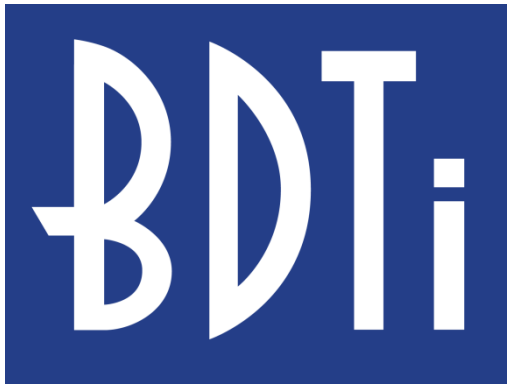

Benchmarking Metrics and Processor Selection for Computer Vision

Workshop: Enabling Computer Vision on ARM

Jeff Bier | May 11, 2015

About BDTI

BDTI provides:

- Best-in-class product development engineering services
  - Emphasis on optimization for performance, cost and power
- Expert, objective benchmarking and evaluation
  - For technology selection, feasibility studies, competitive analysis and proof points
- Licensable benchmark suites and certification services

Focused on:

- Algorithm-intensive applications: vision, video, audio, wireless
- Embedded processors, tools and techniques:
  - CPU, GPU, DSP, FPGA, many-core, etc.

"These guys make a living telling the truth."

– Kevin Morris, Editor in Chief, FPGA Journal

Momentary Fad or Transformative Technology?

Machines are useful mainly to the extent that they **interact with the physical world**

Visual information is the richest source of information about the real world: People, places, and things

Vision is the highest-bandwidth way for machines to obtain info from the real world

Embedded vision can:

- **Boost efficiency and quality**
- **Enhance safety**
- **Simplify usability**
- **Enable innovation**

"Half of the human brain is devoted directly or indirectly to vision."

