ORACLE

Innovating Beyond Moore's Law

Craig Stephen
Vice President, R&D, Oracle Labs

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle’s products remains at the sole discretion of Oracle.

The following is intended to outline our general product
for information purposes only, and may not be incorpora
It is not a commitment to deliver any material, code, or f
not be relied upon in making purchasing decisions. The
and timing of any features or functionality described for
remains at the sole discretion of Oracle.

No Kidding !

Agenda

- Research and Innovation at Oracle
- The End of Moore's Law?
- Challenges Beyond Moore's Law
- Opportunities Beyond Moore's Law
- Your Future Computing Infrastructure

Agenda

- **Research and Innovation at Oracle**
- The End of Moore's Law?
- Challenges Beyond Moore's Law
- Opportunities Beyond Moore's Law
- Your Future Computing Infrastructure

On Innovation

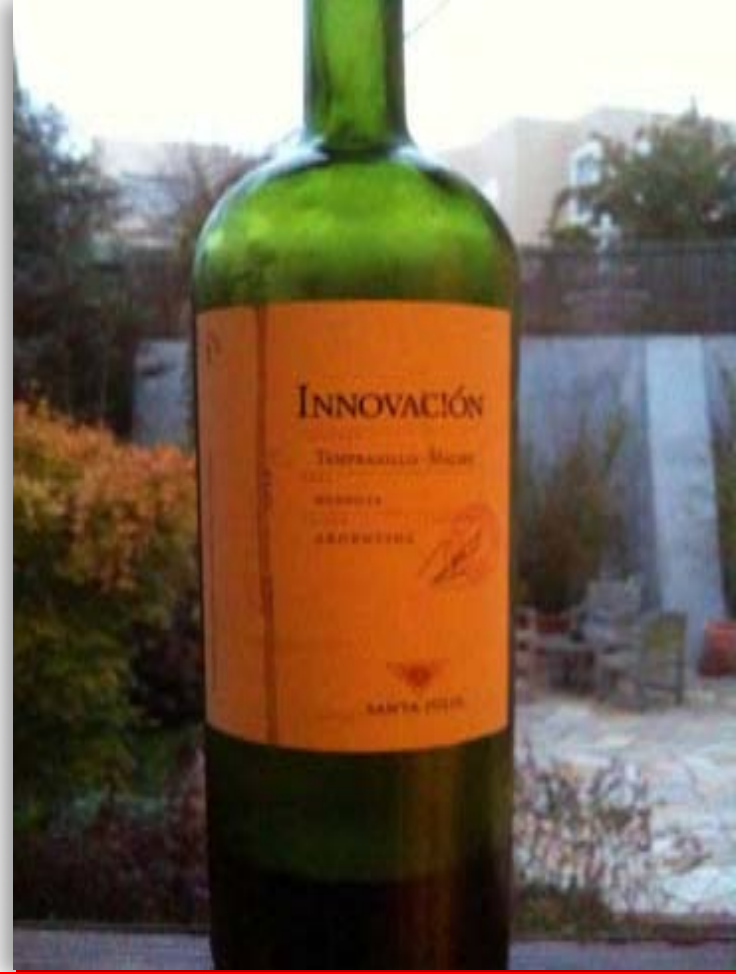

On Innovation at Oracle

Oracle Engineering Culture

- Practical
- Focused on customer value
- Needs predictable timeframes

- Continuous, sustaining innovation baked into process
- **Innovation Engines**

Innovation at Oracle

An Innovation Engine Example

- Database development requires strict cycle times
- If a development project late…
  - the release train won't wait…
  - but another train will come soon
- If a feature is still a little experimental…
  - It can be a runtime option for beta (so long as it doesn't break anything)
- ***Fail fast***, ***fail cheap***, and ***iterate,*** but ***never stop the train***

