Acceleration of Cloud Data Centers: Challenges and future directions

Christoforos Kachris, NTUA, InAccel

Dimitrios Soudris, NTUA

Making Software Projects Easier to Understand

- Challenges
  - Heterogeneous computing
  - The need of accelerators
  - Domain specific architectures

Data Center traffic

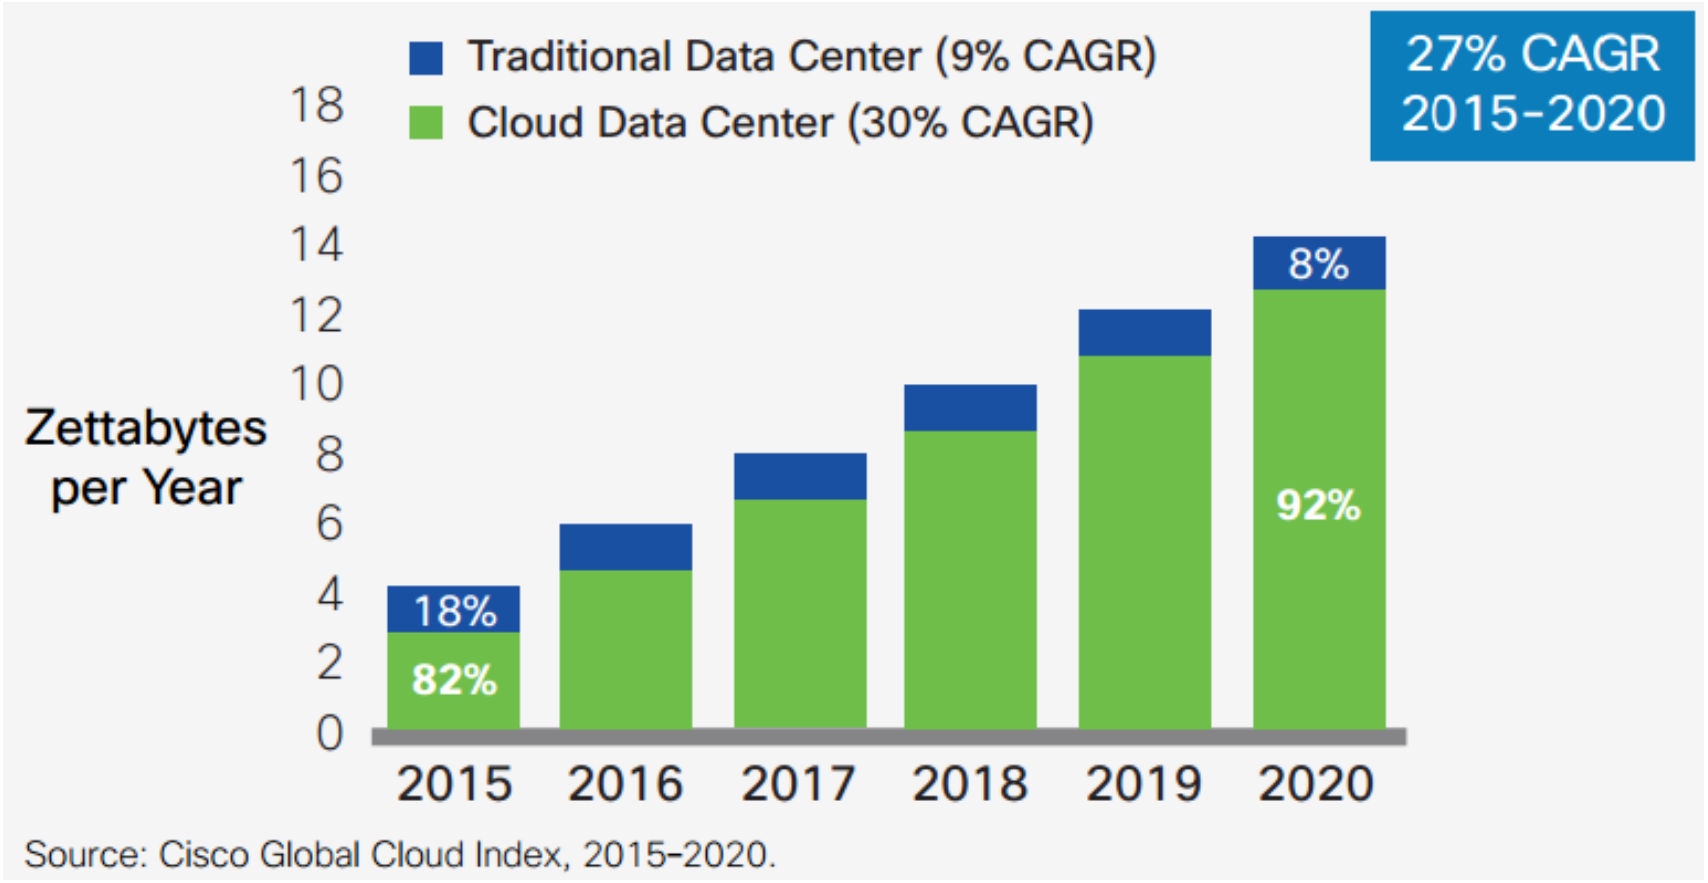

Source: Cisco Global Cloud Index, 2015–2020.

Data Center applications

Accelerators can increase performance at lower TCO for targeted workloads

Financial market

- In the US alone, NYSE processes more than **1.4 billion messages** every day….
- …That's 3.5 times more than the number of internet searches Google handles daily.

