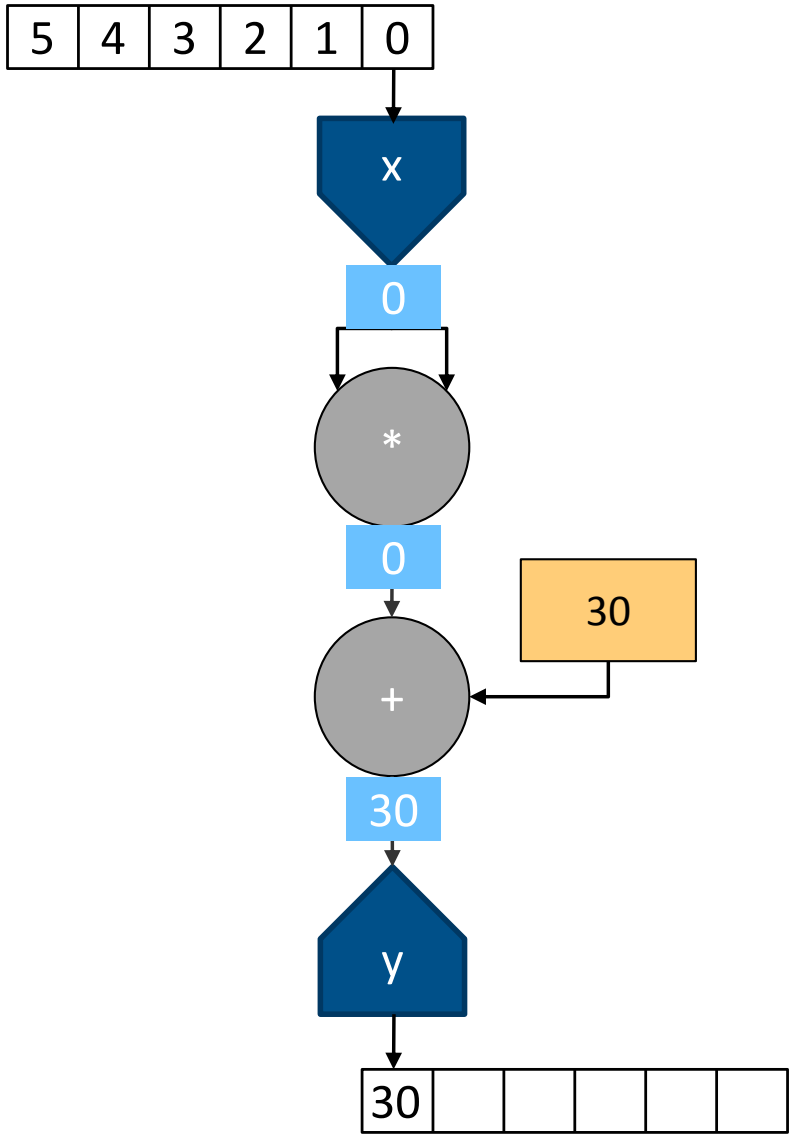

Streaming Data through the Kernel

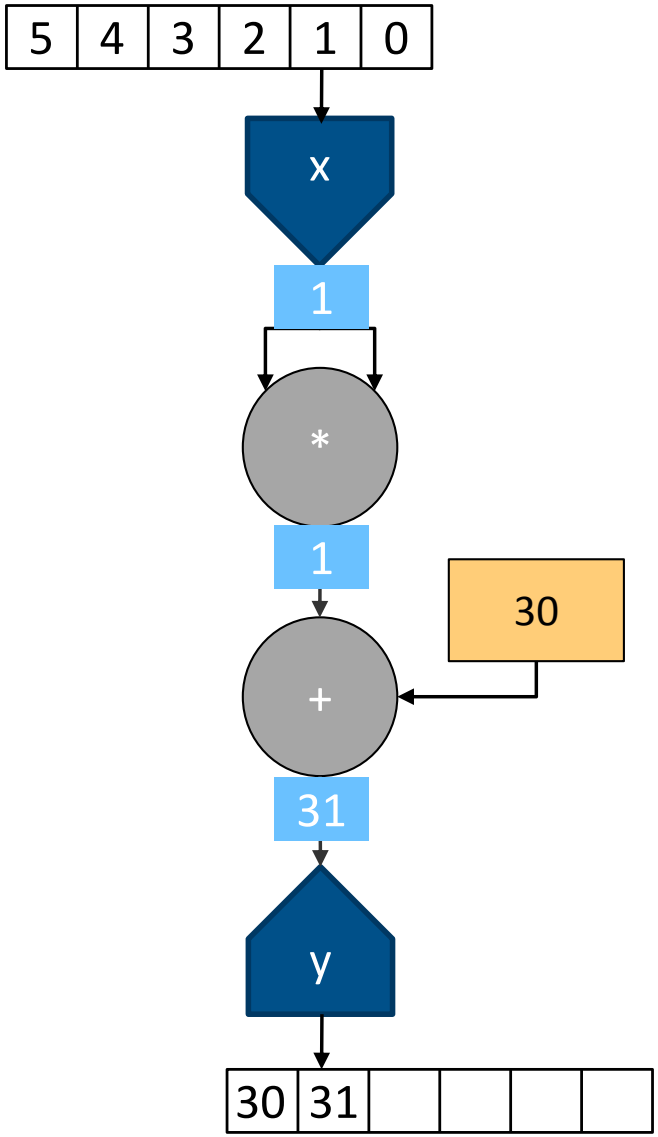

Streaming Data through the Kernel

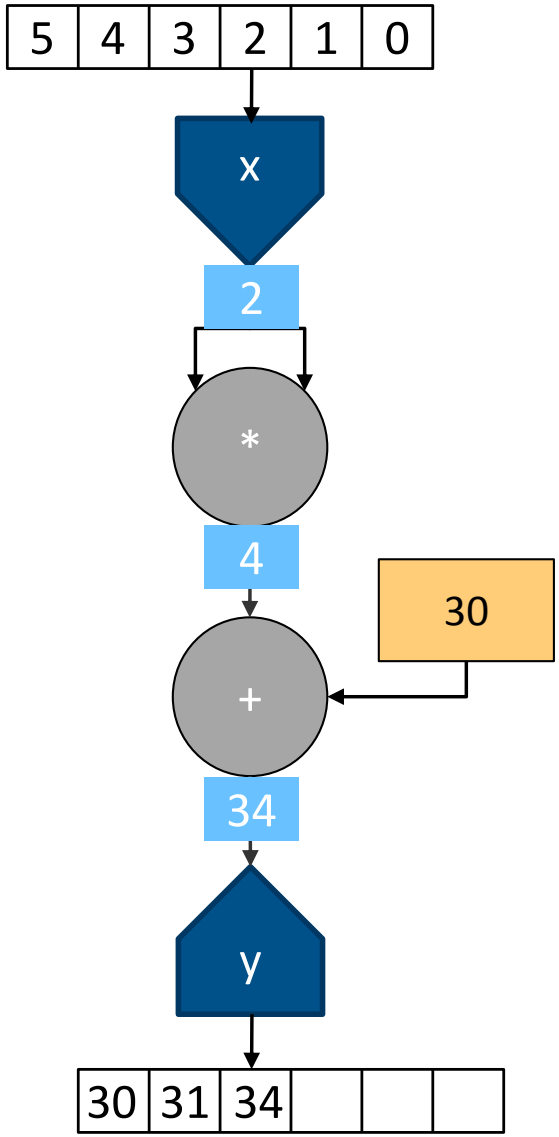

Streaming Data through the Kernel

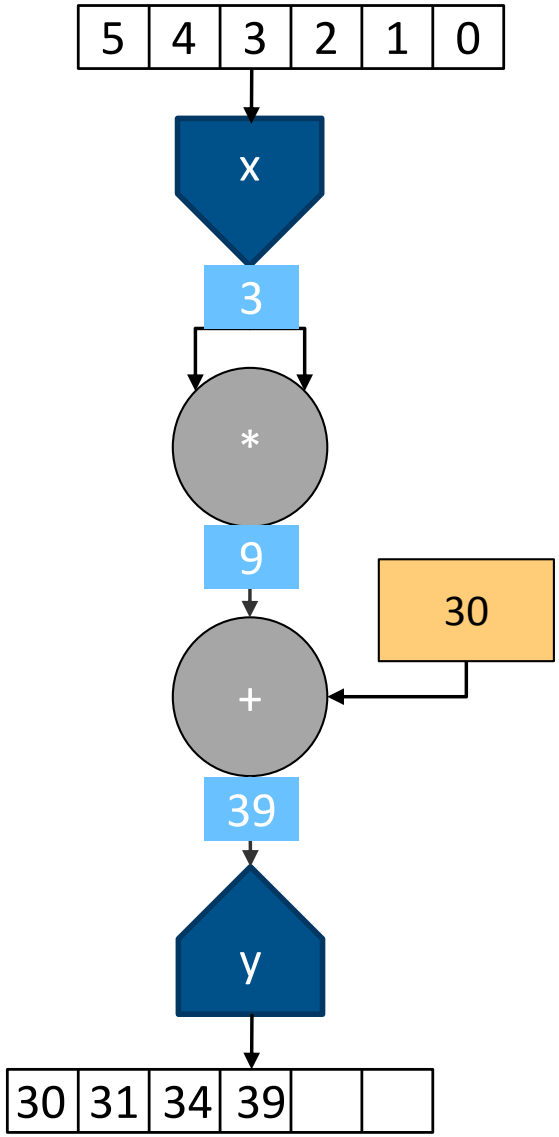

Streaming Data through the Kernel

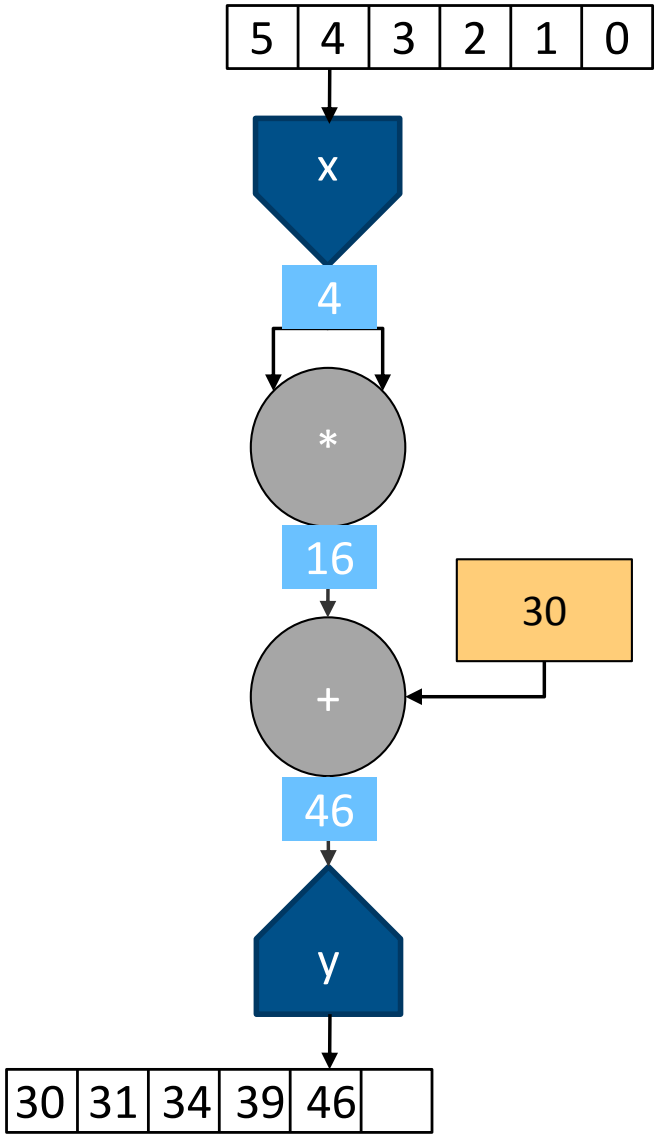

Streaming Data through the Kernel

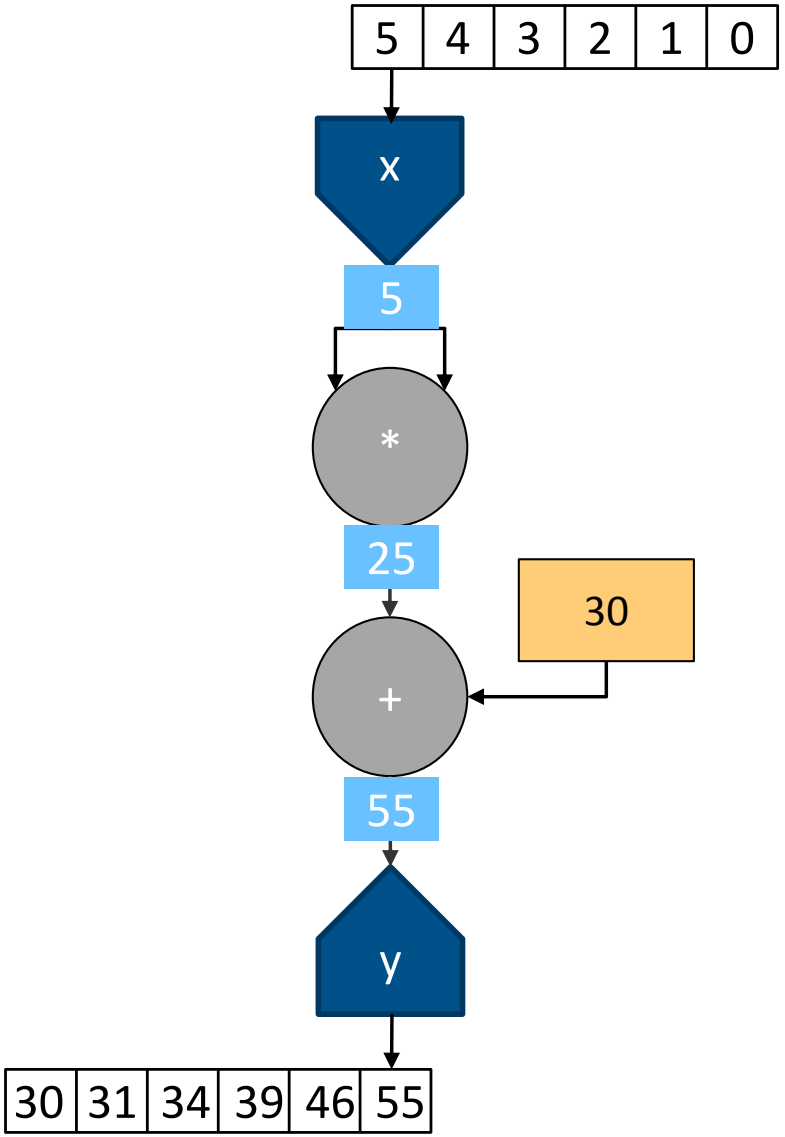

Pipelining Data through the Kernel

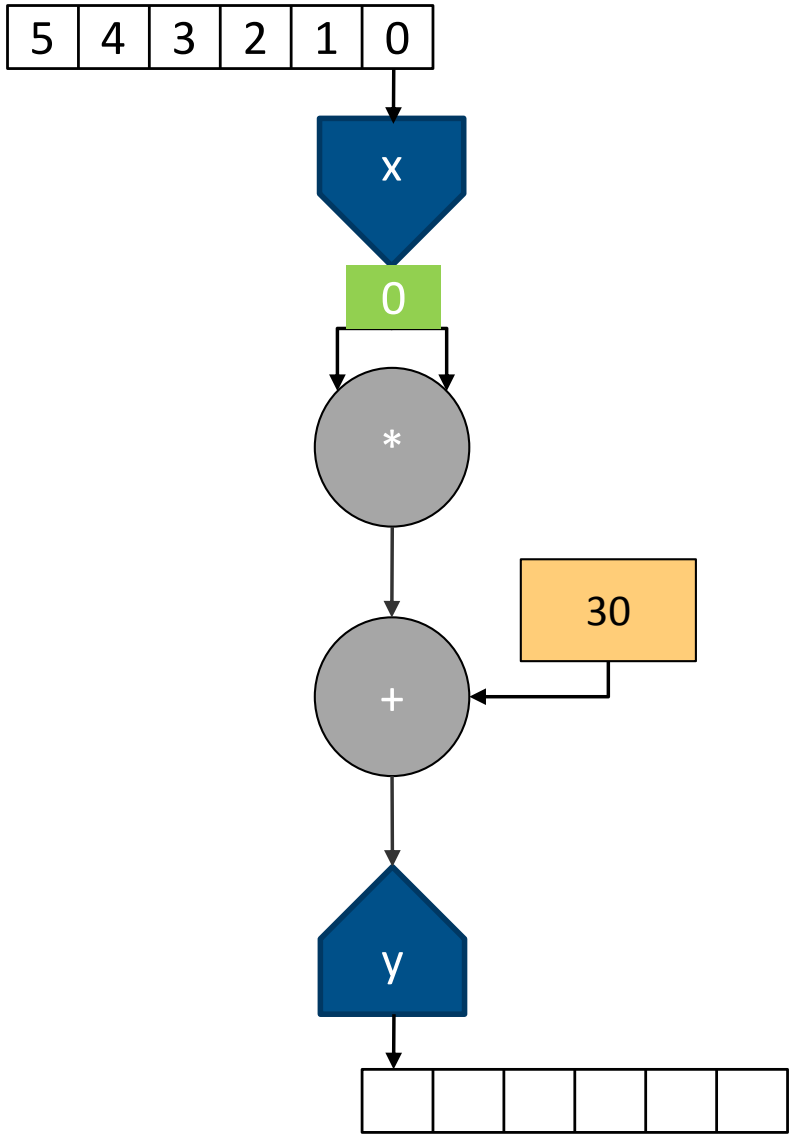

Pipelining Data through the Kernel

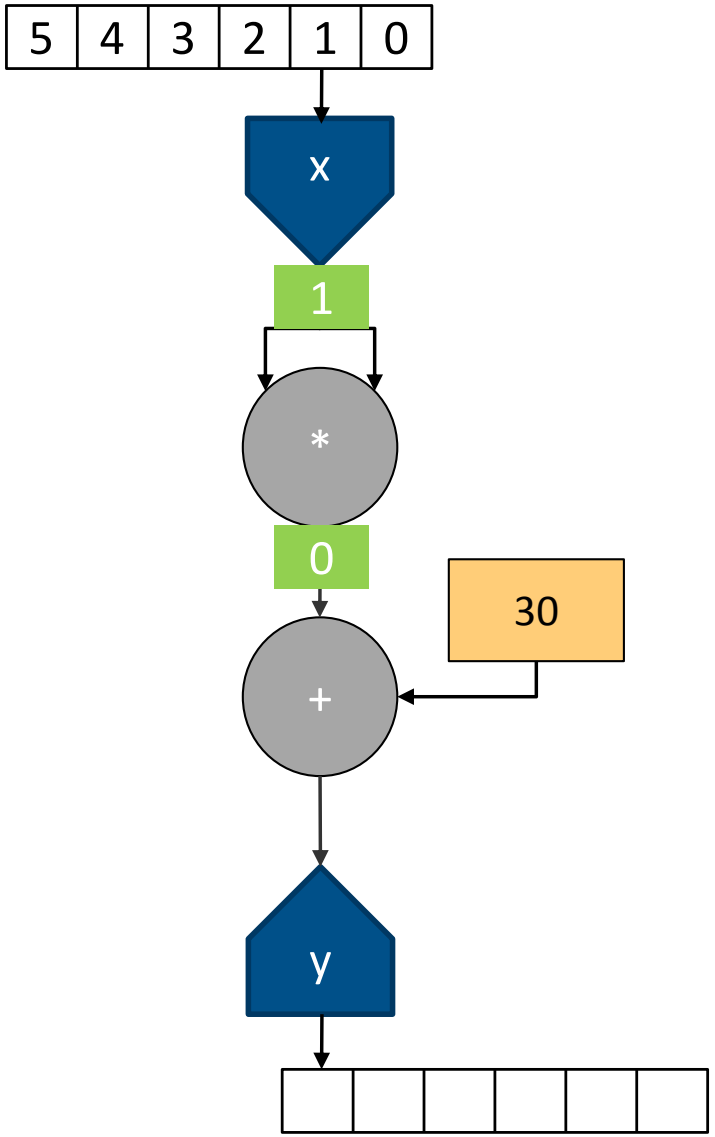

Pipelining Data through the Kernel