– Paraphrased from Prof. Mriganka Sur, MIT

The Processor Selection Challenge

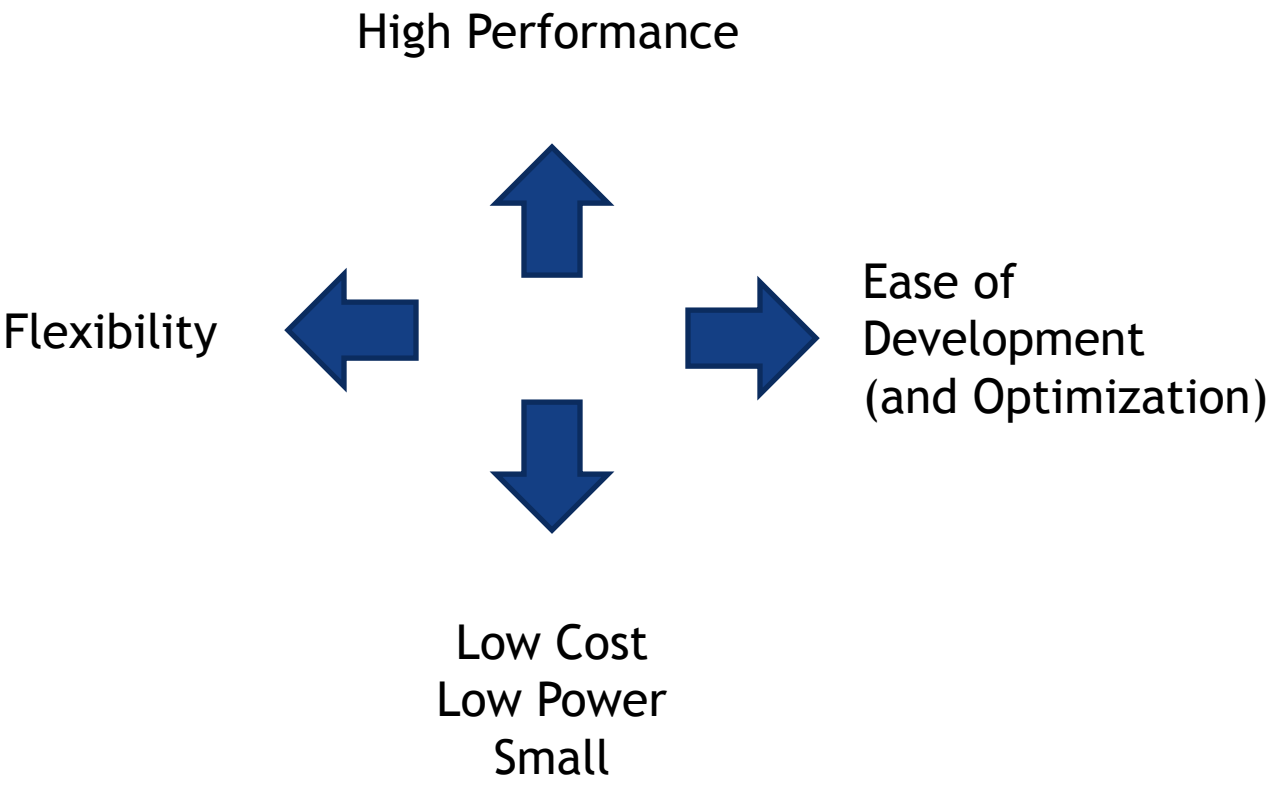

Processing Efficiency vs. Development Effort

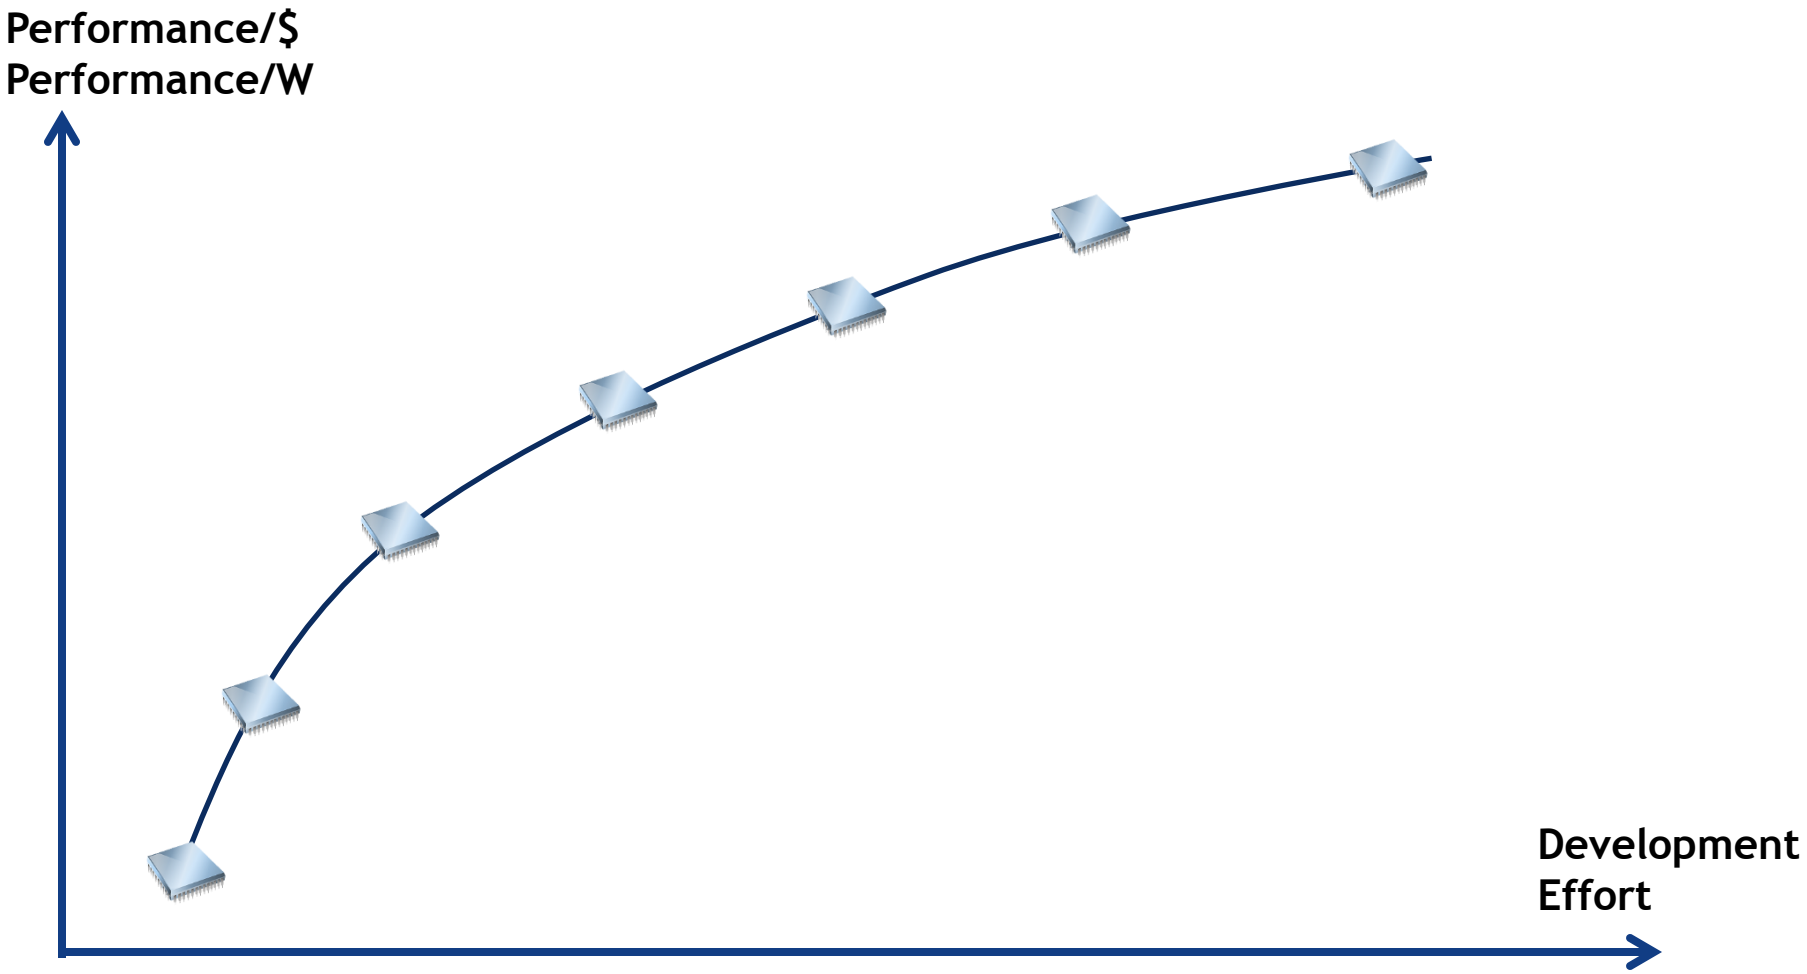

Reality

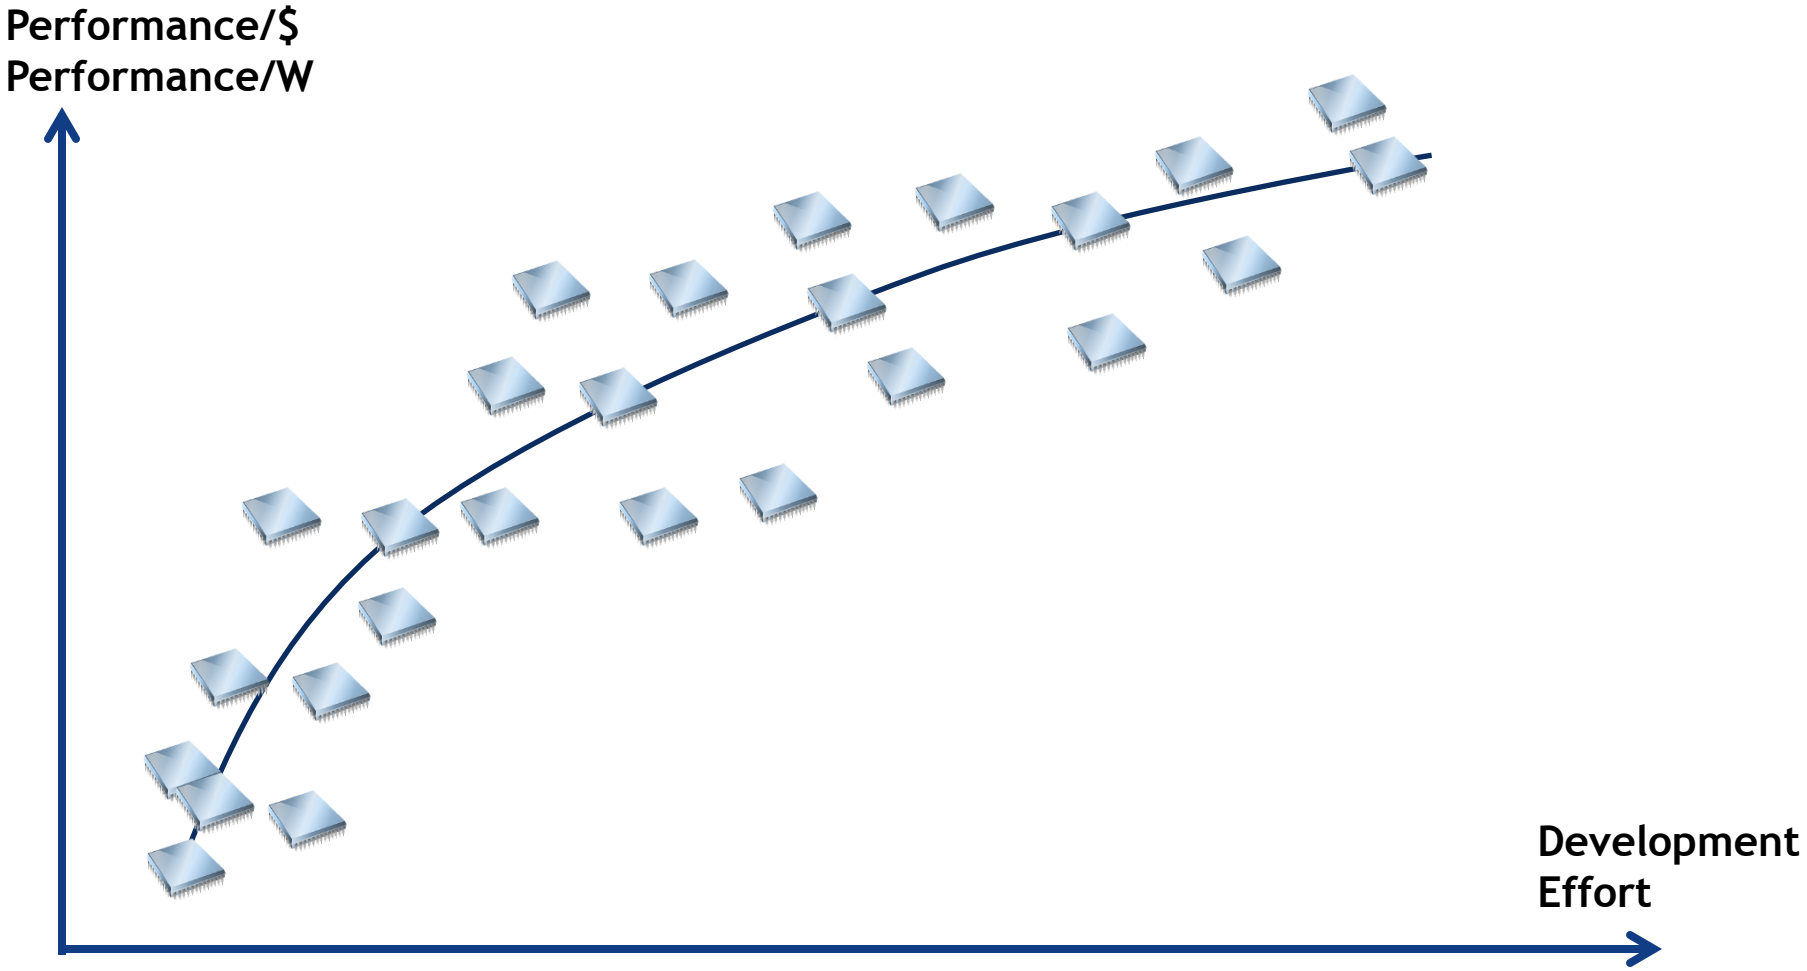