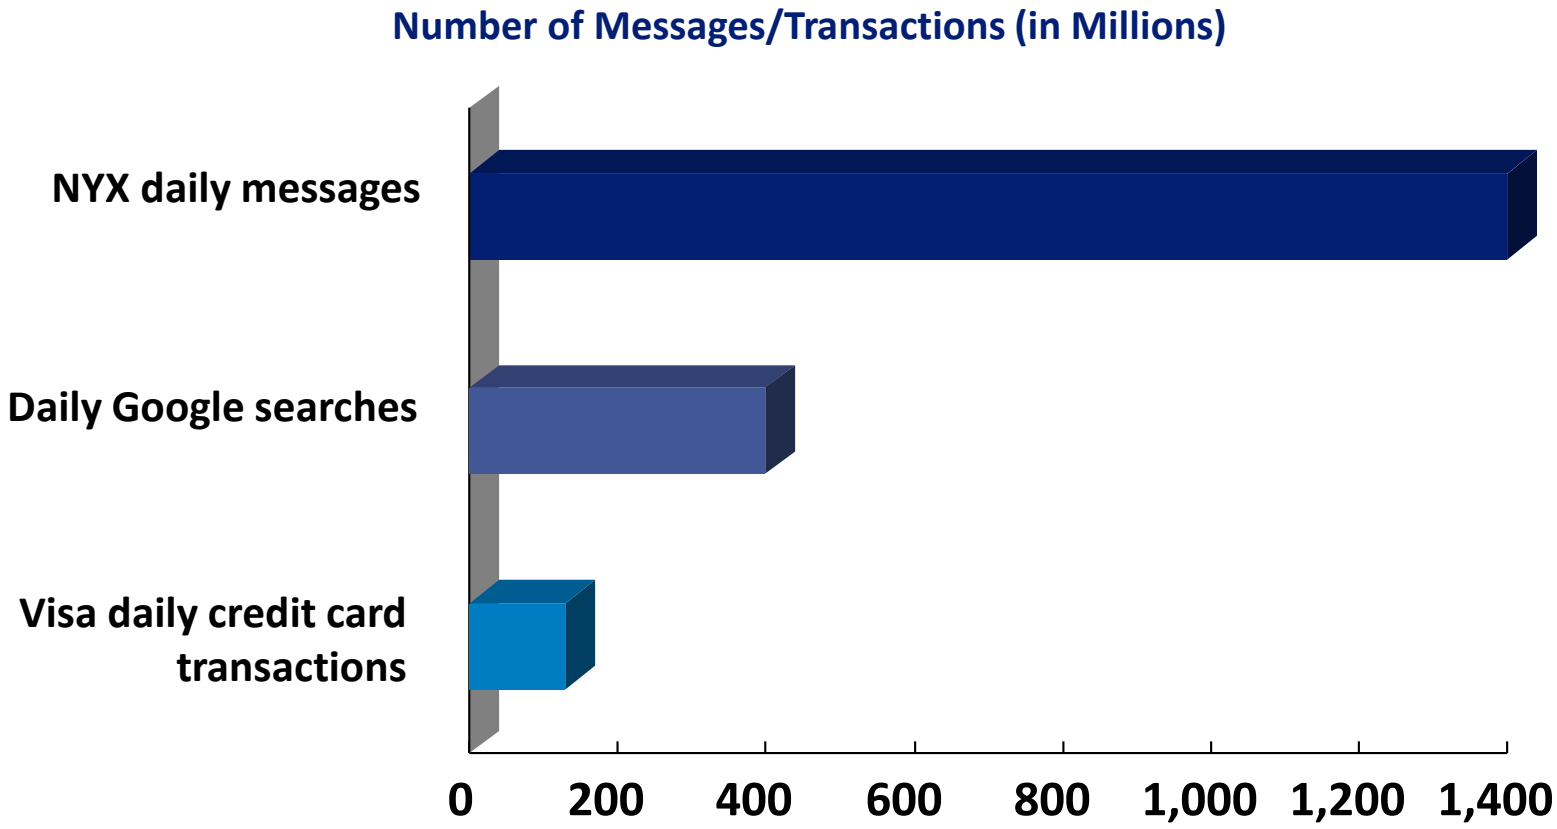

Data Science: need for high computing power

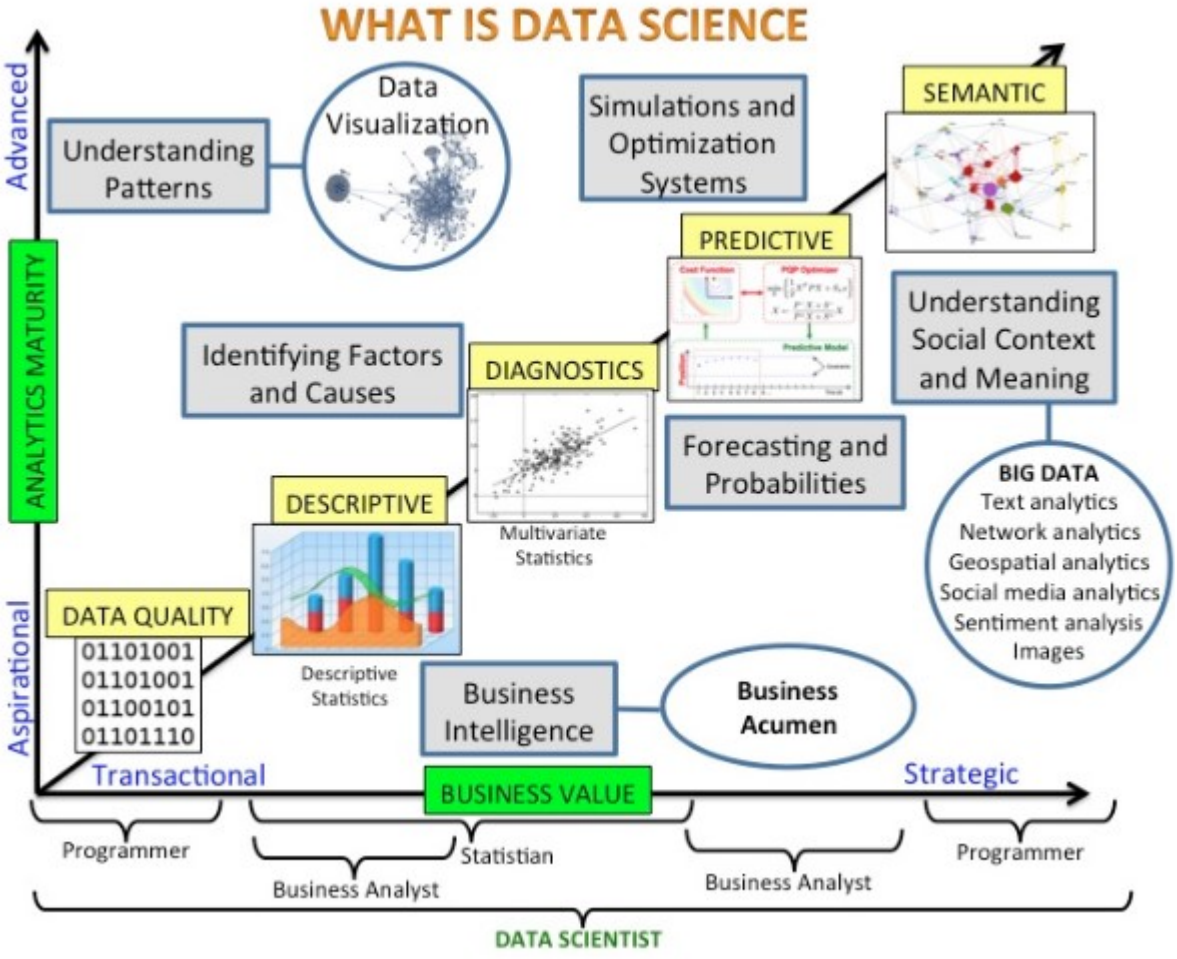

How Big are Data Centers

Data Center Site	Sq ft
Facebook (Santa Clara)	86,000
Google (South Carolina)	200,000
HP (Atlanta)	200,000
IBM (Colorado)	300,000
Microsoft (Chicago)	700,000

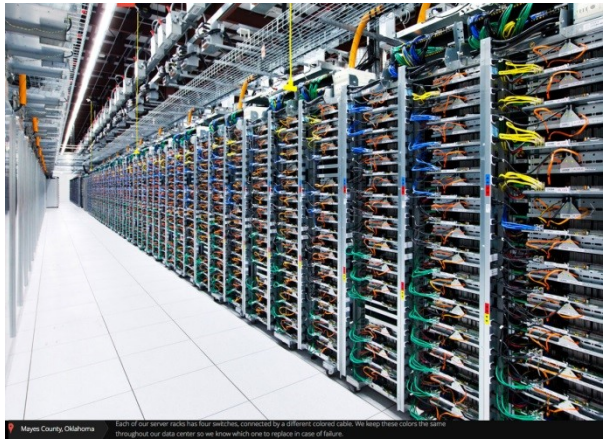

[Source: "How Clean is Your Cloud?", Greenpeace 2011]

Wembley Stadium:**172,000** square ft

Google data center

Apple's Data Center

Data Centers Power Consumption

- Data centers consumed **330 Billion KWh** in 2007 and is expected to reach **1012 Billion KWh** in 2020

	2007 (Billion KWh)	2020 (Billion KWh)
Data Centers	**330**	**1012**
Telecoms	293	951
Total Cloud	**623**	**1963**