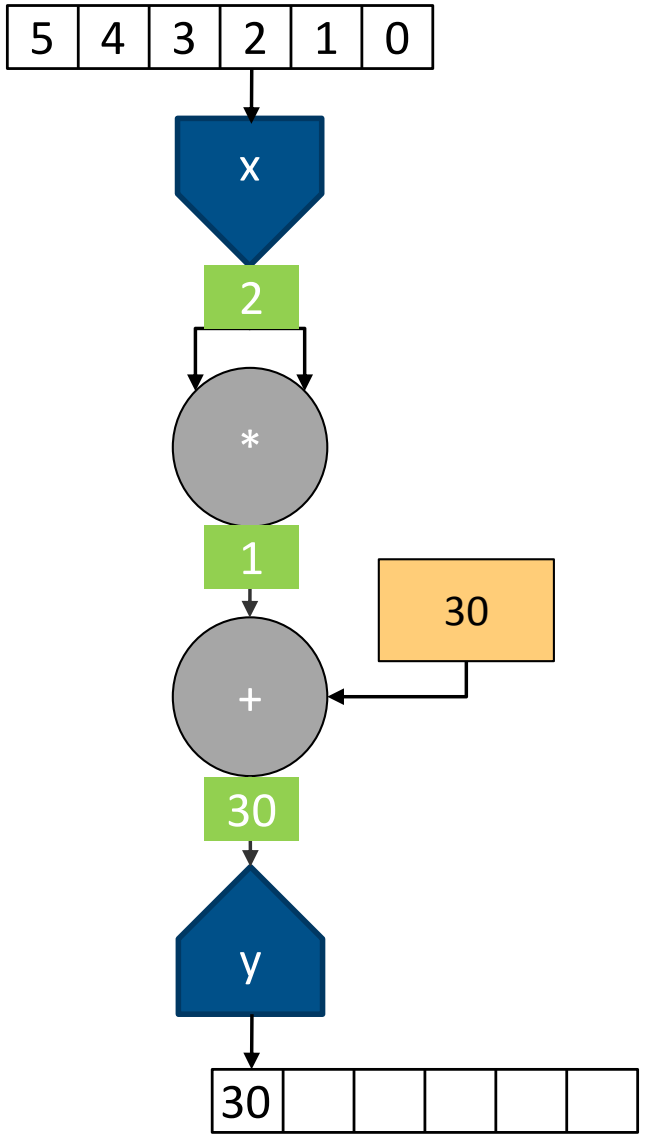

Pipelining Data through the Kernel

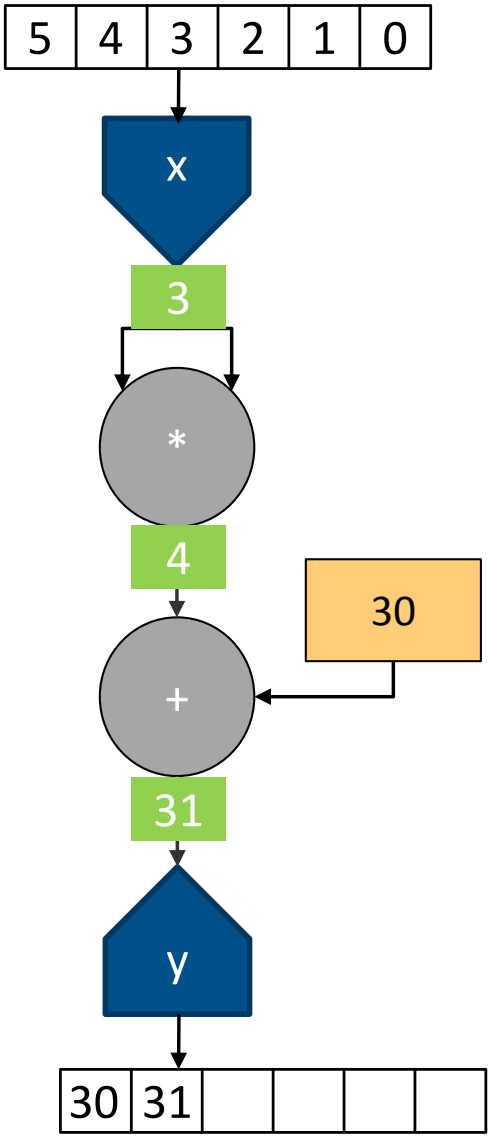

Pipelining Data through the Kernel

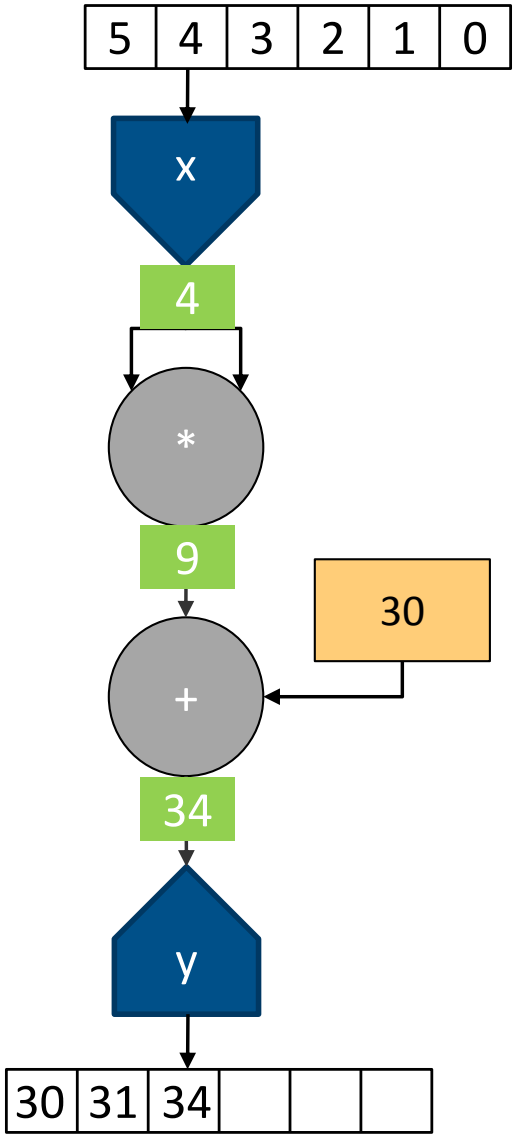

Pipelining Data through the Kernel

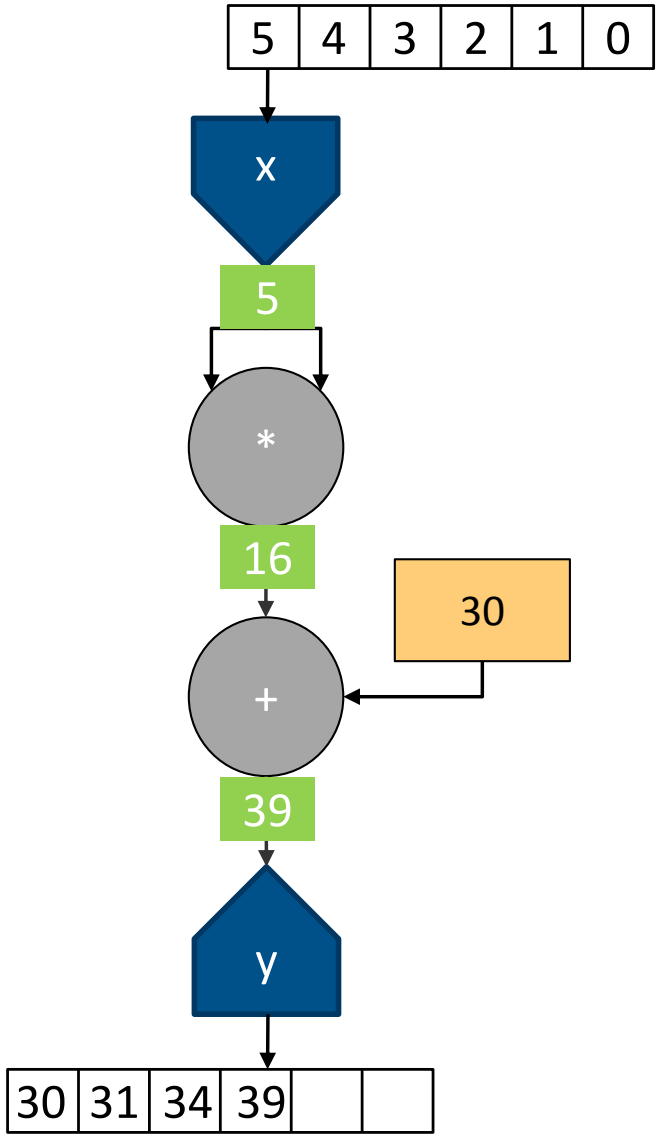

No registers – low latency

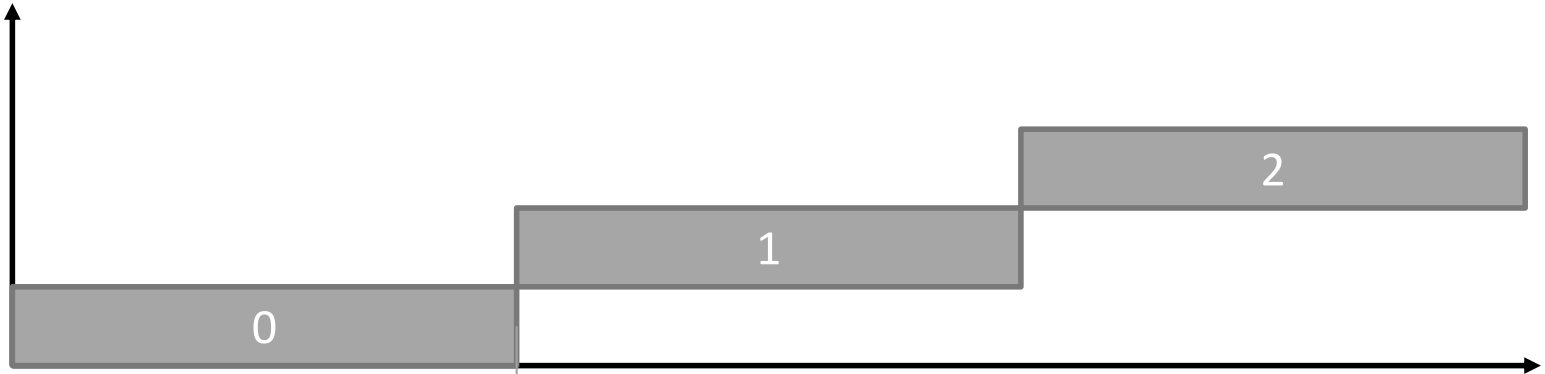

Pipelined – highly parallelised

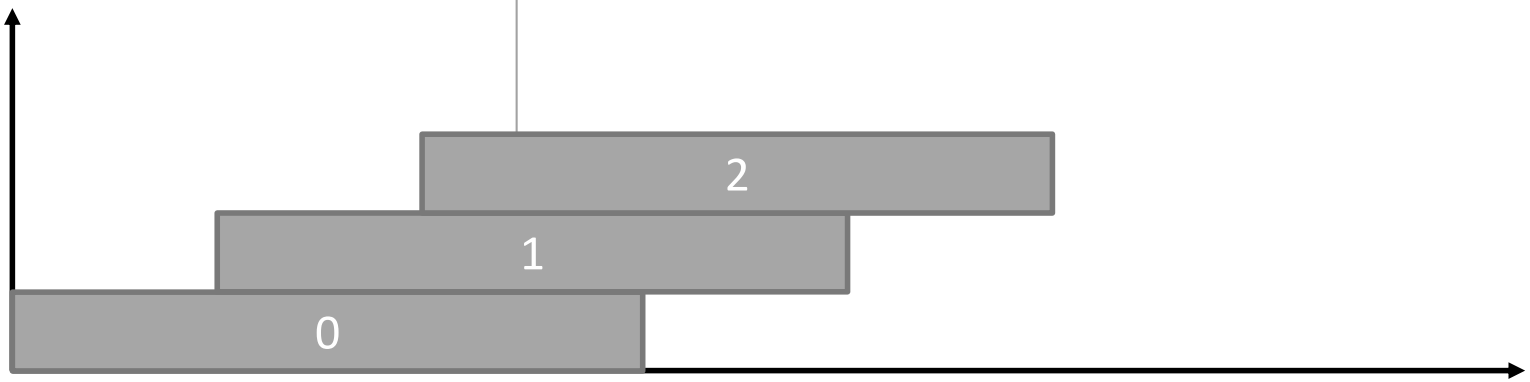

Traditionally FPGA were for specialists

- Can't use FPGA on its own …
  - Need interface to Computer & Data
  - Need interface to local memory
  - Need to design bespoke HW