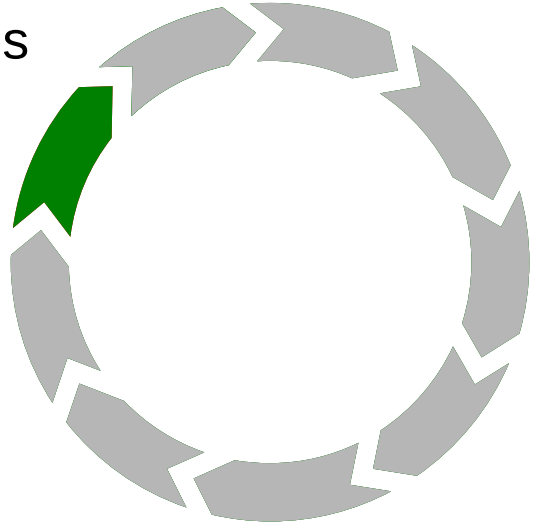

Oracle Labs

Another path to innovation

Oracle Labs

Another path to innovation

Oracle Labs

Another path to innovation

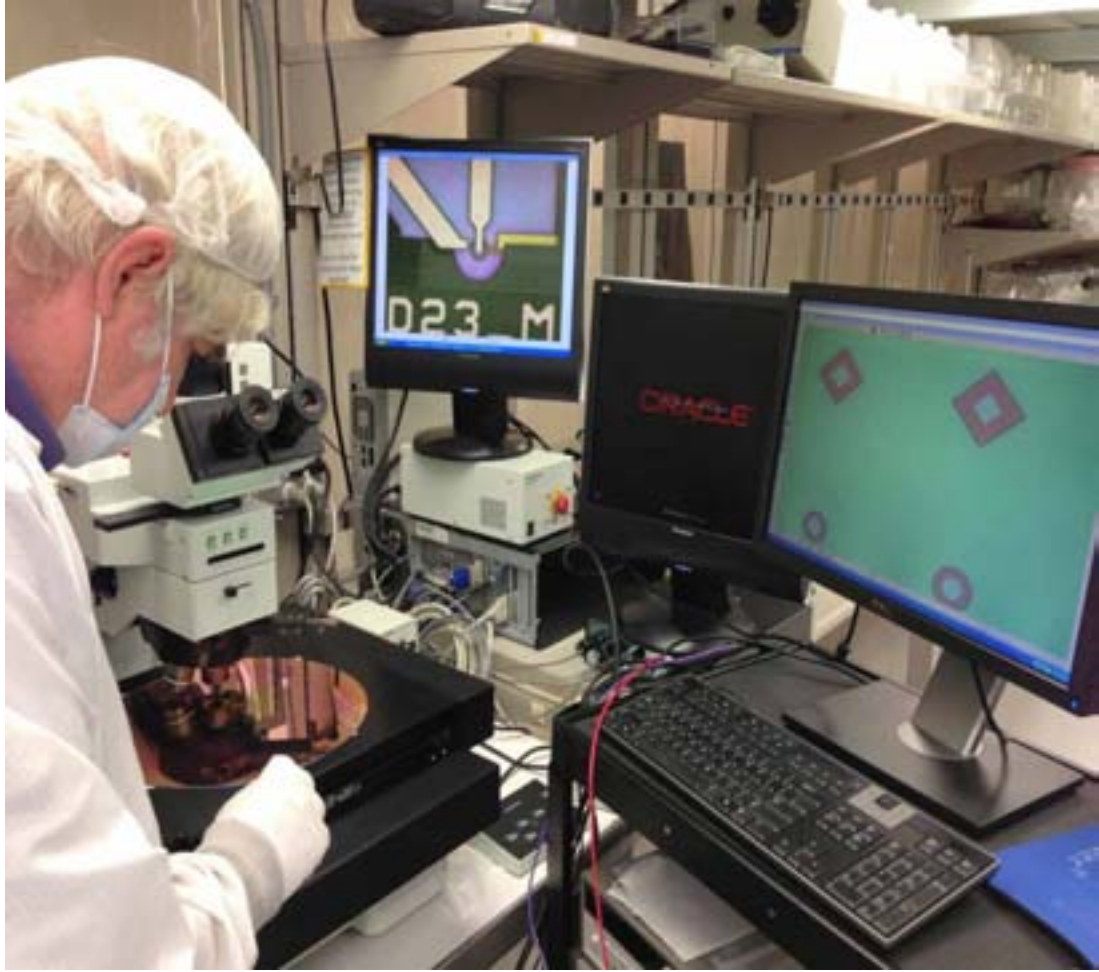

"The mission of the Labs at Oracle is straightforward – to identify, explore, and transfer new technologies that have the potential to substantially advance Oracle's business."

Edward Screven

Chief Corporate Architect

Oracle Labs

“ Identify, explore, and transfer new technologies
that have the potential to substantially advance Oracle’s business ”

- Clear connection to Oracle’s business
- Plausible technology transfer path
- Real risk of failure
- Time horizons may not be compatible with Product Development
- Opportunity may be out of scope for existing product organizations

Oracle Labs Provides Innovation Options

Labs Project Criteria

- Connected to business
- Path to tech transfer
- Not addressable by Engineering

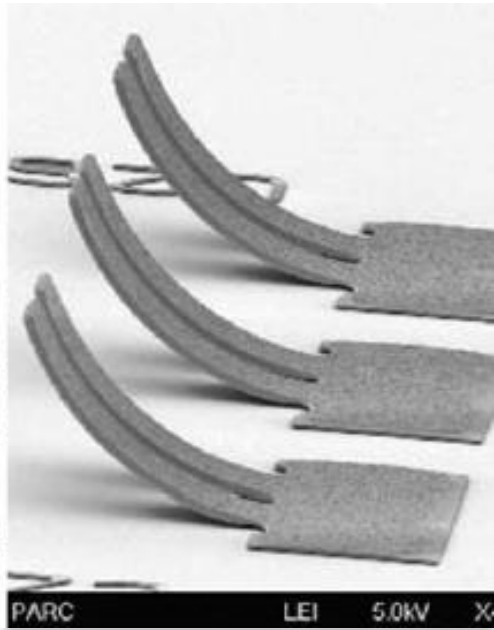

Agenda

- Research and Innovation at Oracle
- **The End of Moore's Law?**
- Challenges Beyond Moore's Law
- Opportunities Beyond Moore's Law
- Your Future Computing Infrastructure

"Everybody talks about the weather,
but no one ever does anything about it."

Mark Twain

by popular attribution

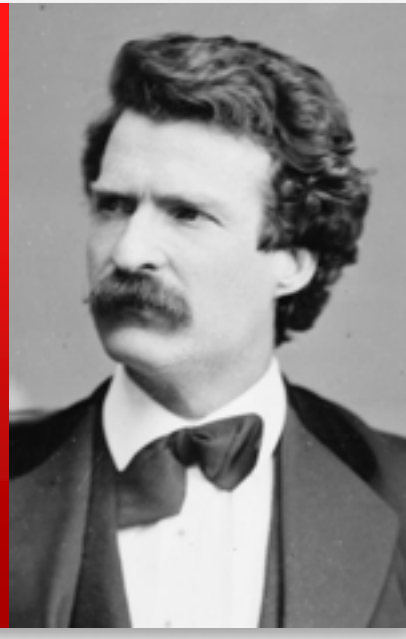

“The coldest winter I ever spent was a summer in San Francisco.”

Mark Twain

by popular attribution

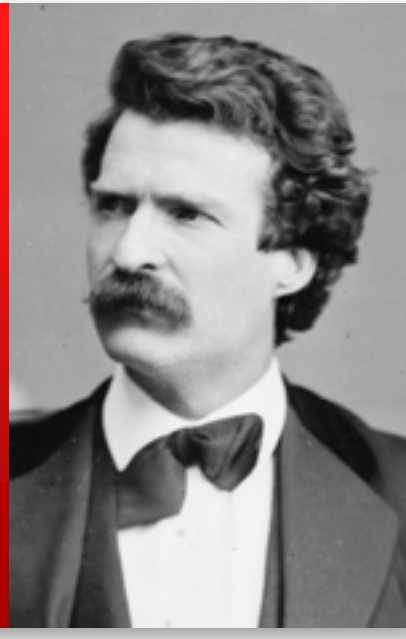

What Does Moore's Law Mean, Exactly?