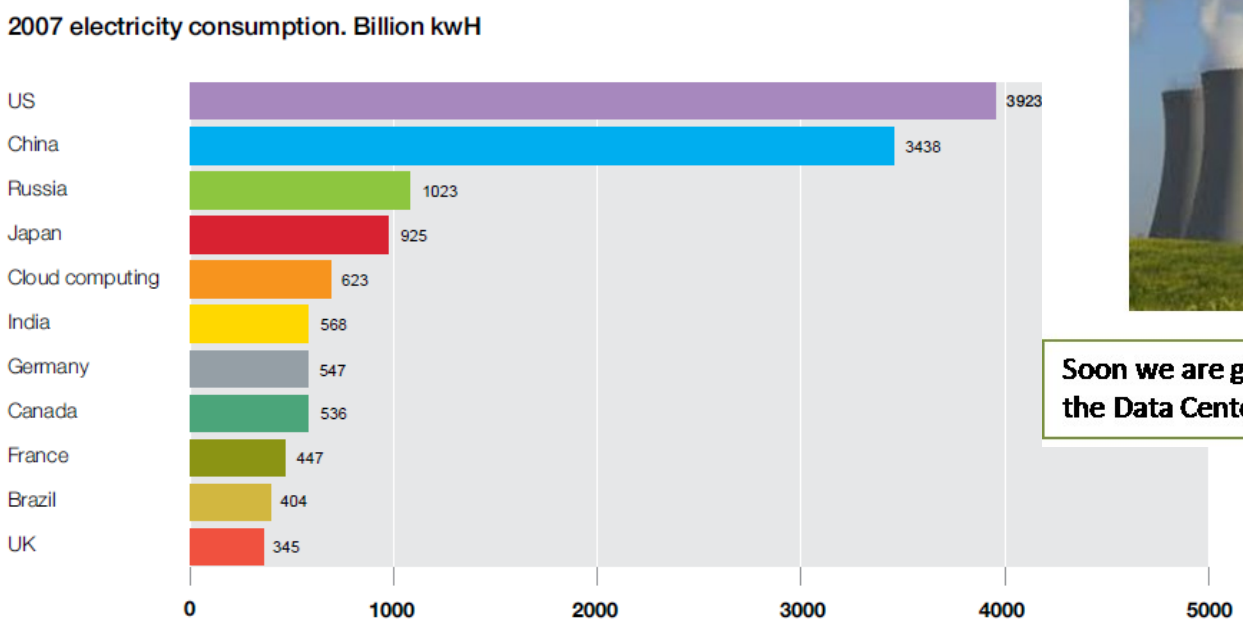

Soon we are going to need a power plant next to the Data Centers

[Source: How Clean is Your Data Center, Greenpeace, 2012

Carbon Emissions of ICT Sector

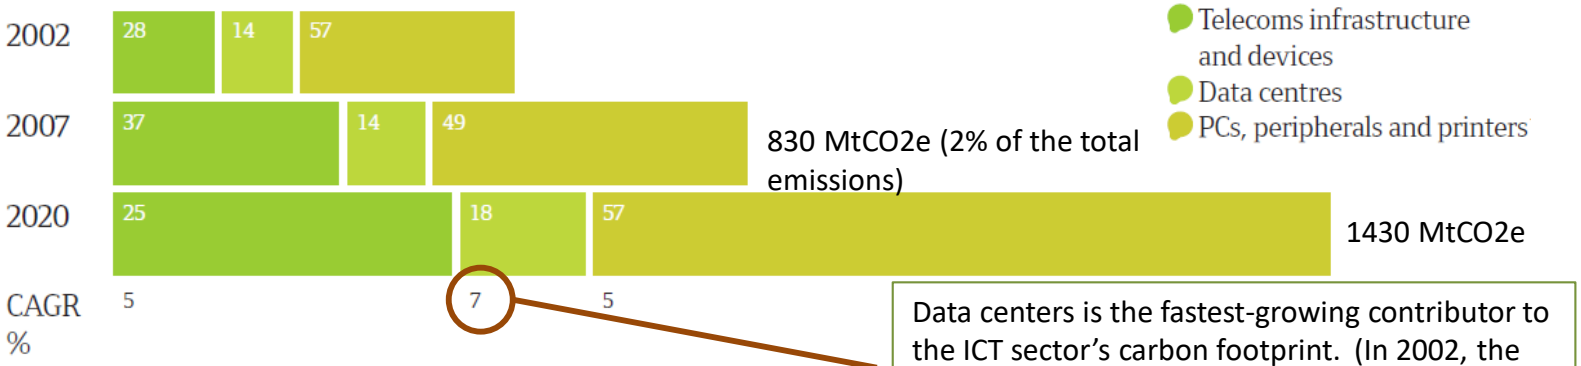

[Source: SMART 2020: Enabling the low carbon economy in the information age , Greenpeace]

Data centers is the fastest-growing contributor to the ICT sector's carbon footprint. (In 2002, the global data centre footprint, including equipment use and embodied carbon, was 76 MtCO2e and this is expected to more than triple by 2020 to 259 MtCO2e)

The CO2 emissions of 10,000 Google searches is equal to a five mile trip in the average U.S. automobile [Source: Google, Inc.]