- FPGAs difficult to program …
  - Specialist languages VHDL, Verilog
  - Need Electronics training & understand FPGAs
  - Simulation only at HW level, modelsim
  - Need also to engineer the interfaces
- FPGAs difficult to use …
  - Need low level drivers to reconfigure and setup

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

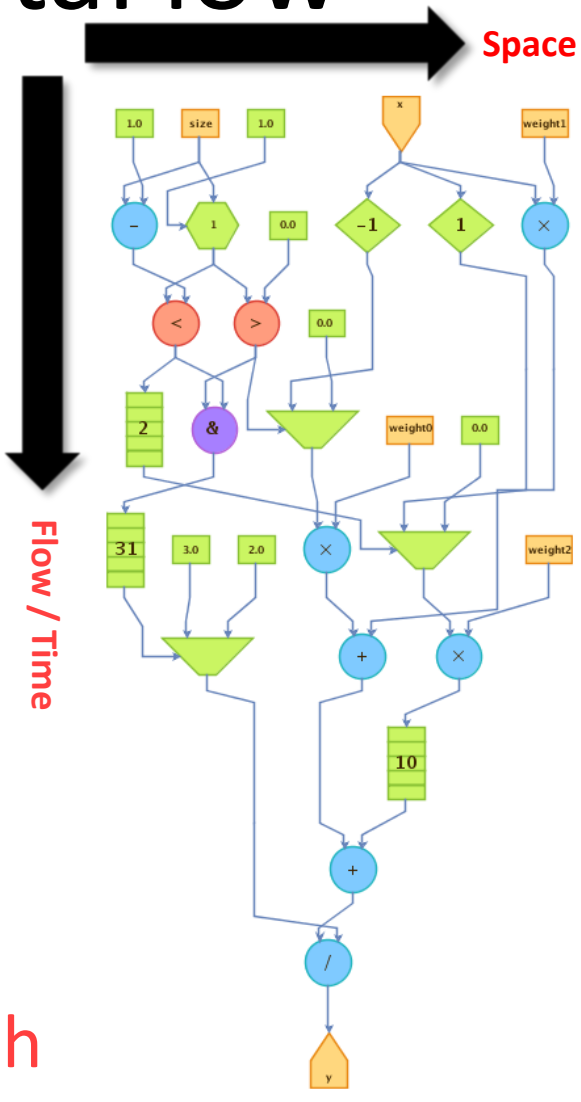

Optimal Solution: Execution Graph

Data flow graph as generated by compiler 4866 nodes

Each node represents an operator in MaxJ code with area time parameters. Each line (edge) represents a DFEVar in MaxJ code.