A Rising Tide

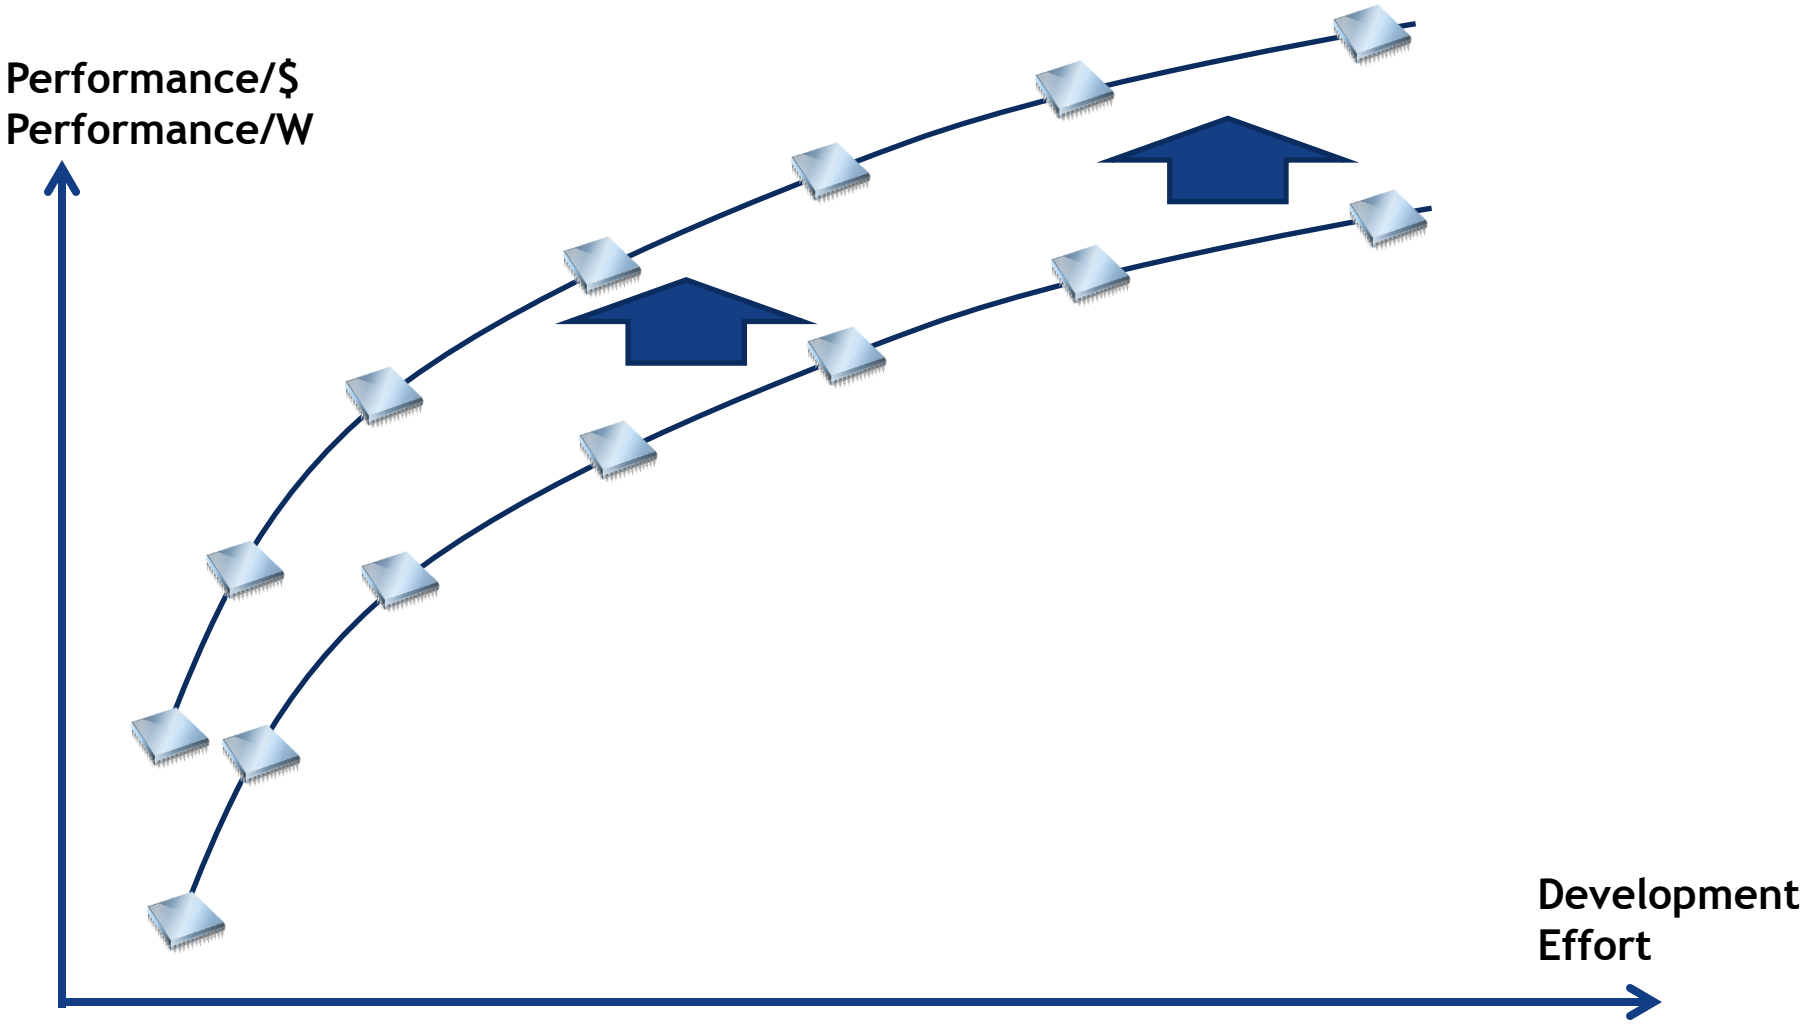

Trend: Heterogeneous Architectures

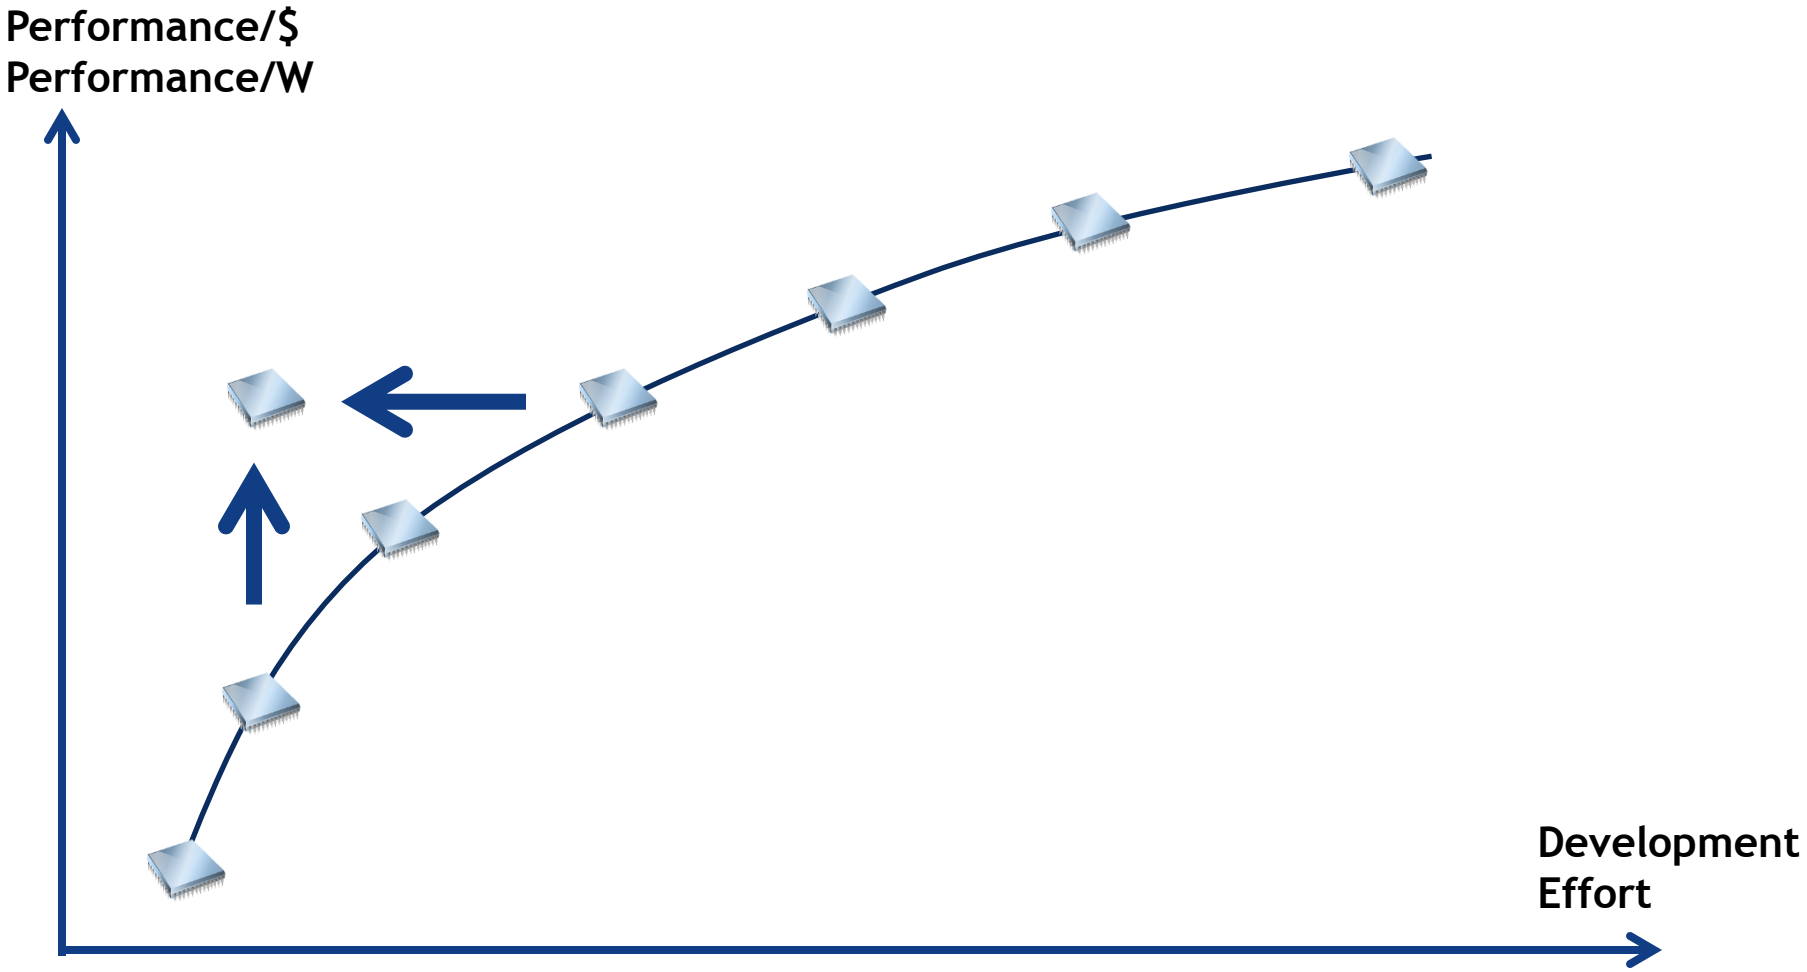

Trend: Mobile Application Processors

- Very heterogeneous processors
- Benefit from huge investments by suppliers, because mobile market is huge
  - Hardware performance, efficiency, integration
  - Application development infrastructure
- Mobile apps have become a primary locus of software development
- APs can be difficult to buy and use for embedded applications
- APs are used in some embedded applications (sometimes in mobile device form, sometimes via a system-on-module)

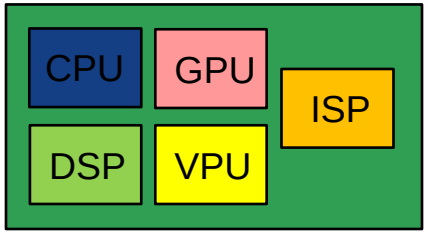

Trend: General-Purpose GPU

- Graphics processing units (GPUs) are massively parallel machines
- Over the past decade, GPUs and their tools have evolved to support non-graphics workloads ("general-purpose GPU" or "GPGPU")
  - Widely used in demanding workstation and data center applications in which data parallelism is abundant
    - E.g., Photoshop, FFTs
- NVIDIA pioneered this concept with CUDA
- Others have joined via OpenCL and RenderScript
- Important recent developments:
  - Now in mobile application processors and embedded processors
  - Expanding support via libraries, code examples, optimized middleware
  - OpenCL support
  - HSA (Heterogeneous System Architecture)