End of Growth of single program speed

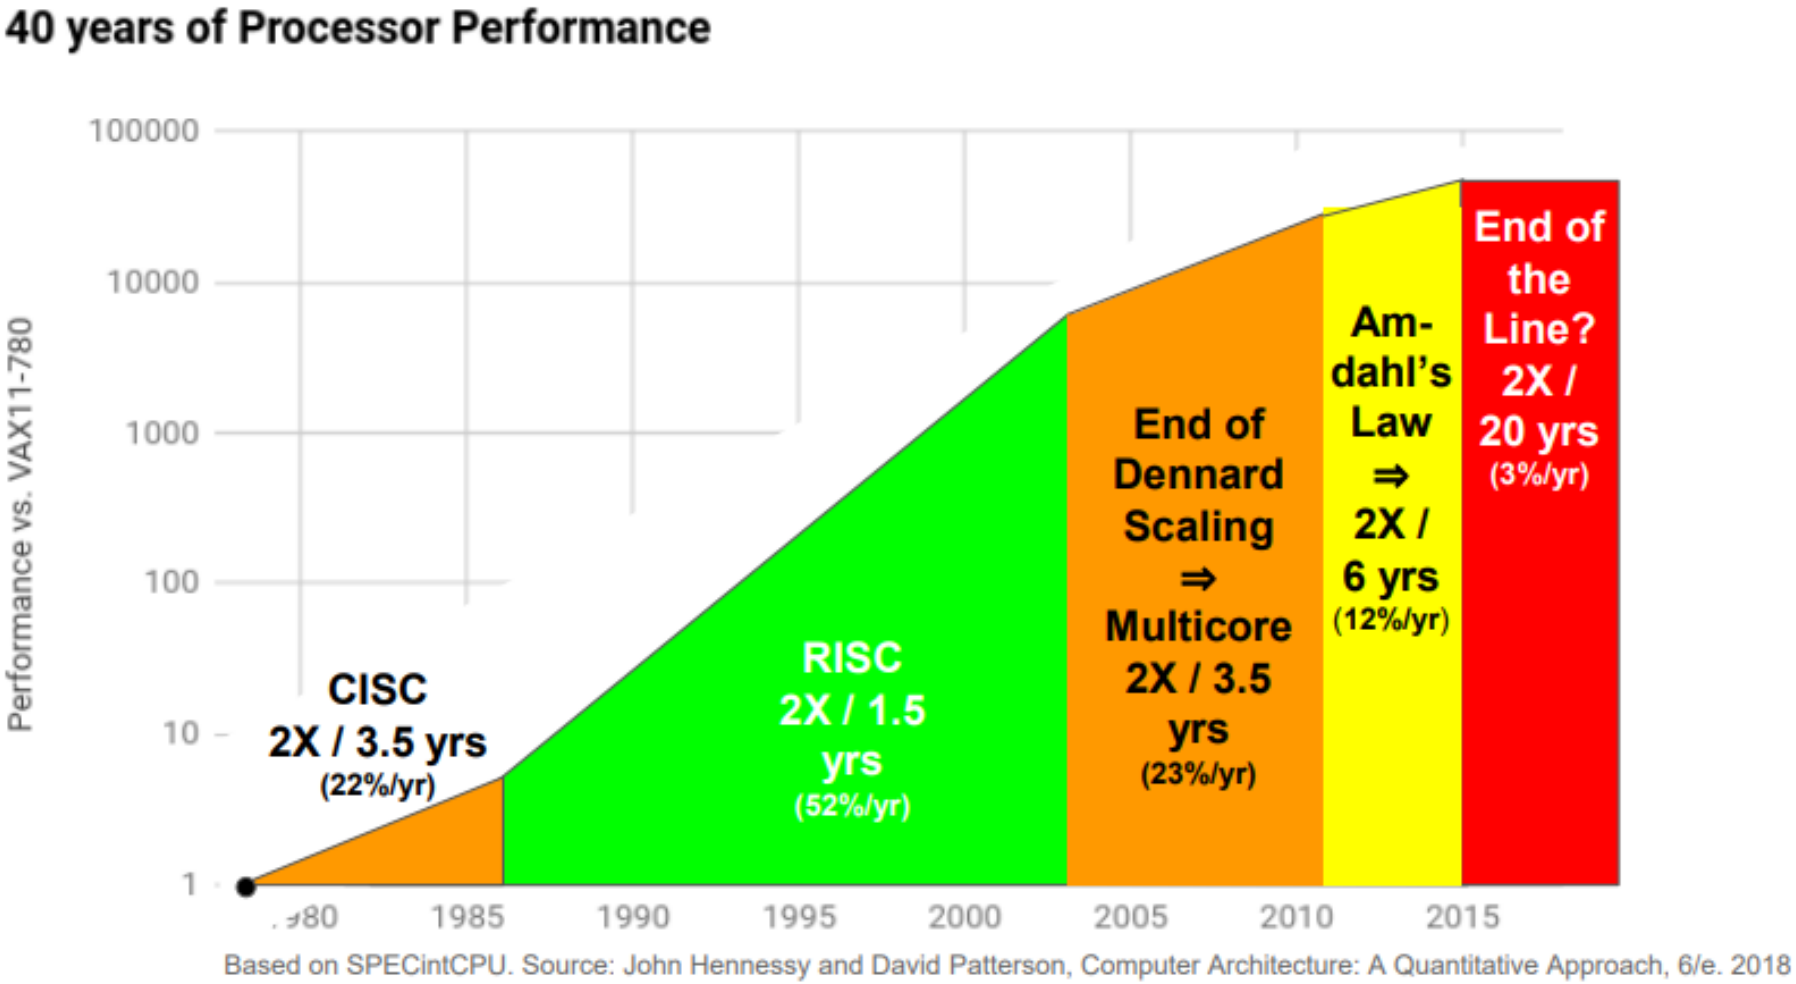

Hardware acceleration

Hardware acceleration is the use of specialized hardware components to perform some functions faster (10x-100x) than is possible in software running on a more general-purpose CPU.

- Hardware acceleration can be performed either by **specialized chips (AICS)** or
- By programmable specialized chips (**FPGAs**) that can be configured for specific applications

Hardware accelerators

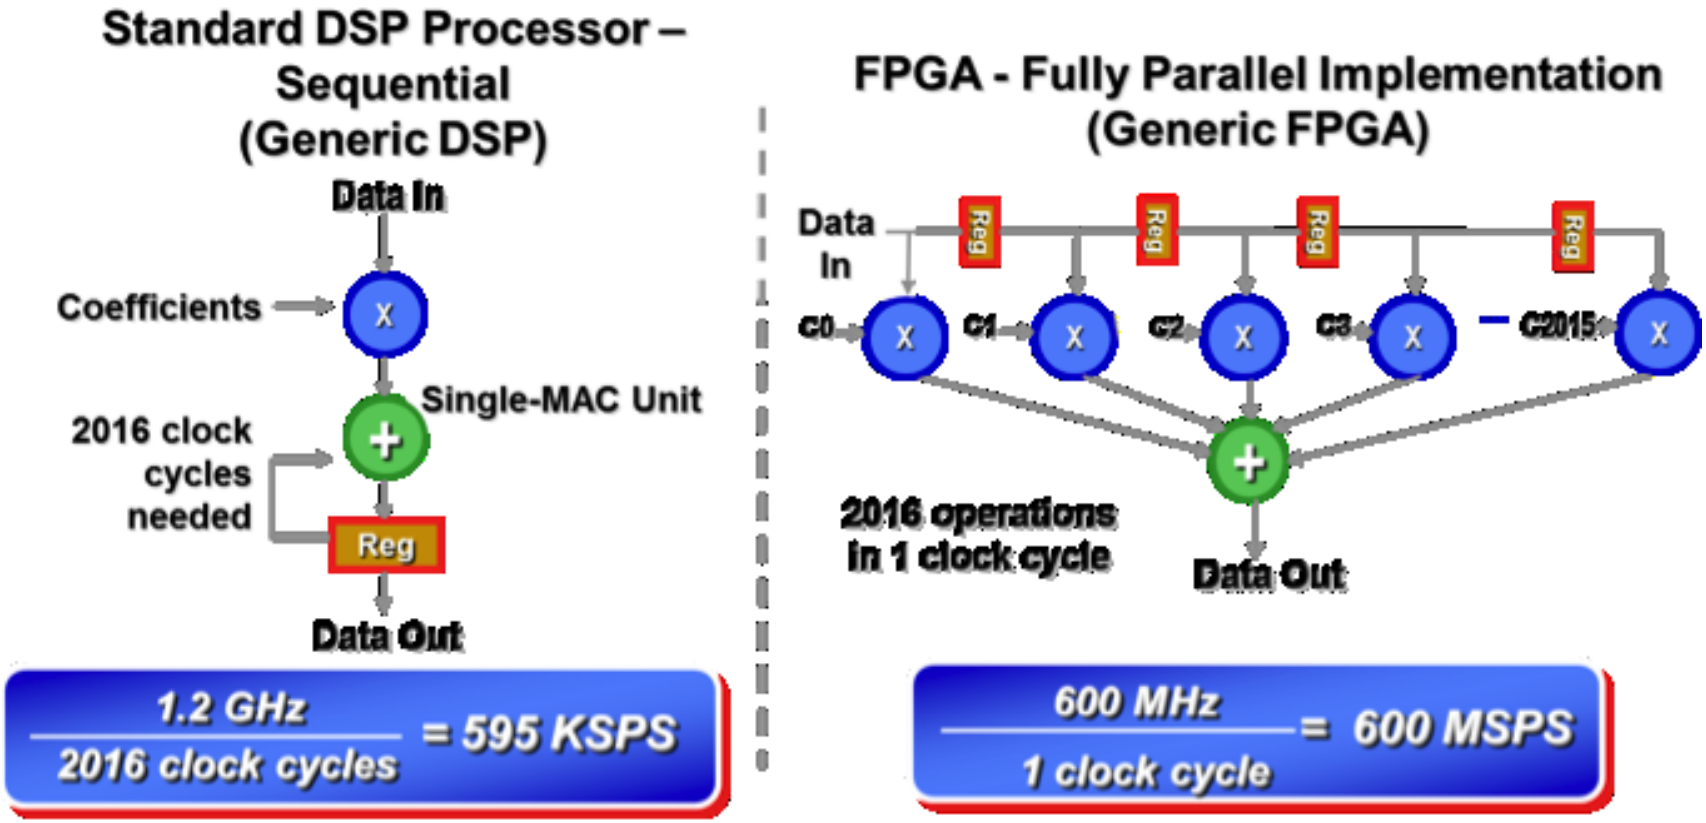

Accelerators

A GPU is effective at processing the same set of operations in parallel – single instruction, multiple data (SIMD). A GPU has a well-defined instruction-set, and fixed word sizes – for example single, double, or half-precision integer and floating point values.