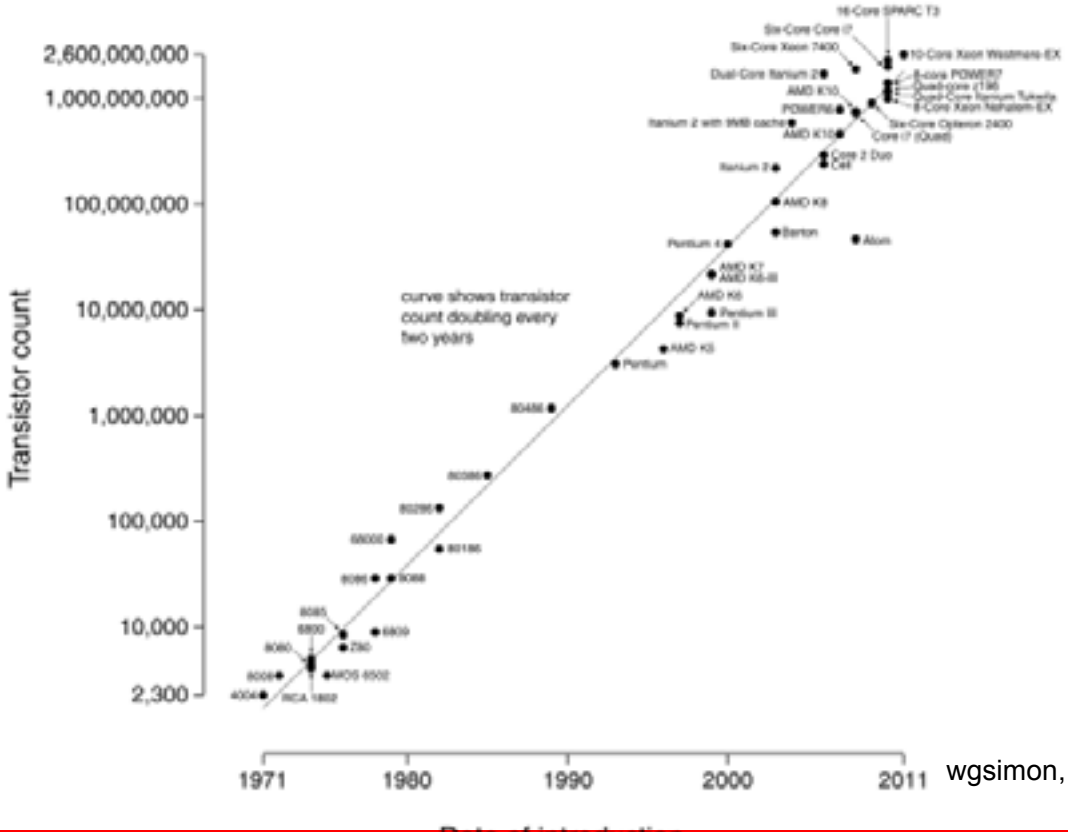

wgsimon, wikipedia

- 2x transistors / unit area
- …every two years
- We're still on track

"Prediction is very difficult, especially if it's about the future."

Neils Bohr

How Long Do We Have?

We can set some limits on information density

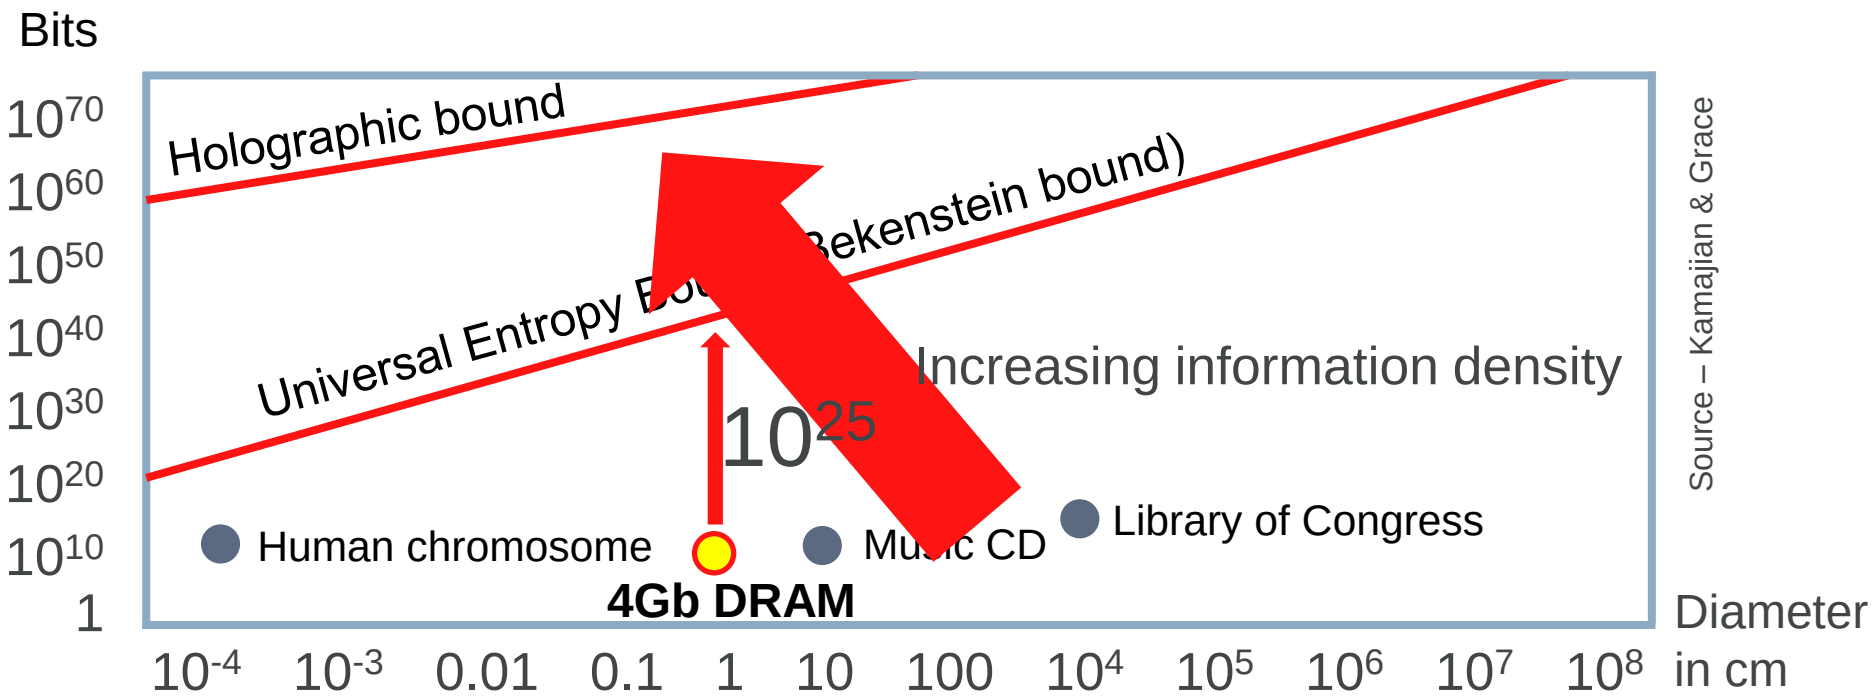

How Long Do We Have?

We have room to improve by 1,000,000,000,000,000,000,000,000 x !