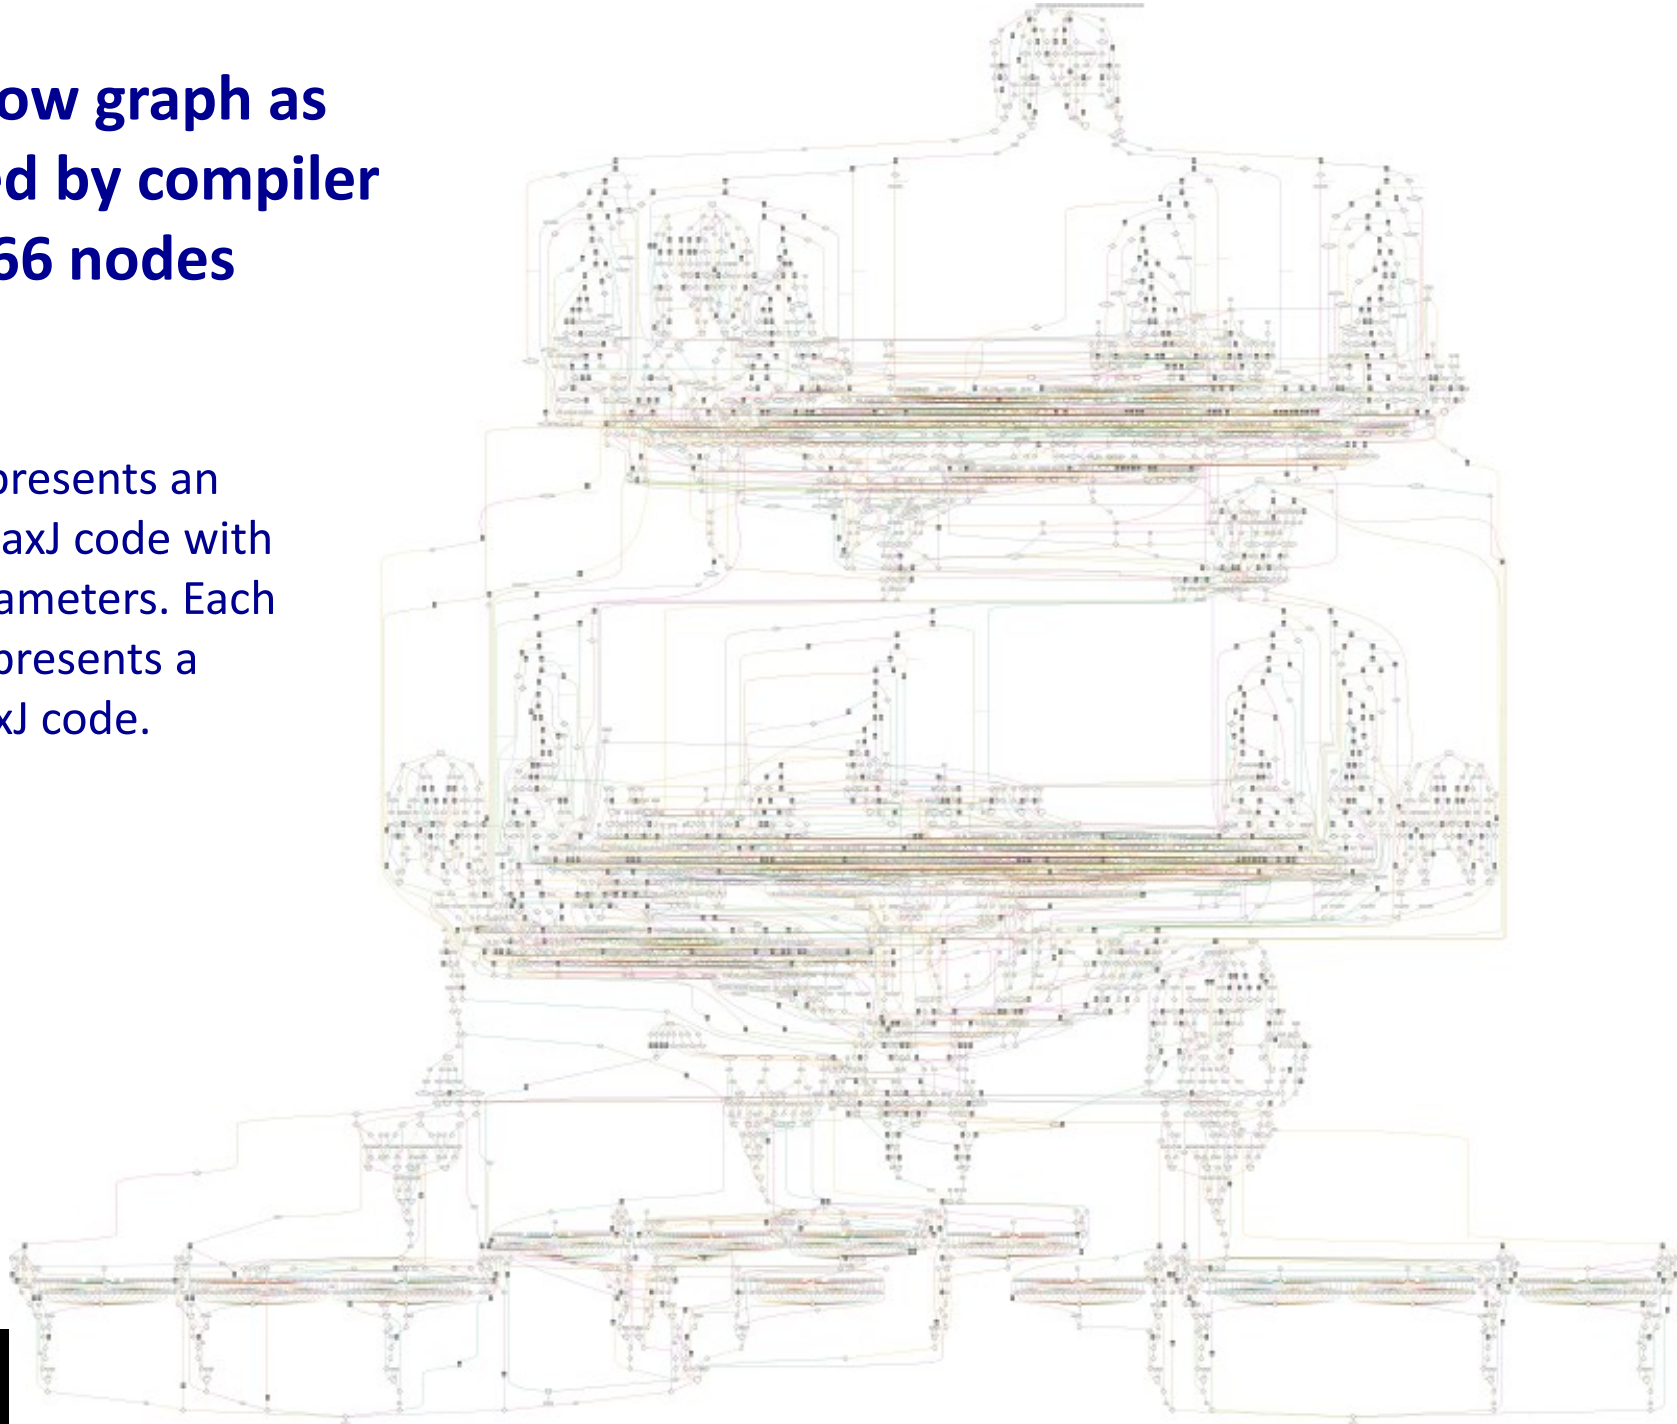

The Maxeler Generic Architecture Application

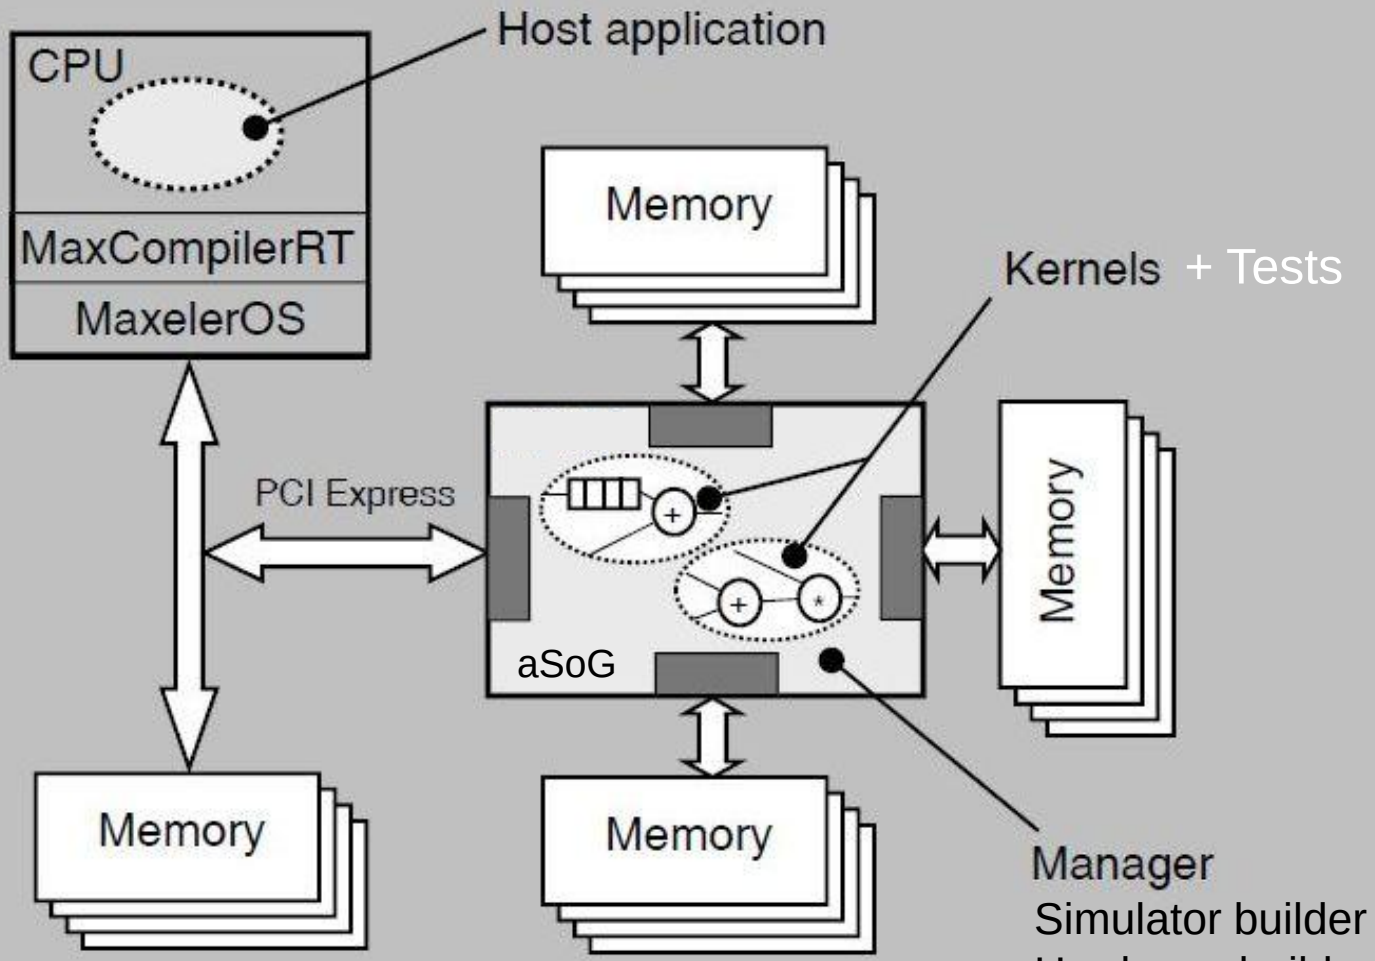

MAXELER Technologies
MAXIMUM PERFORMANCE COMPUTING

Important: Supporting any CL and any OS!

2n+3

MaxCompiler

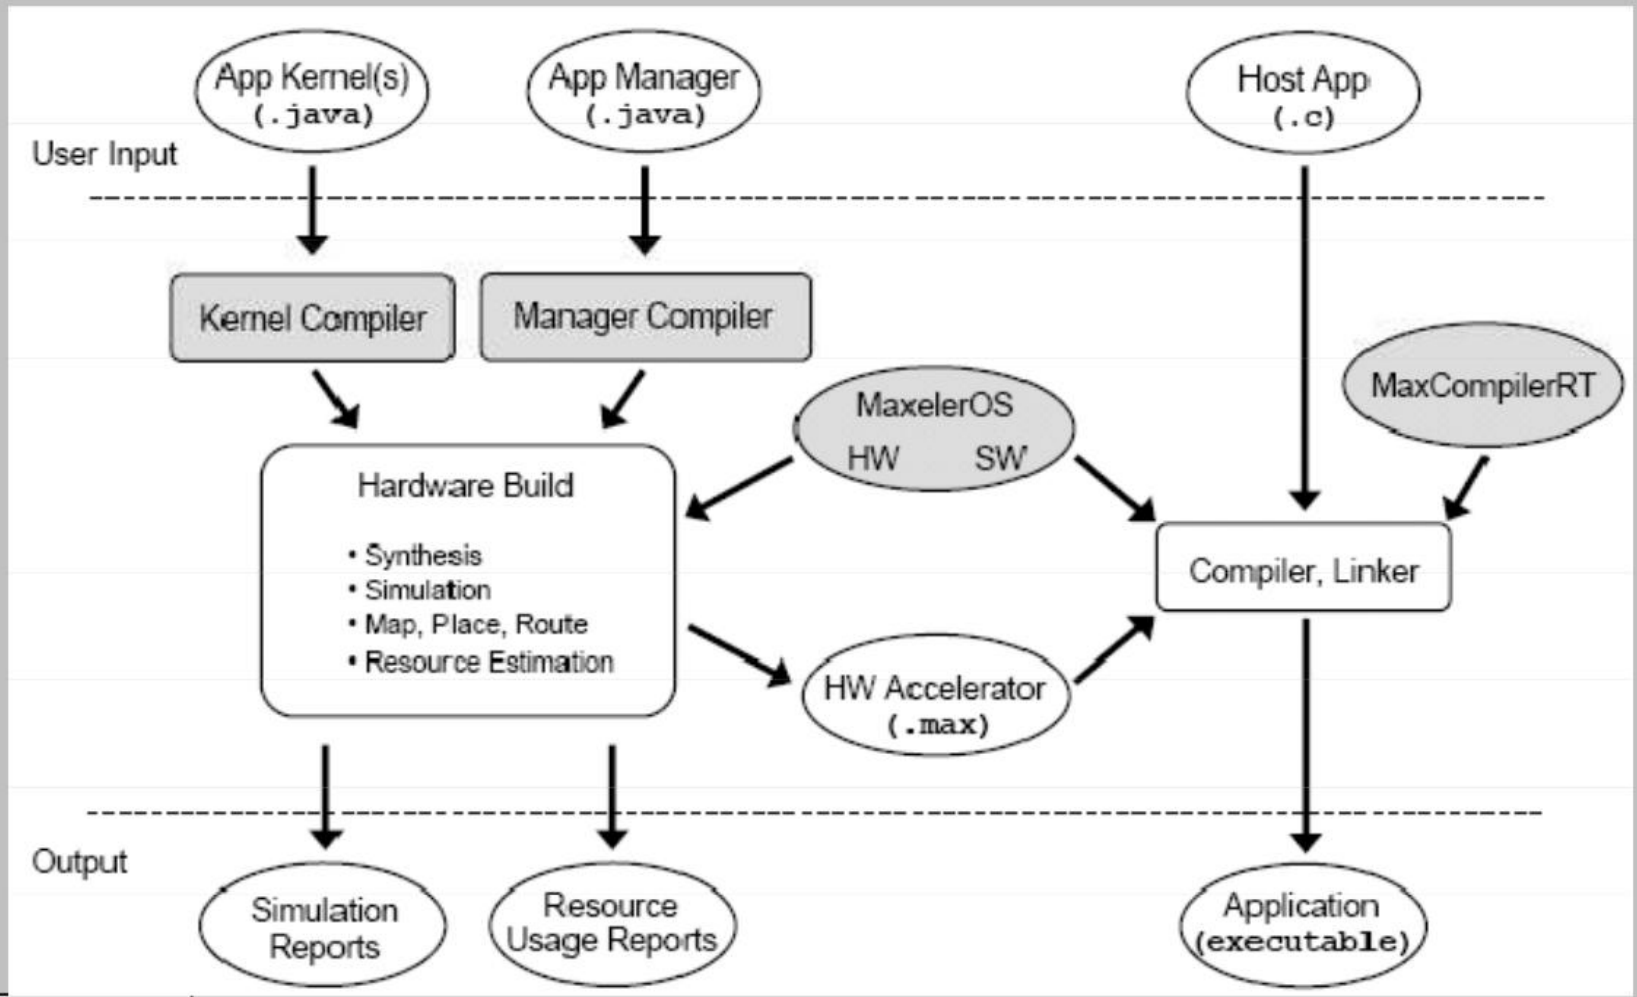

MAXELER Technologies
MAXIMUM PERFORMANCE COMPUTING

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!