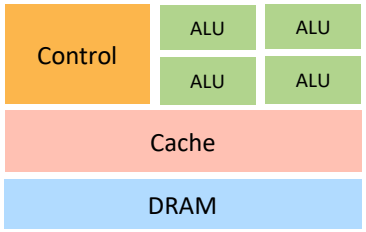

CPU
(one core)

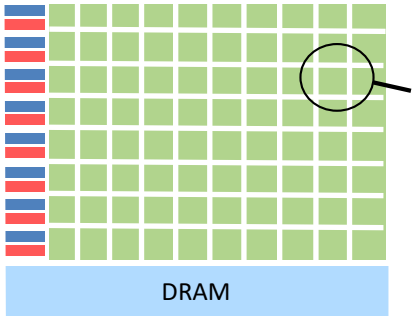

Each GPU has 2880 of these cores

GPU

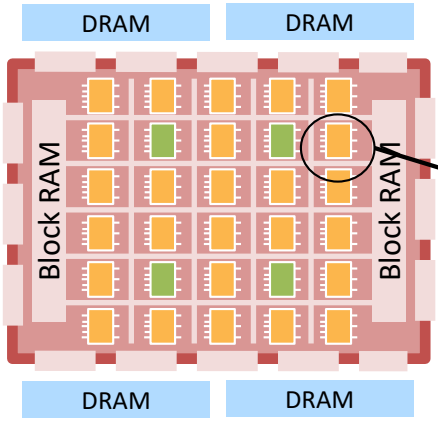

Each FPGA has more than 2M of these cells

FPGA

An FPGA is effective at processing the same or different operations in parallel – multiple instructions, multiple data (MIMD). An FPGA does not have a predefined instruction-set, or a fixed data width.

Processing Platforms

- HW acceleration can be used to reduce significantly the execution time and the energy consumption of several applications (10x-100x)

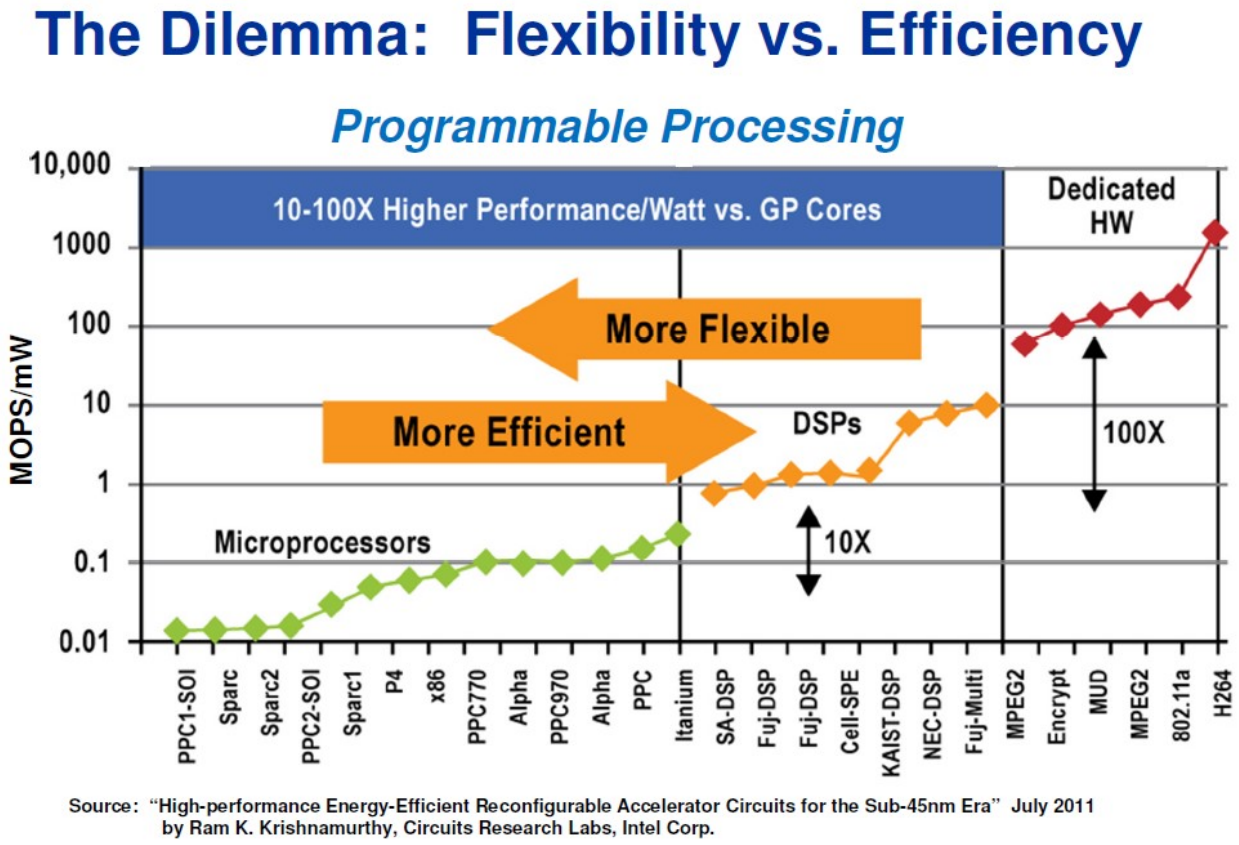

Source: "High-performance Energy-Efficient Reconfigurable Accelerator Circuits for the Sub-45nm Era" July 2011 by Ram K. Krishnamurthy, Circuits Research Labs, Intel Corp.

Hardware accelerators

CPU: High flexibility, low throughput

Best of 2 worlds

Chip Accelerators: High throughput, high NRE
Initial one-off costs

Programmable accelerators (FPGAs) can provide massive parallelism and higher efficiency than CPUs for certain categories of applications

InAccel offers novel accelerators as IP cores to boost the performance of their applications (similar to specialized containers)

FPGAs in the cloud

- Altera's acquisition by Intel

- Microsoft's catapult for Bing search

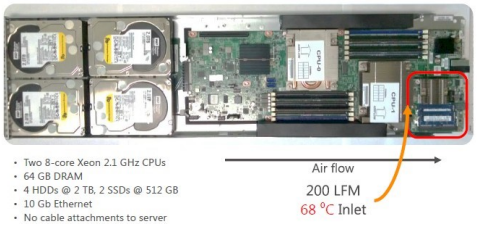

- IBM open power – CAPI interface with FPGAs

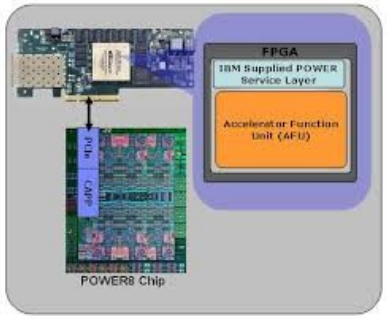

Intel Xeon + FPGAs

Software Development for Accelerating Workloads using Xeon and coherently attached FPGA in-socket