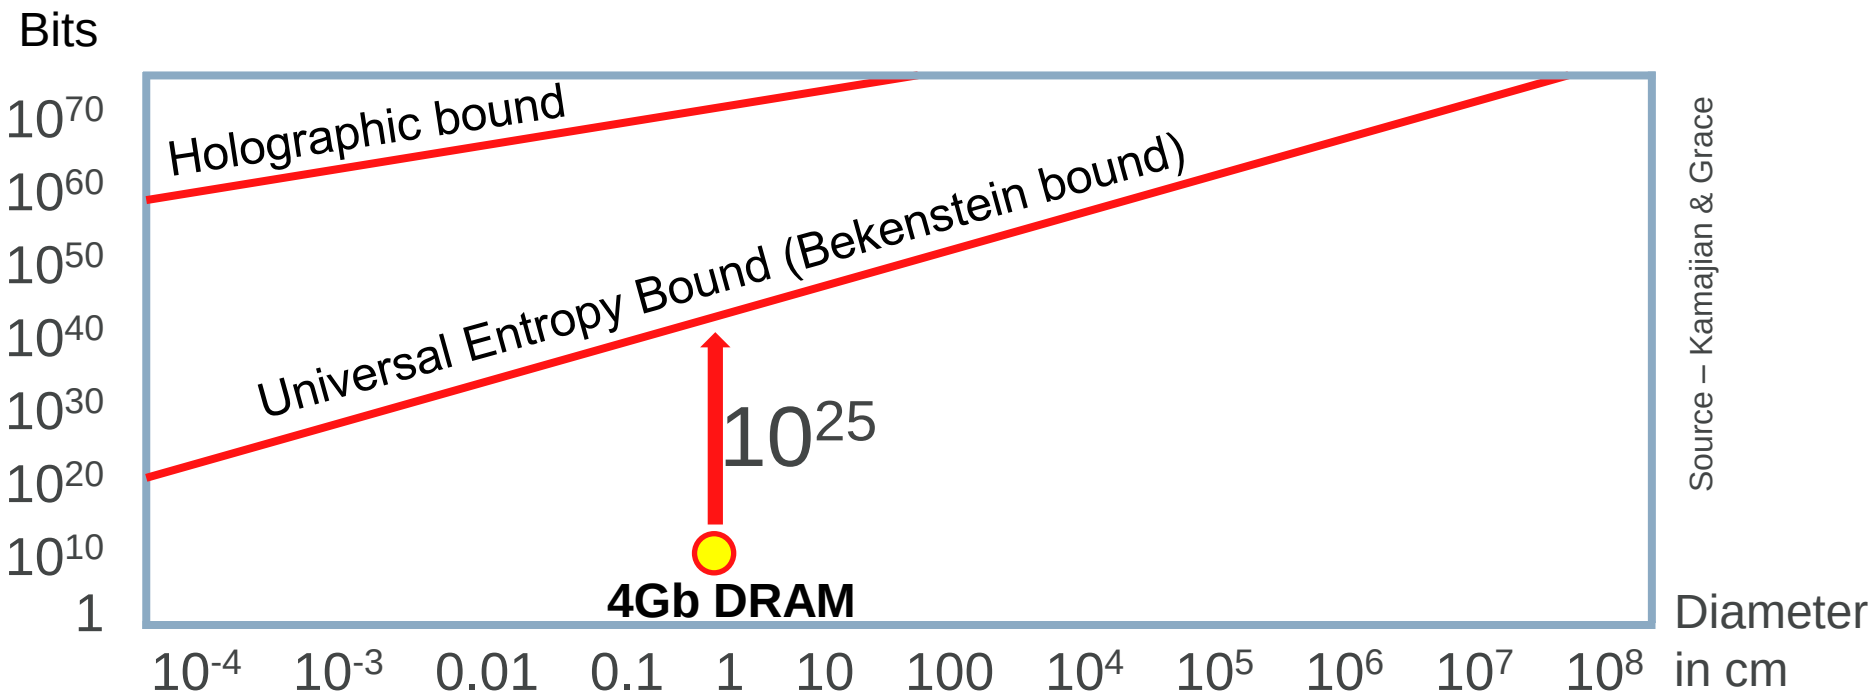

How Long Do We Have?

We can set some more practical limits

- Semiconductor physics & process engineering

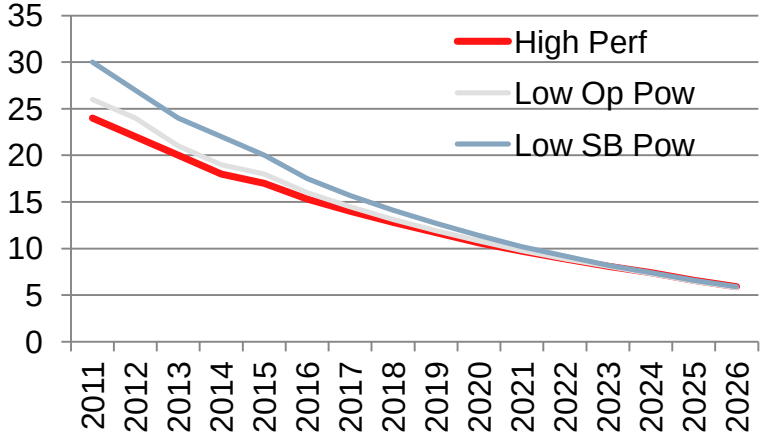

5.5nm
20x density

- Economics

Intel Fab 42	**$5.2** B, 2011
TSMC Fab 15	**$9.3** B, 2010
GlobalFoundries Fab 8	**$4.6** B, 2009

Source – International Technology Roadmap for Semiconductors

Is it the End of Moore's Law?

Wrong question?

- Some runway left…
- Increasingly expensive…
- The challenges transcend Moore's Law

Agenda

- Research and Innovation at Oracle
- The End of Moore's Law?
- **Challenges Beyond Moore's Law**
- Opportunities Beyond Moore's Law
- Your Future Computing Infrastructure