Maxeler Dataflow Appliance

- Software Based Solution
- Dataflow Computing in the Datacentre

MPC C Series

The CPU

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

MPC X Series

The Dataflow Appliance

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

MPC N Series

The Networking Appliance

Intel Xeon CPUs and 4 DFEs 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**

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)

- 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

Technology ignites our business. Be the spark.

J.P.Morgan

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 984 MFlops/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 very simple sample:

$f(x) = 2x + x$

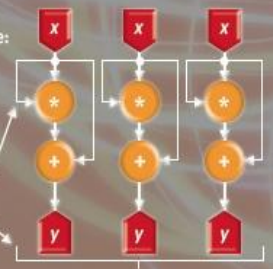

Moving from a single calculation to a fine grained parallelism

A slightly more complex example:

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

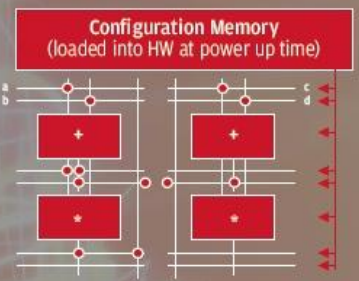

Migrating algorithms from C++ to FPGAs involves doing a Fourier Transform from time domain execution to spatial domain execution in order to maximize computational throughput. **It's a paradigm shift to stream computing that provides acceleration of up to 1,000x compared to an Intel CPU.**

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 Ticker	Bloomberg Ticker	DSF Pricing					Timestamp
		Price	Coupon	PV01	NPV	Implied Rate	
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.