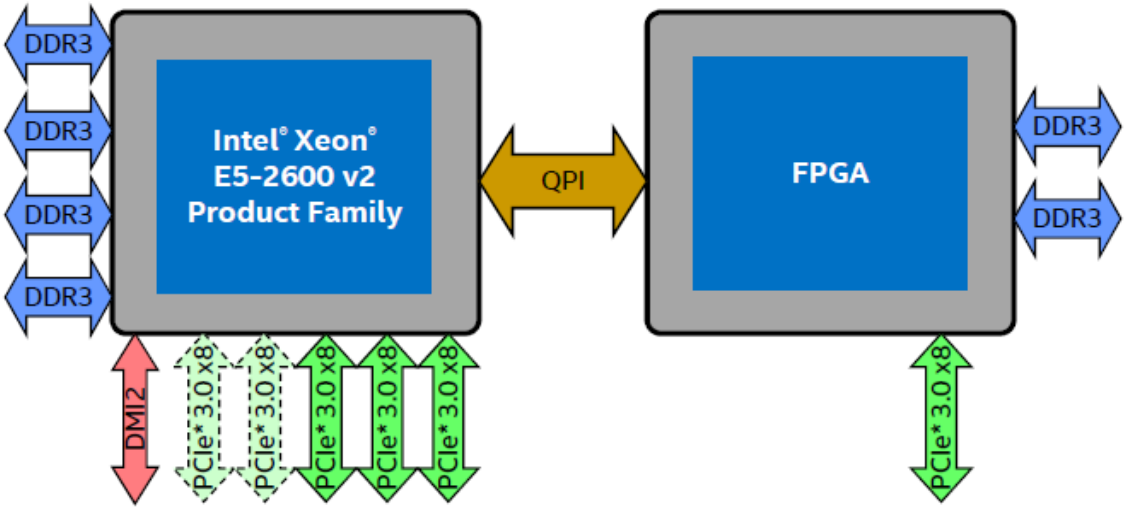

Processor	Intel® Xeon® E5-26xx v2 Processor
FPGA Module	Altera Stratix V
QPI Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR3 (up to 64 GB)
Expansion connector to FPGA Module	PCIe 3.0 x8 lanes - maybe used for direct I/O e.g. Ethernet
Features	Configuration Agent, Caching Agent,, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