Some Challenges Transcend Moore's Law

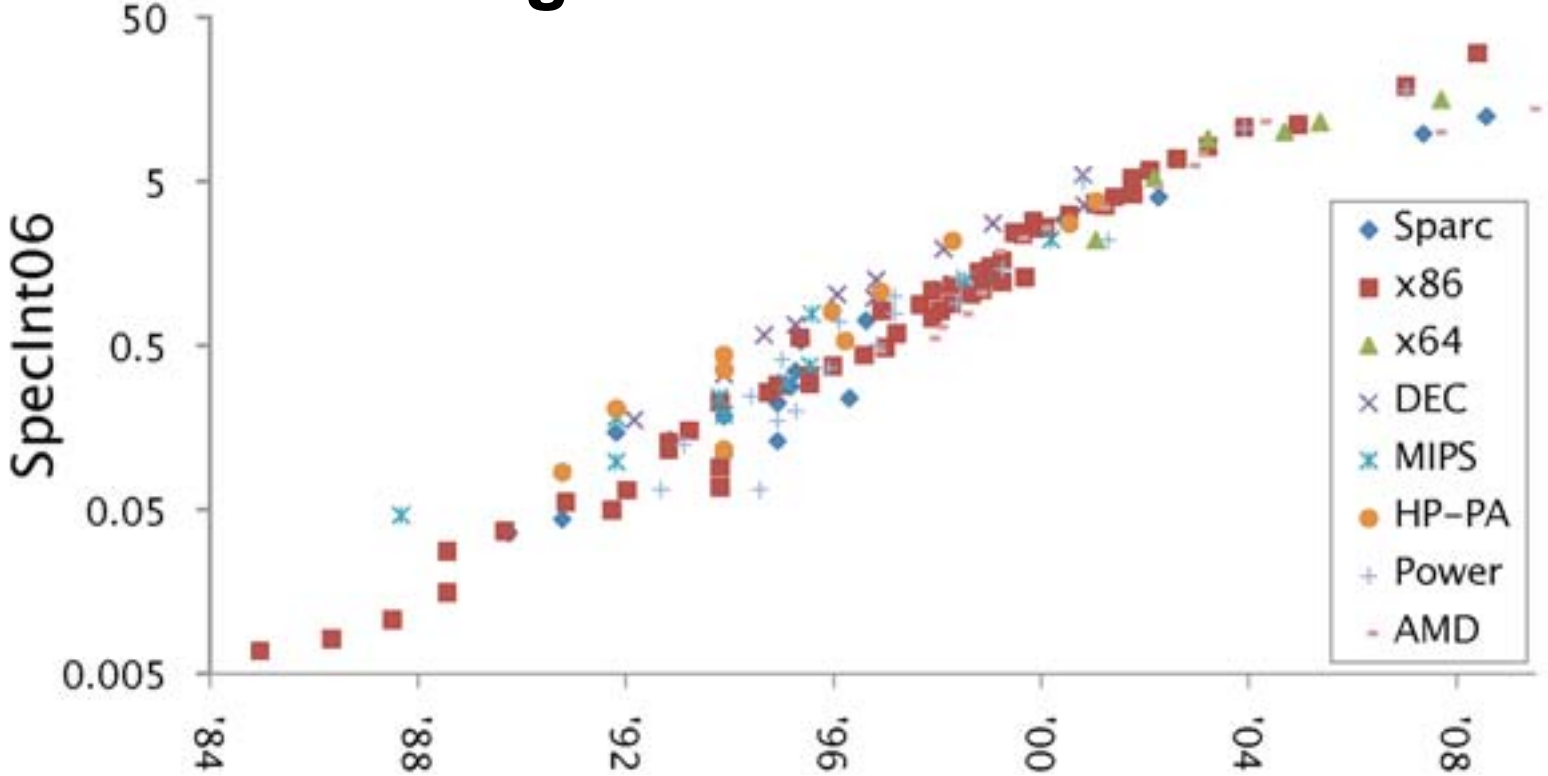

Ho et al
IEEE Design and Test, 2010

35% CAGR in computer performance – that's impressive!

Where Did Performance Growth Come From?

Faster transistors… and more of them

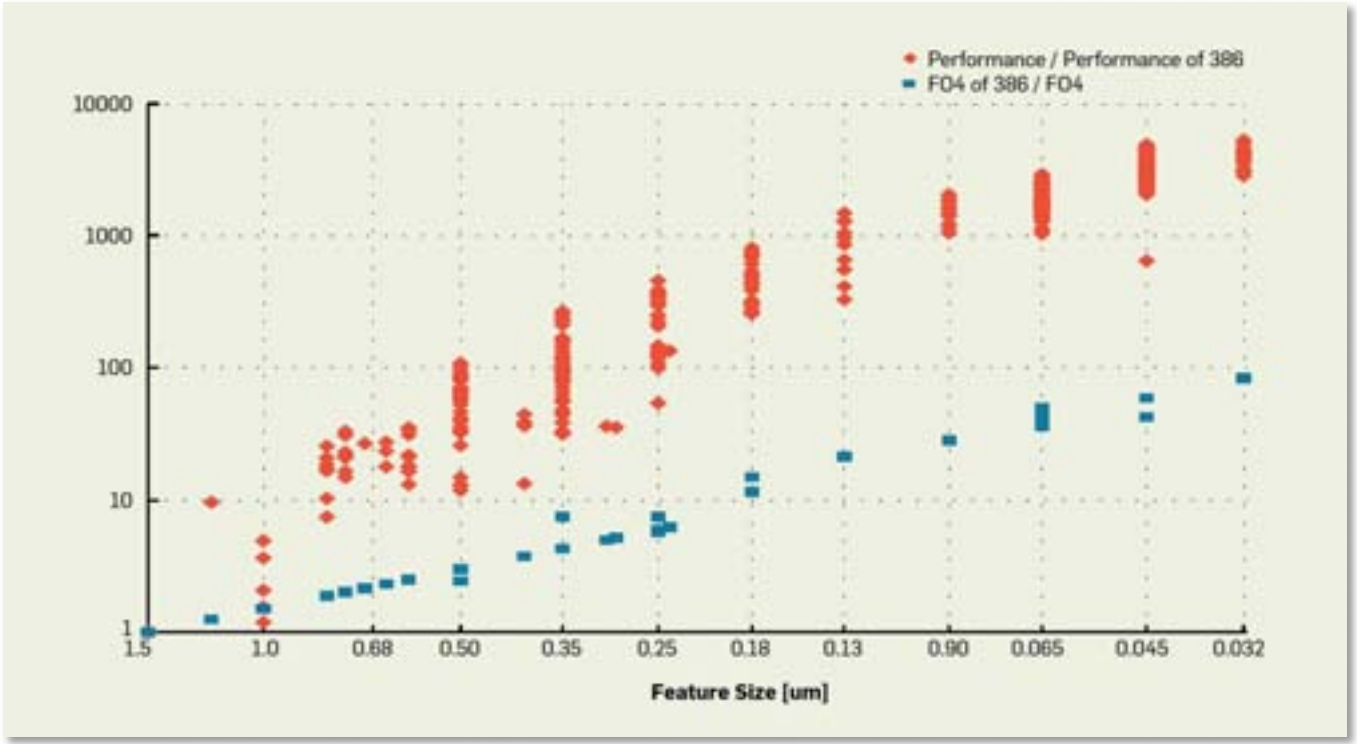

Danowitz et al.
ACM QUEUE

Performance, defined

Performance = instructions/cycle * cycles/second

Performance, defined

Performance = instructions/cycle * frequency

Performance, defined

Performance = instructions/cycle * frequency

But growth here is stalled

Frequency Scaling Has Tailed Off….