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

MAXELER Technologies

Seismic Data Acquisition

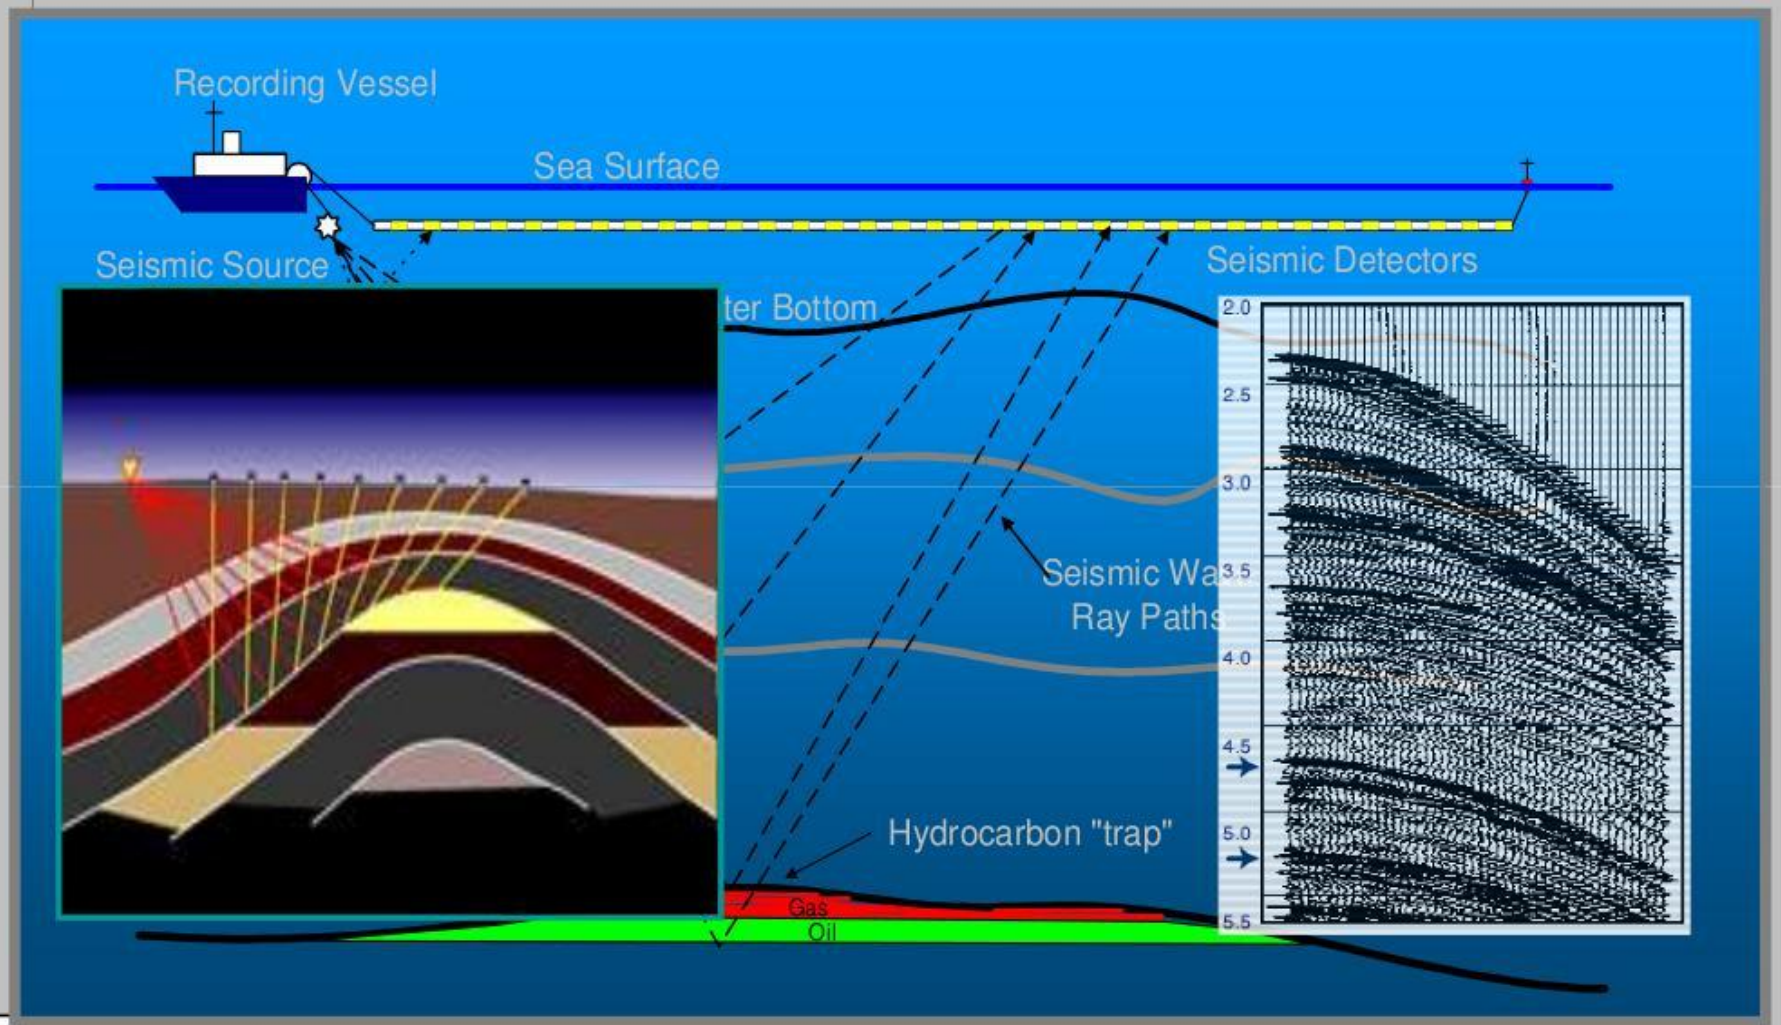

Courtesy of Schlumberger

MAXELER Technologies
MAXIMUM PERFORMANCE COMPUTING

Real-Time Data Filtering

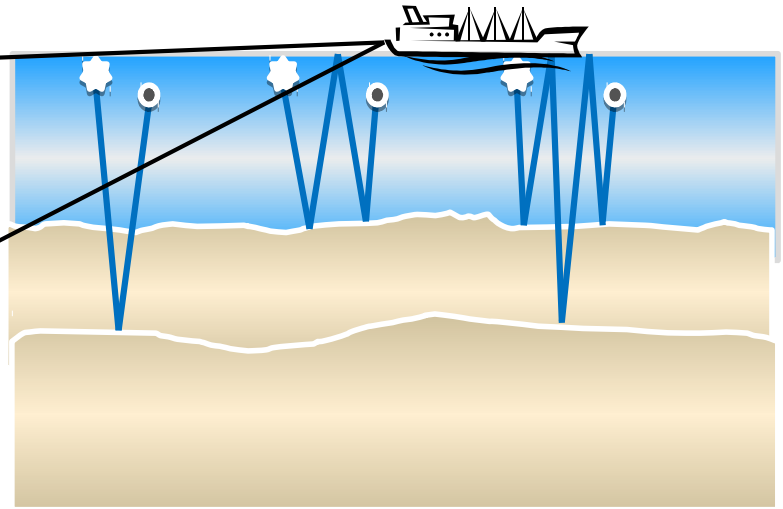

Low-Power Cluster

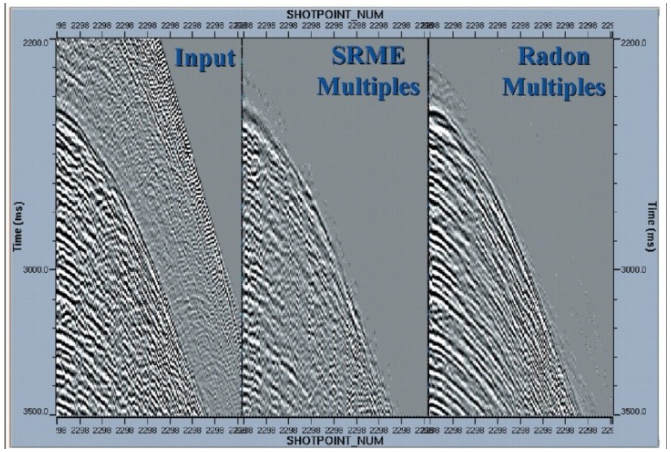

MAXELER Technologies

Seismic Imaging

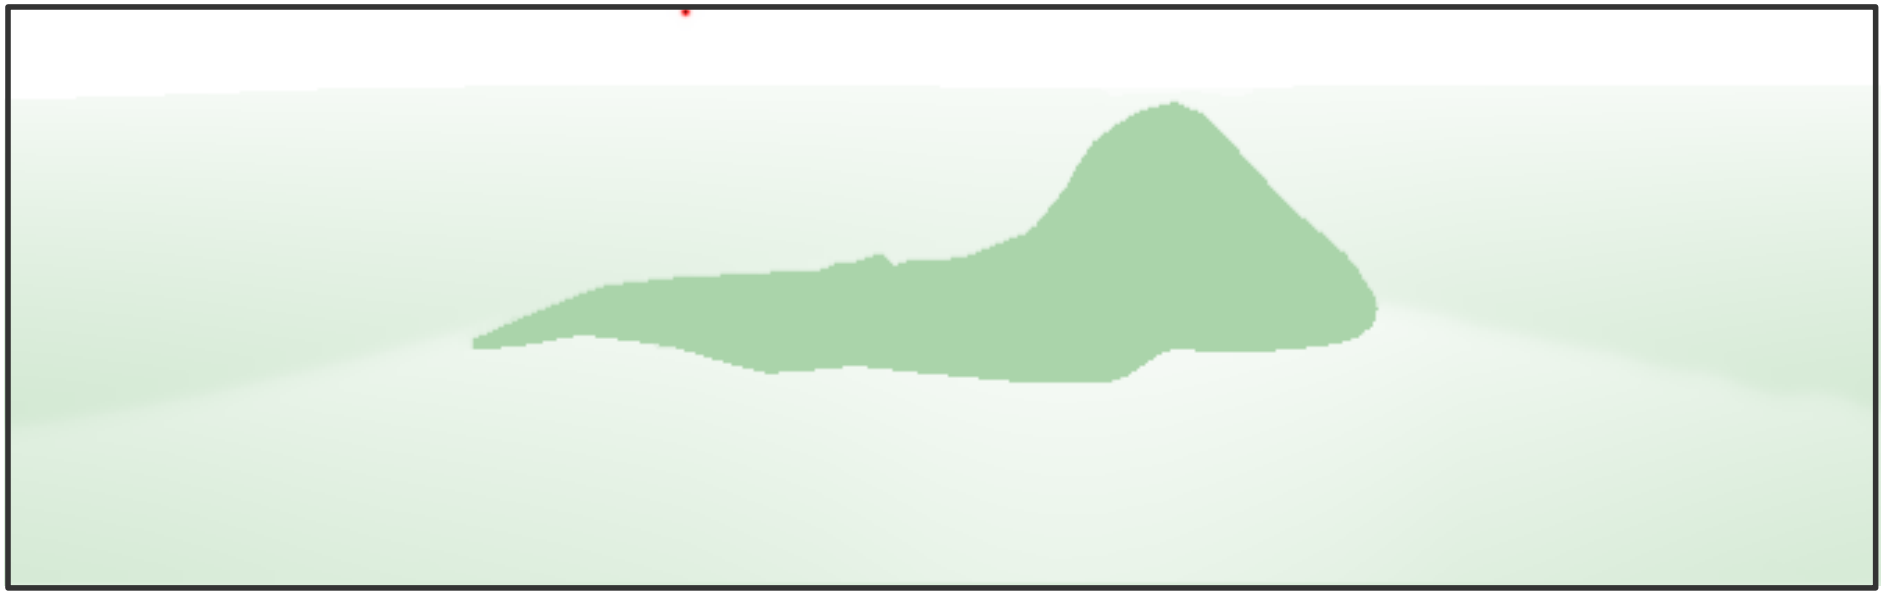

- Running on MaxNode servers
  - 8 parallel compute pipelines per chip
  - 10x less power: 150MHz vs 1.5GHz
  - 30x faster than microprocessors