Programming Interface

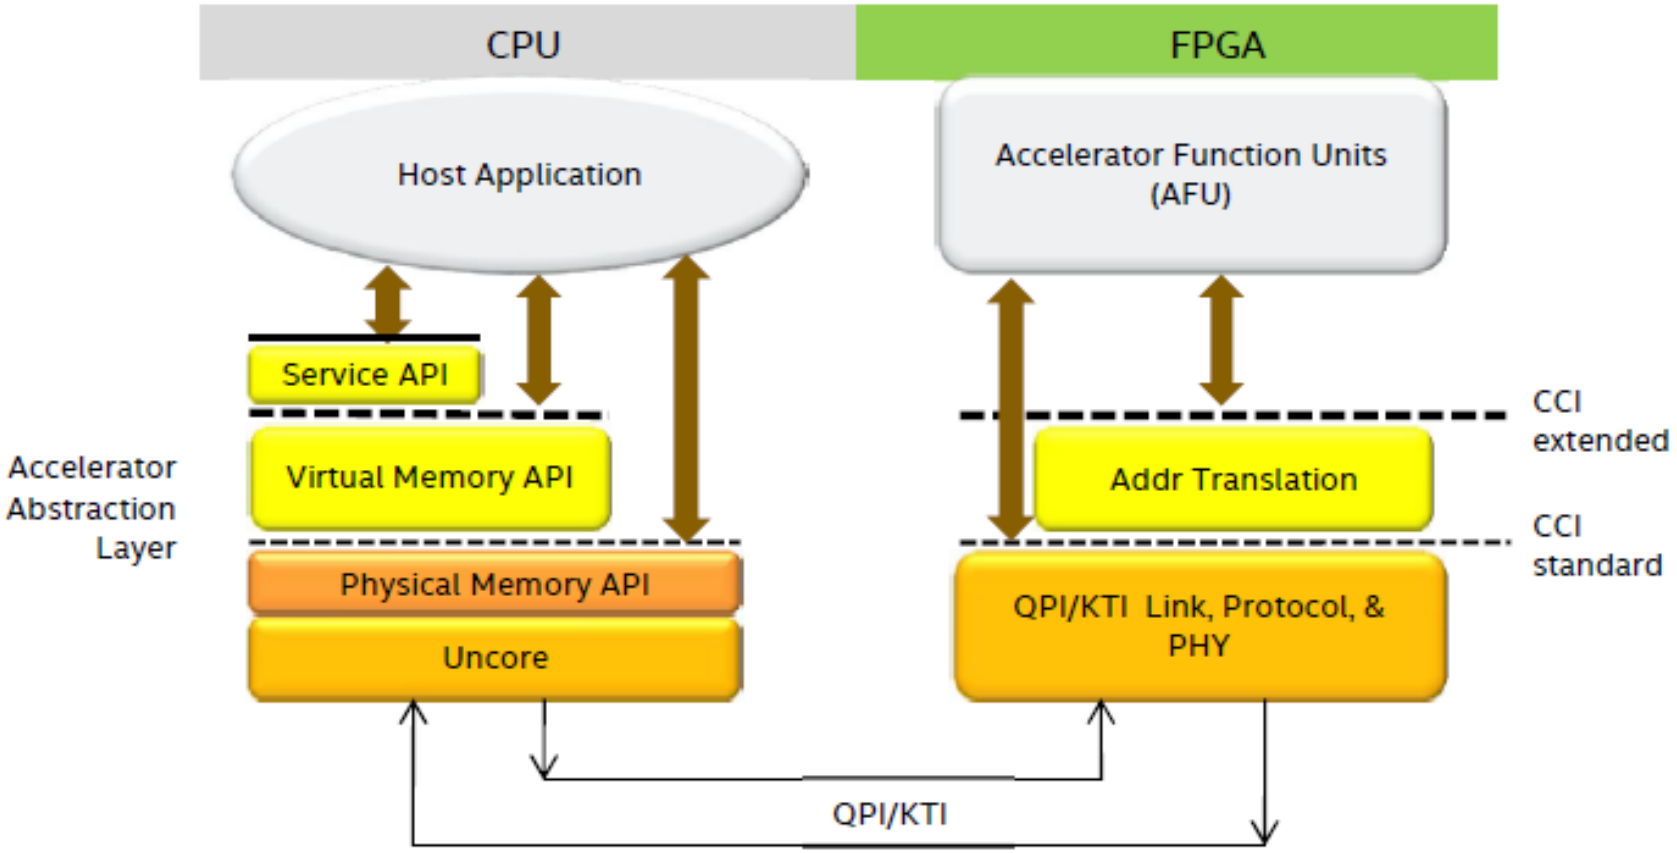

Catapult FPGA acceleration Card

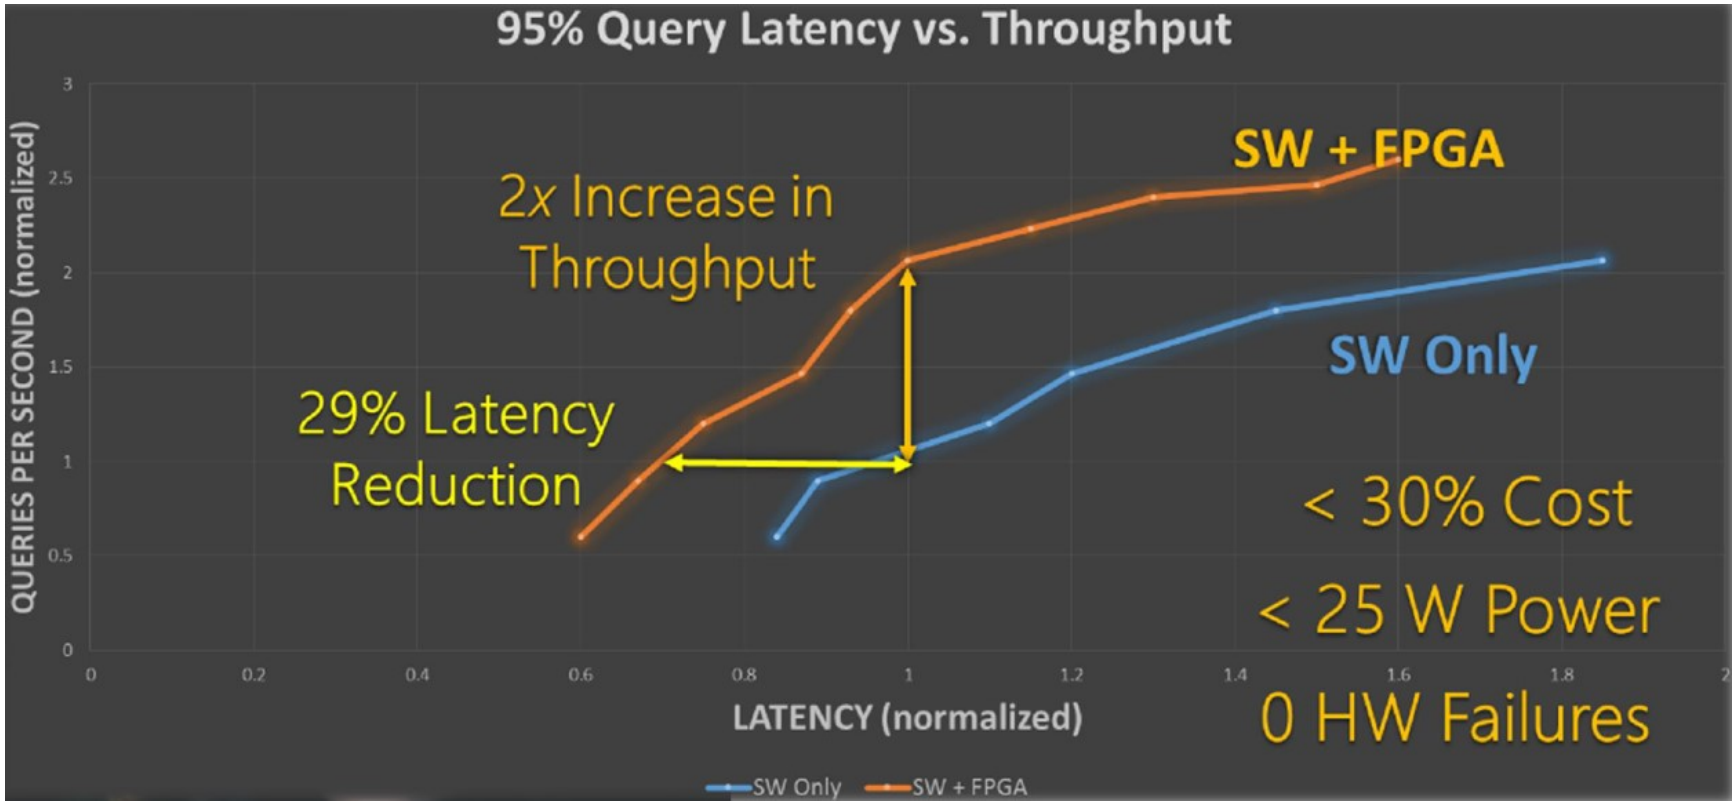

VINEYARD

VINEYARD Framework

- **Accelerators stored in an AppStore**
- **Cloud users request accelerators based on applications requirements**
- **Decouple Hardware – Software designers**