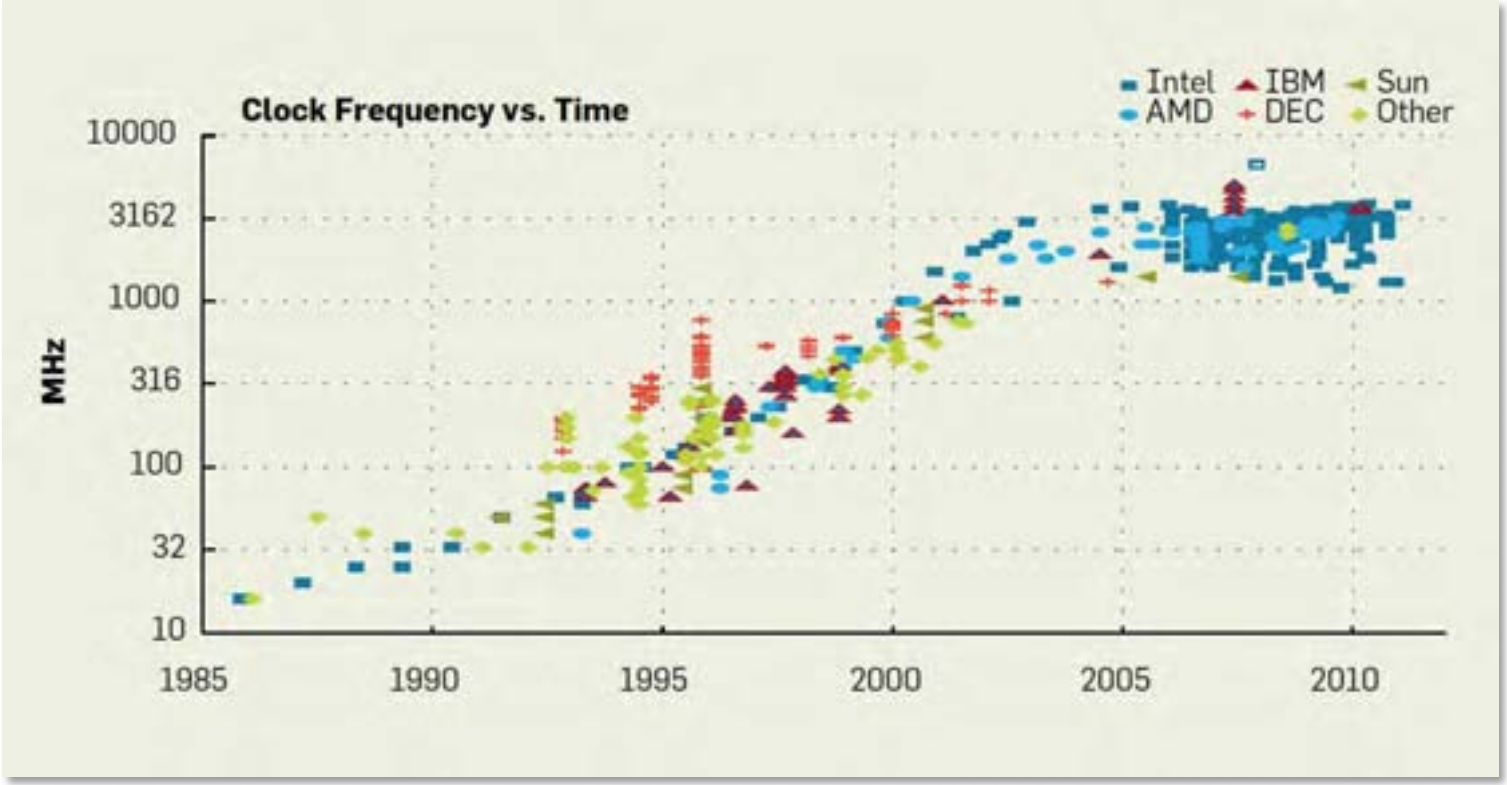

Danowitz et al. ACM QUEUE, 2012

Frequency Scaling Has Tailed Off….