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

Real-Time Seismic Imaging Appliance

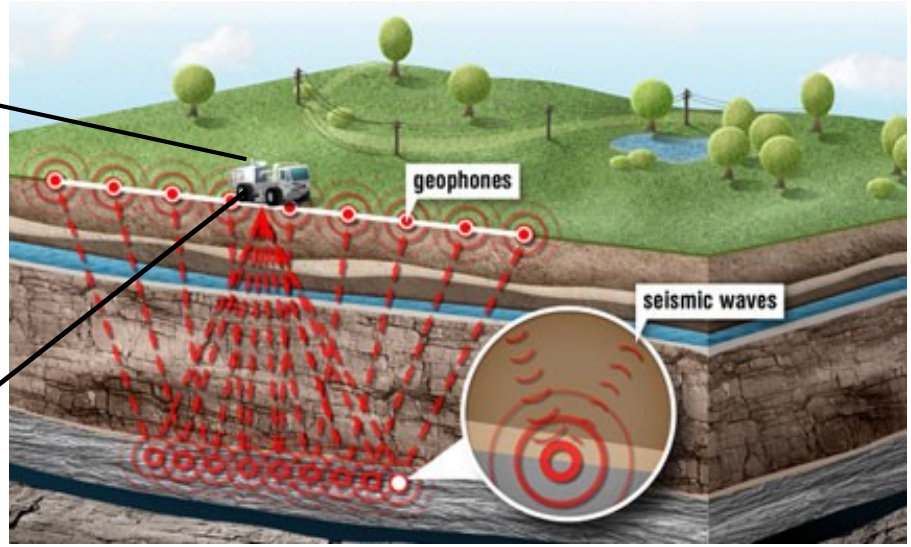

Seismic Imaging Appliance

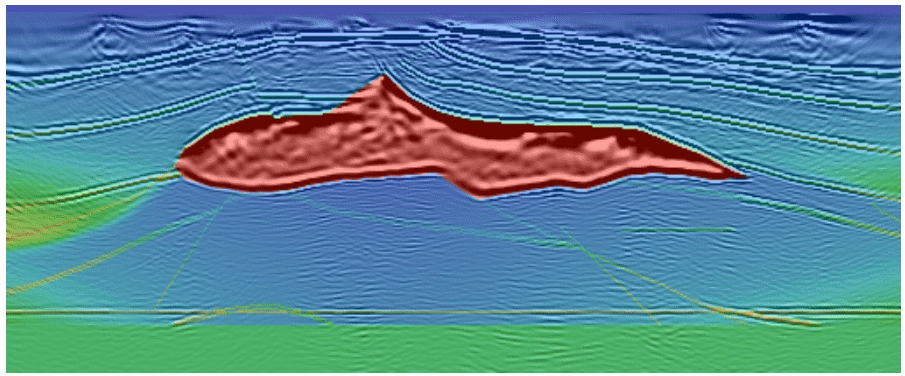

Real Time Monitoring

Fracturing monitoring during high pressure injection

Passive seismic

4D seismic imaging

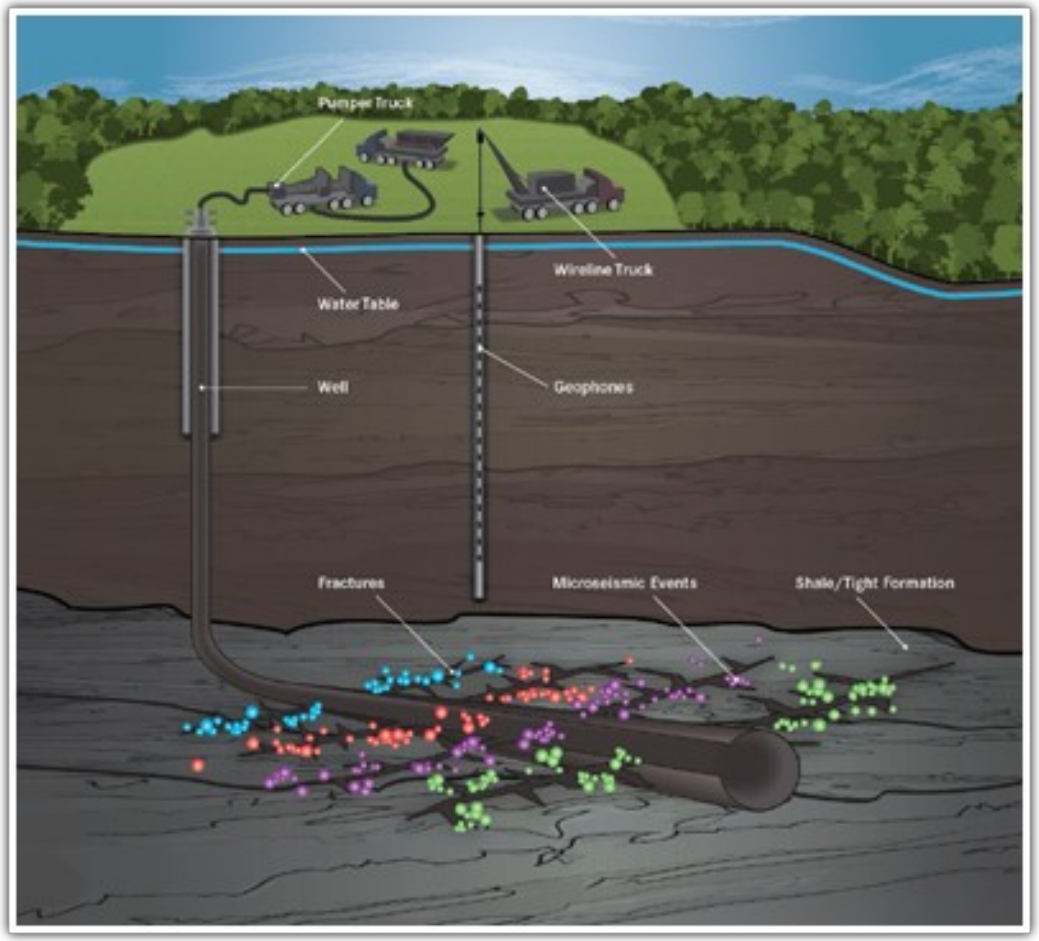

Global Weather Simulation

- Atmospheric equations

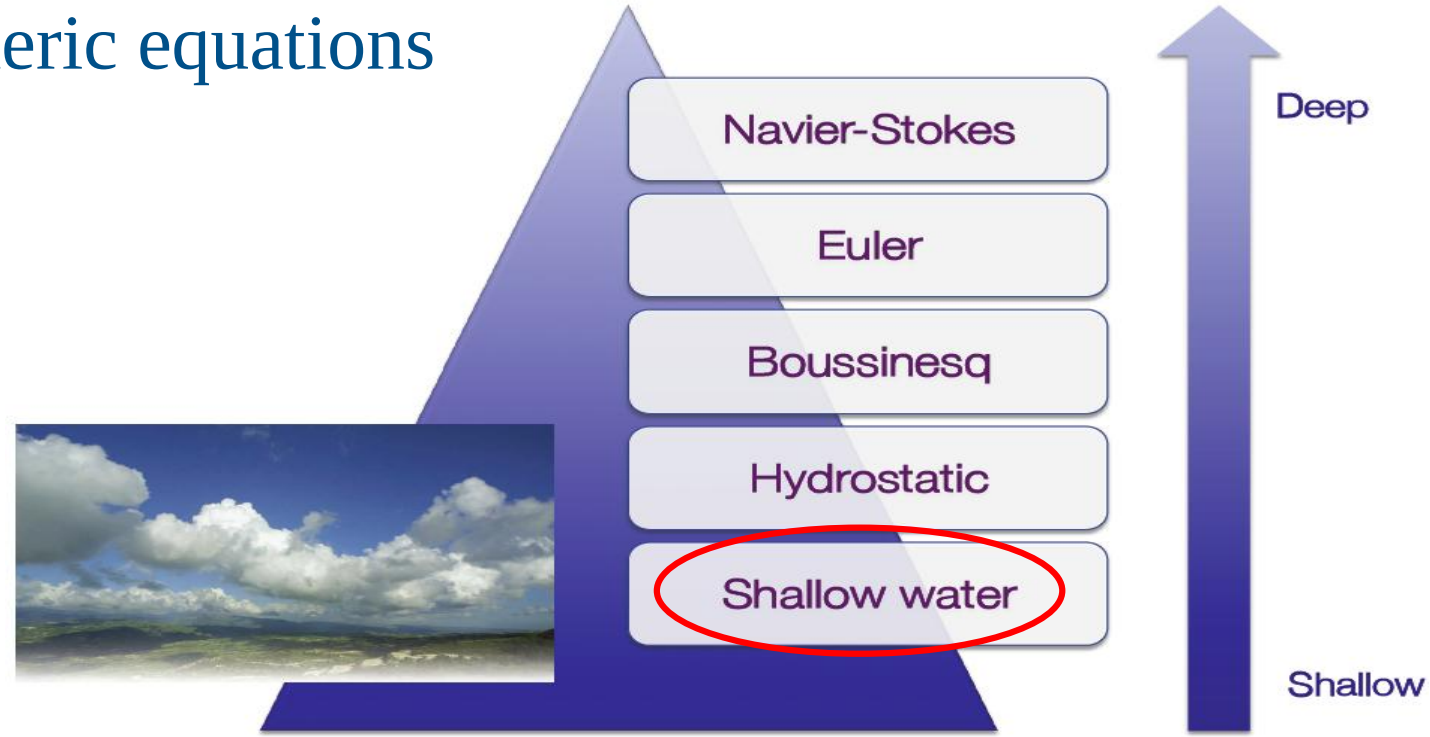

- Equations: Shallow Water Equations (SWEs)

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