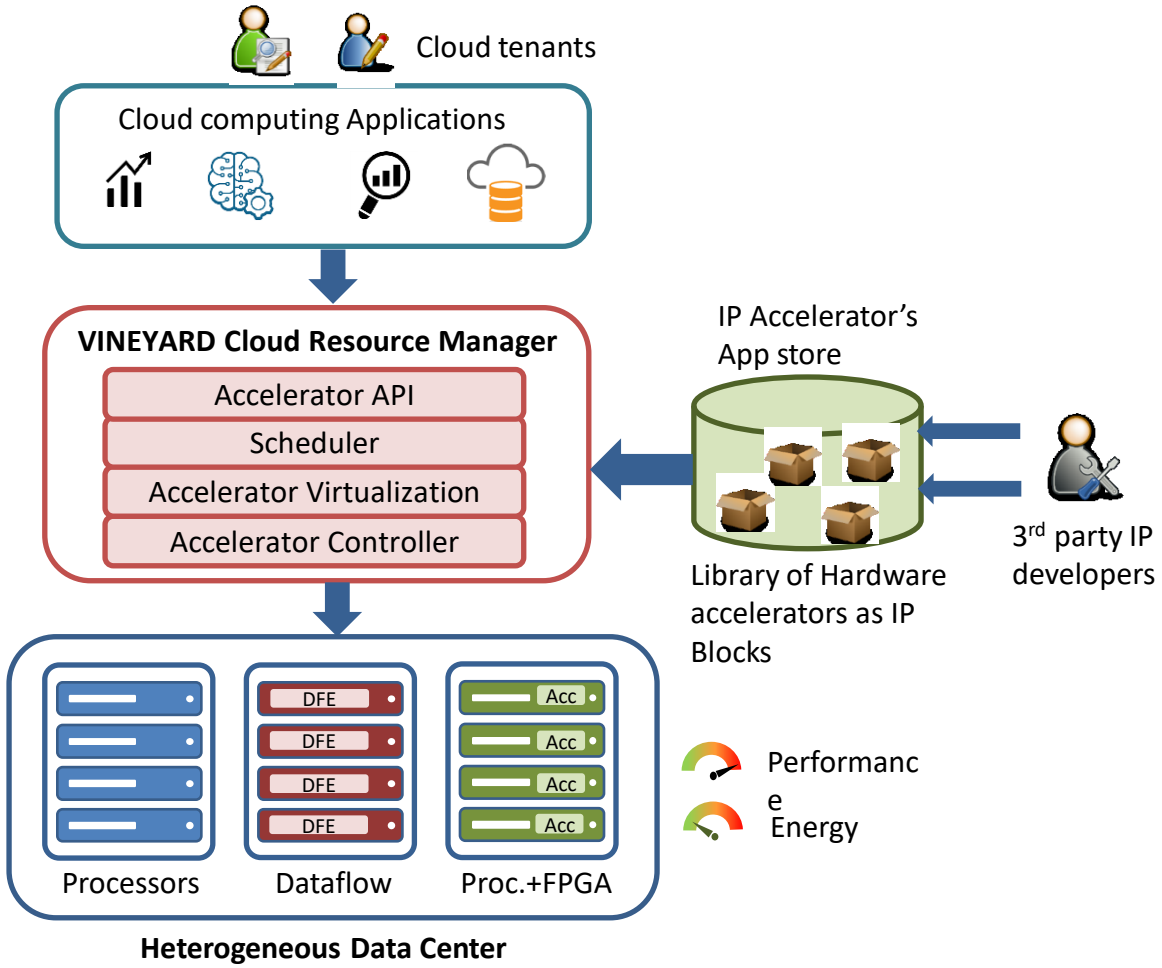

VINEYARD

FPGA as a Service

- Amazon EC F1's Xilinx FPGA

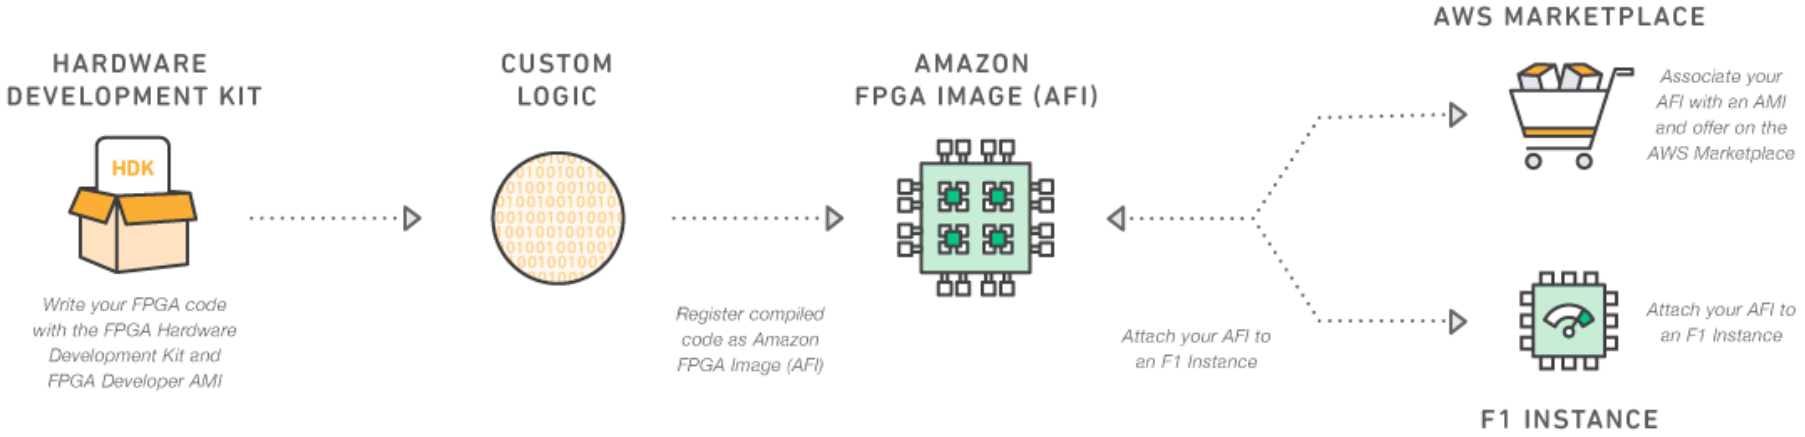

Commercialization

https://aws.amazon.com/marketplace/pp/B07KXQLTJ3

InAccel Products

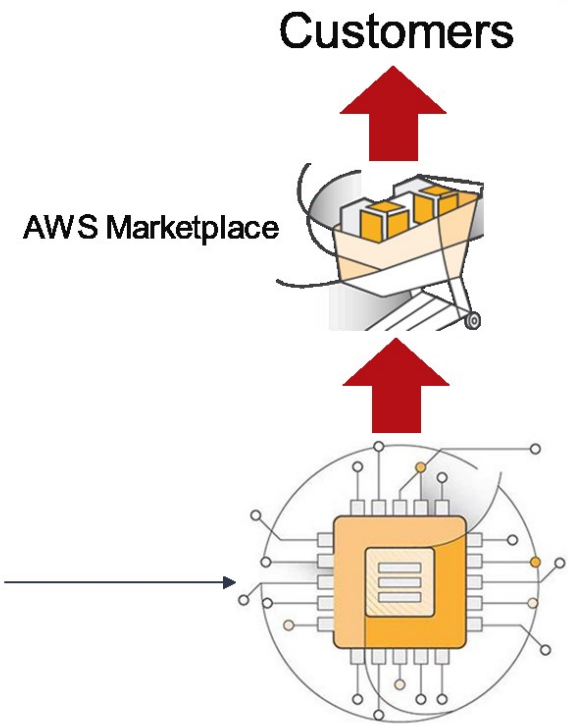

Amazon EC2 FPGA Deployment via Marketplace

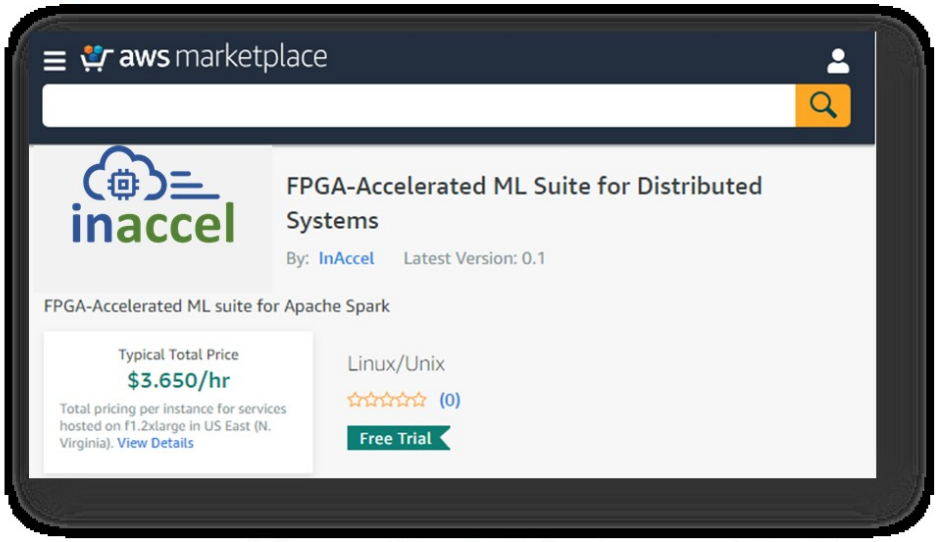

Scalable to worldwide market

First to provide accelerators for Spark

Heterogeneous made easy

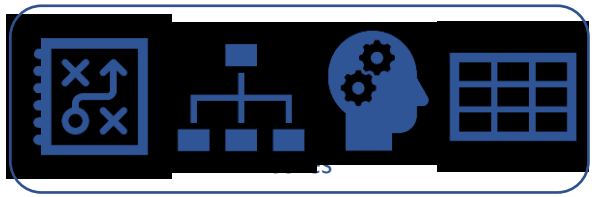

Scalable
Resource Manager

Alveo U200

On-premise

Cloud

3x-10x speedup

4x Lower Cost

Zero-code changes

Acceleration of Spark

- **Spark is the most widely used framework for Data Analytics**
- **Develop hardware components as IP cores for widely used applications**
  - Spark
    - Logistic regression
    - Recommendation
    - K-means
    - Linear regression
    - PageRank
    - Graph computing

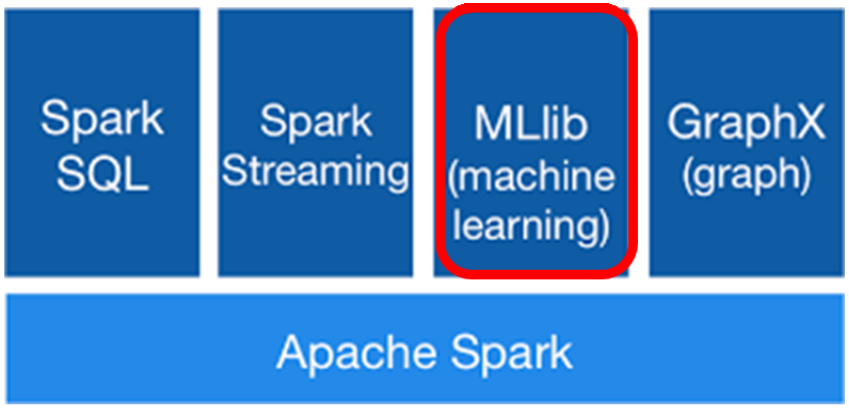

Seamless integration with Zero-code changes:

Without inaccel

```
> spark-submit
```

With inaccel

```
> spark-submit --inaccel
```

Acceleration made easy

Now we're ready to build a pipeline and fit it. This puts the data through all of the feature processing, model tuning & training we described in a single call.

In [*]:
```
from pyspark.ml import Pipeline

pipeline = Pipeline(stages=[labelIndexer, featuresScaler, cv, indexToLabel])

%time model = pipeline.fit(train)
```

The next cell converts the test set from LibSVM to **Parquet** memory format, in order to serve as our streaming source.

Partneships

XILINX
ALLIANCE PROGRAM
MEMBER

aws partner network

awsmarketplace

Alibaba Cloud

Microsoft Partner

intel FPGA Design Solutions Network

Conclusions

- Future Data Center will have to sustain huge amount of network traffic
- However the power consumption will have to remain almost the same
- FPGA acceleration as a promising solution in providing
  - **high throughput**,
  - **low latency** and
  - **energy efficient processing**

Questions?

Thank you!

Christoforos Kachris

chris@inaccel.com