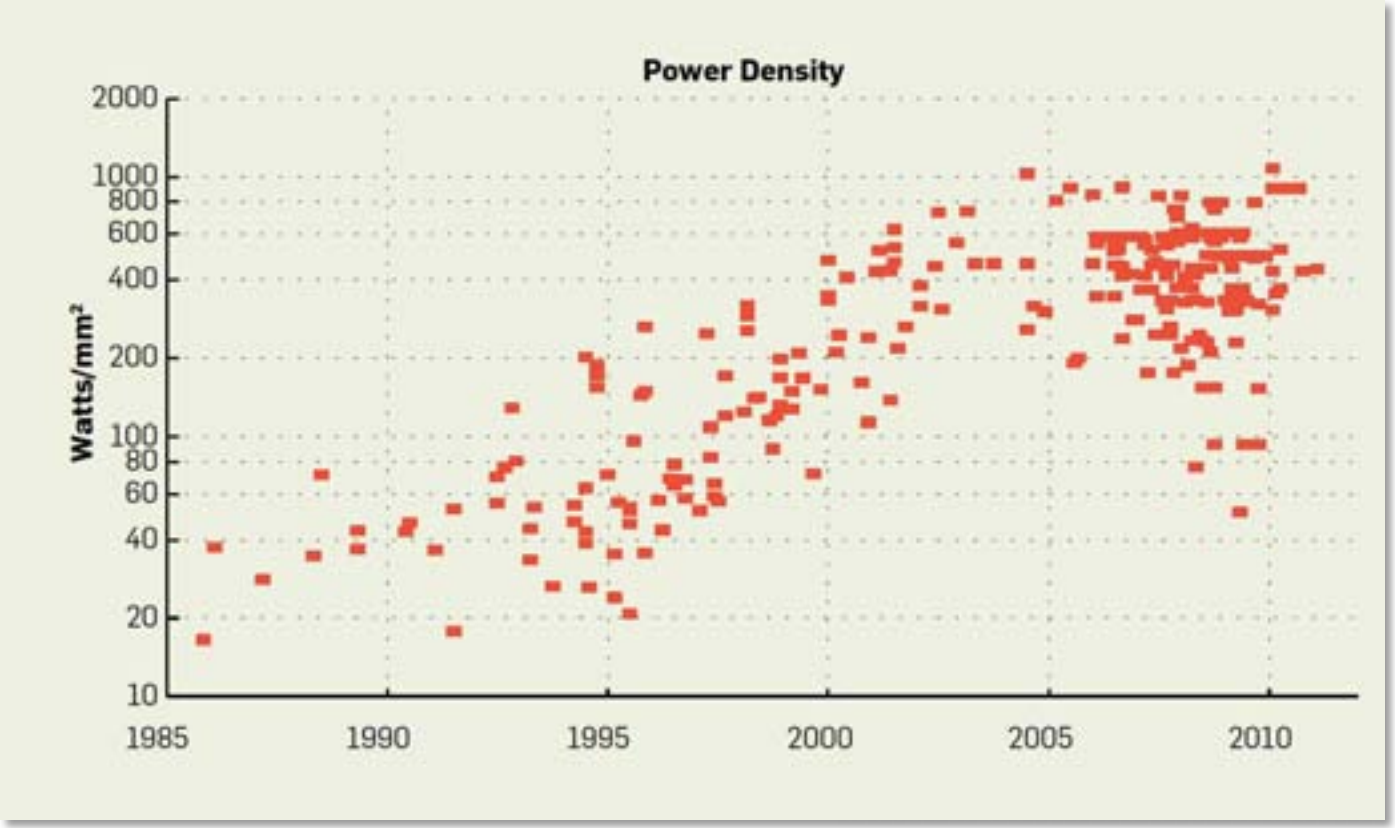

Danowitz et al.
ACM QUEUE
2012

Performance, defined

Performance = instructions/cycle * frequency

If growth here is stalled…

Performance, defined

Performance = instructions/cycle * frequency

We need to look here!

Performance, defined

Performance = **parallelism** * frequency

We need to look here!

Industry's solution - add more cores

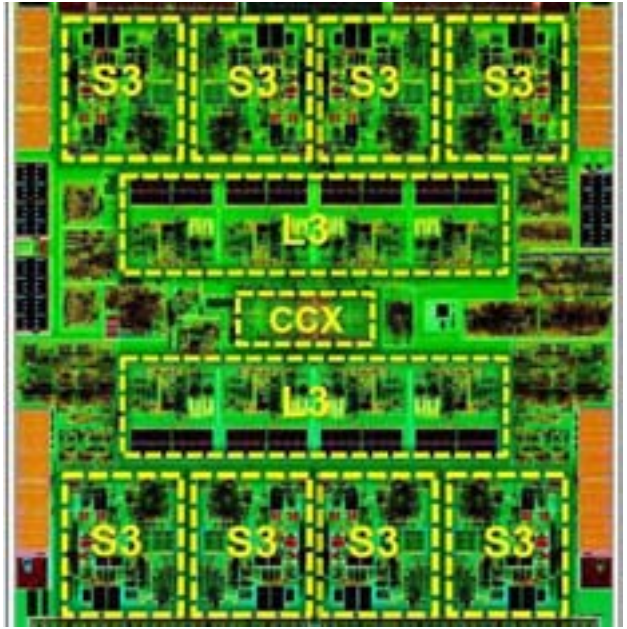

Regardless of Moore's Law

Daunting Challenges Remain

- Processor performance scaling
- Interconnect
- Programming

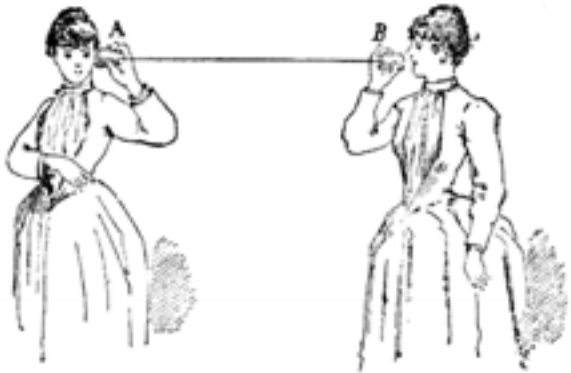

Fig. 76. Trådtelefon.

For (i = 1 to n)
{…}

Agenda

- Research and Innovation at Oracle
- The End of Moore's Law?
- Challenges Beyond Moore's Law
- **Opportunities Beyond Moore's Law**
- Your Future Computing Infrastructure

Four Courses of Action

- Make do with less – buy time from Moore
  1. Use compute cycles wisely
  2. Make every transistor count
- Build more – exploit parallelism, with or without Moore
  3. Build massive scaleout compute
  4. Harness massive scaleout compute