[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]

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 × 1024 × 6

MaxNode speedup over Tianhe node: 14 times

14x

Imperial College London

ISCAS

MAXELER Technologies

Weather Model -- Power Efficiency

Platform	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 × 1024 × 6

MaxNode is 9 times more power efficient

9 x

Imperial College London

ISCAS

MAXELER Technologies

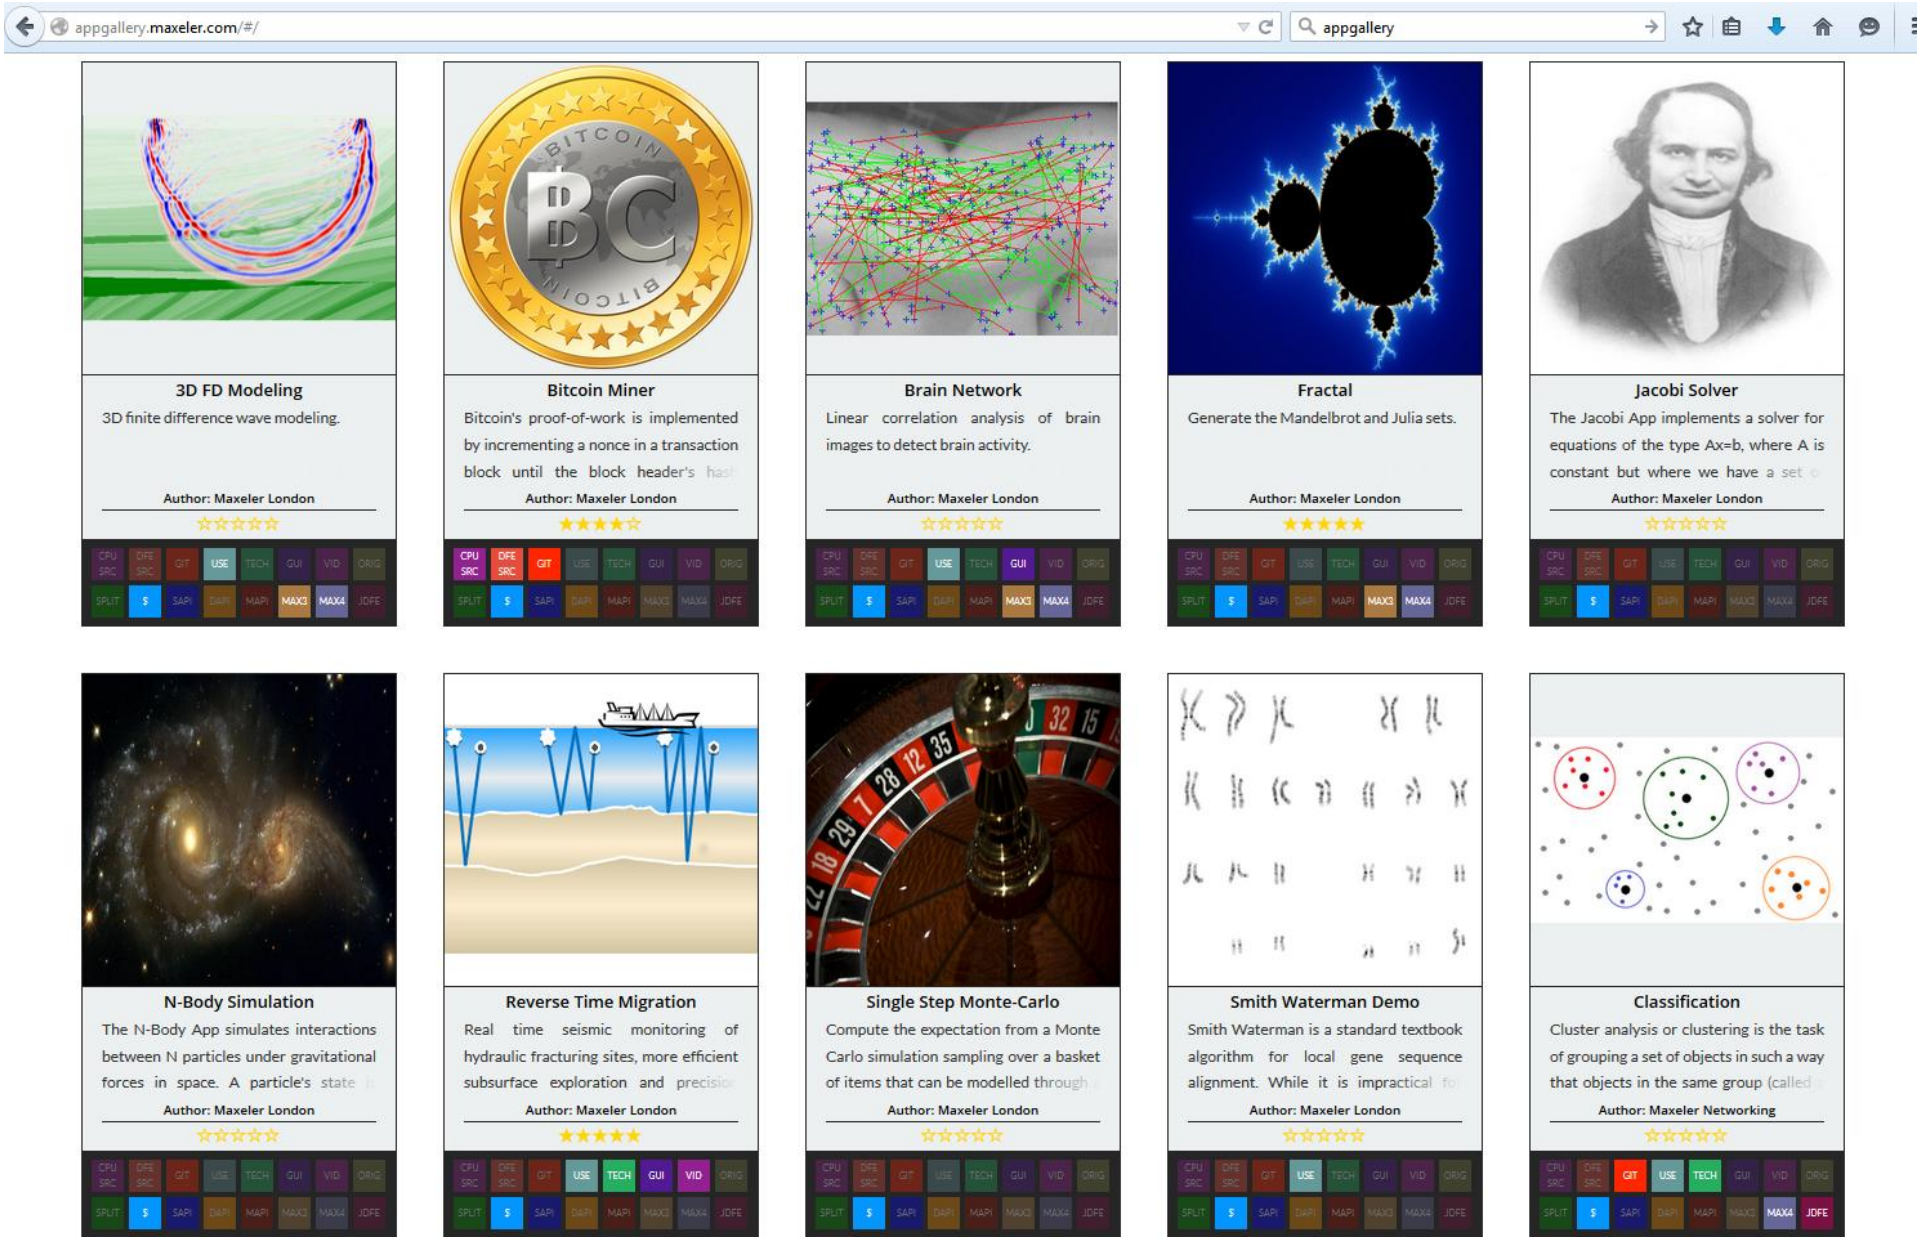

appgallery.maxeler.com

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

Ačiū už dėmesį