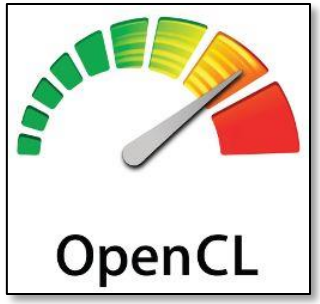

The Benchmark Design Challenge

Development Effort,
Implementation Effort

Breadth
of Applicability

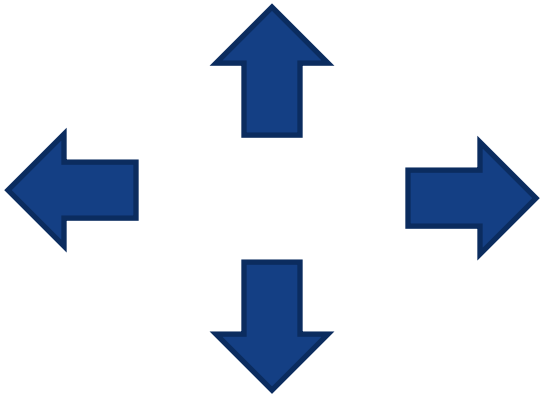

Accuracy,
Relevance

Ease of Use
of Results

Long-Term Trends

- Application, system, software capability and complexity increase
- System functionality becomes more heterogeneous
- Processors become more complex and heterogeneous
- Proprietary algorithms are key differentiators
- Programming approaches, languages, tools are changing more rapidly

➔ **Collectively, these trends create enormous challenges for benchmarking**

Addressing Complexity and Heterogeneity

As systems, applications and processors become more complex, it becomes difficult to approximate the performance of the whole from that of the parts

Kernel benchmarks are of diminishing value for evaluating systems

The natural solution is to benchmark at the system level, using full applications

This approach has its own challenges:

- Creating realistic synthetic full-application benchmarks is very costly
- Access to the real application code is often impossible
- Porting and optimizing the real application code to multiple system architectures is very costly

Vision IP Core Benchmark Results (Normalized)

Vendor	Max Clk Freq. (MHz)	Dyn. Power	HOG		LK-OF		Face Detection		Resize
			Speed	Energy	Speed	Energy	Speed	Energy	Speed
Vendor 1	900	0.38	1.00	0.11	0.63	0.26	1.00	0.13	1.00
Vendor 2	1000	0.25	0.90	0.08	1.00	0.11	0.86	0.09	0.92
Vendor 3	500	1.00	0.30	1.00	0.42	1.00	0.34	1.00	0.33

[1] *Entries are normalized*

Lessons Learned

- Different processors are better suited for different applications
  - Applications are diverse, and so are processors

- Selecting a processor is a multi-dimensional optimization problem with incomplete data
  - Allow time to evaluate your options
- Macro trend: Increasing industry investment in processors for vision
  - Including mobile application processors
- You can benefit by riding on the coattails of high-volume applications
  But, progress is uneven and unpredictable ➔ Risk for users

- A processor is only as good as its tools, libraries, etc.

Embedded Vision Summit: May 12, 2015 — Santa Clara, CA

The only industry event focused on enabling creation of systems and apps that "see"

- *"Good balance of technical content and application-driven examples."*

Ren Wu, Baidu

Embedded Vision Summit 2015 highlights:

- Inspiring keynotes by leading innovators
- Three tracks of practically-oriented technical and business talks
- Demos of the latest apps and technologies
- In-depth pre- and post-Summit workshops

Mike Aldred, Dyson

Registration open at www.EmbeddedVisionSummit.com

Empowering Product Creators to Harness Embedded Vision

The Embedded Vision Alliance (www.Embedded-Vision.com) is a partnership of 50 leading embedded vision technology suppliers

Mission: Inspire and empower product creators to incorporate visual intelligence into their products

- The Alliance provides low-cost, high-quality technical educational resources for engineers
- Member companies position themselves as leaders to thousands of product creators via the Alliance web site and conferences
- Members meet quarterly to develop business partnerships and gain insights into markets and technologies
- We secure frequent press coverage on embedded vision topics, gaining exposure for our members as thought leaders