Use Cycles Wisely

Performance, Thanks to Moore, Again…

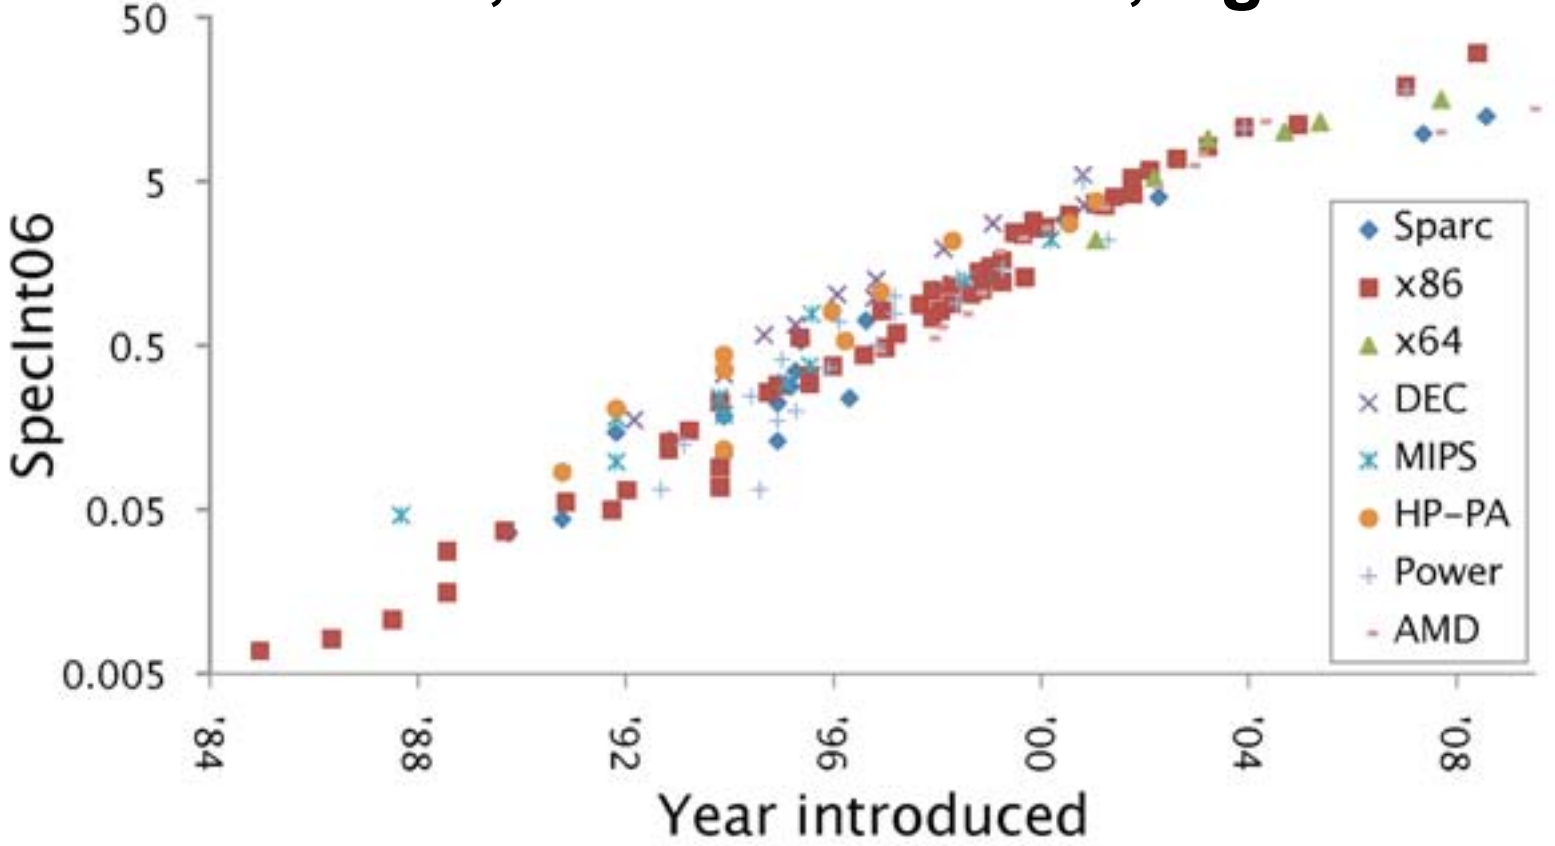

Performance, Thanks to Moore, Again…

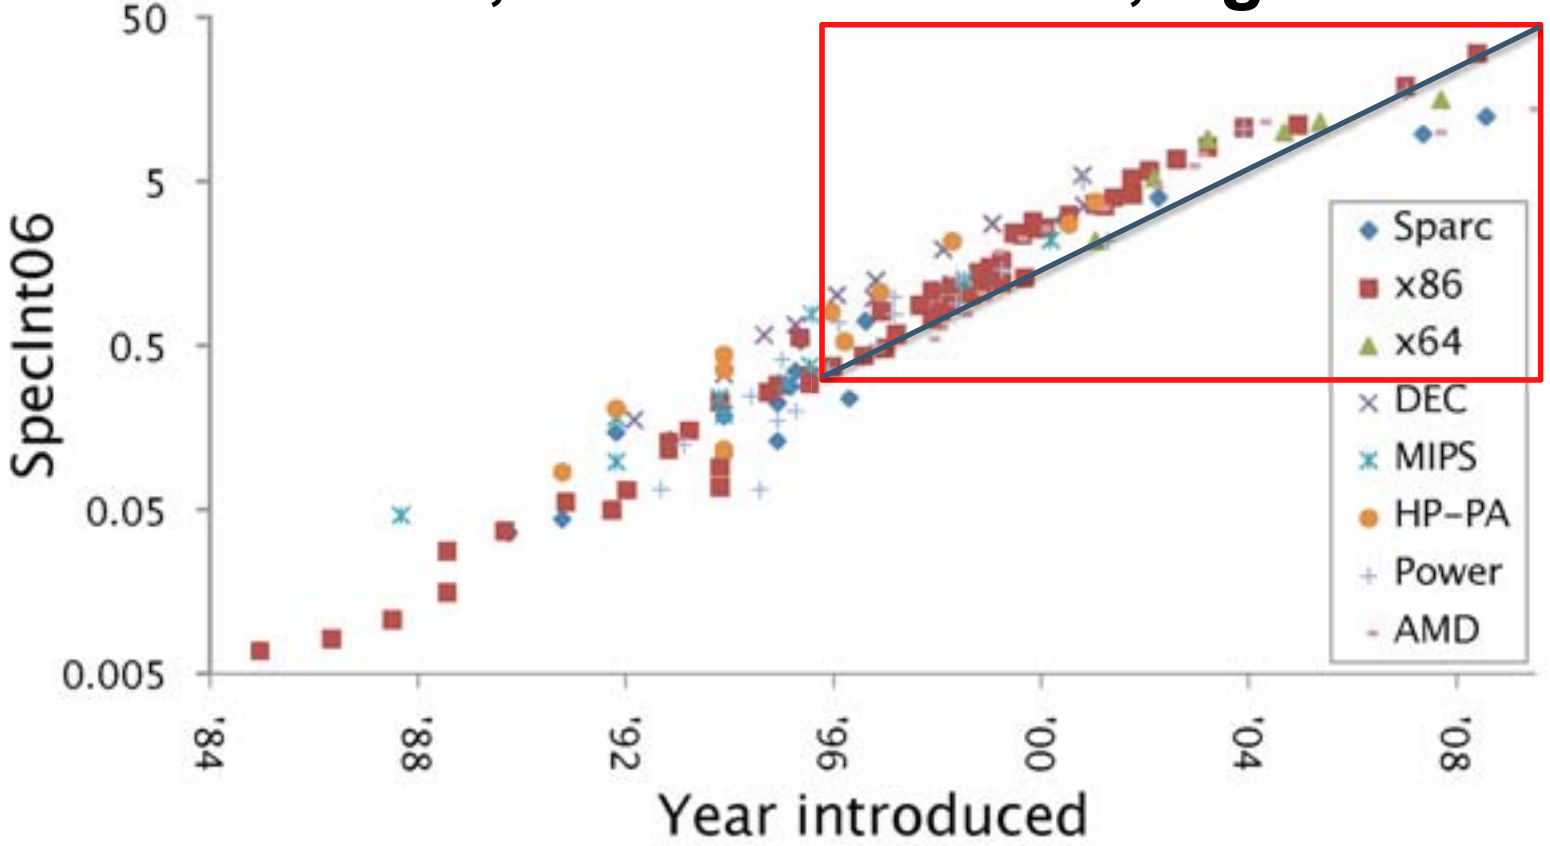

Performance Comparison

What's going on here?

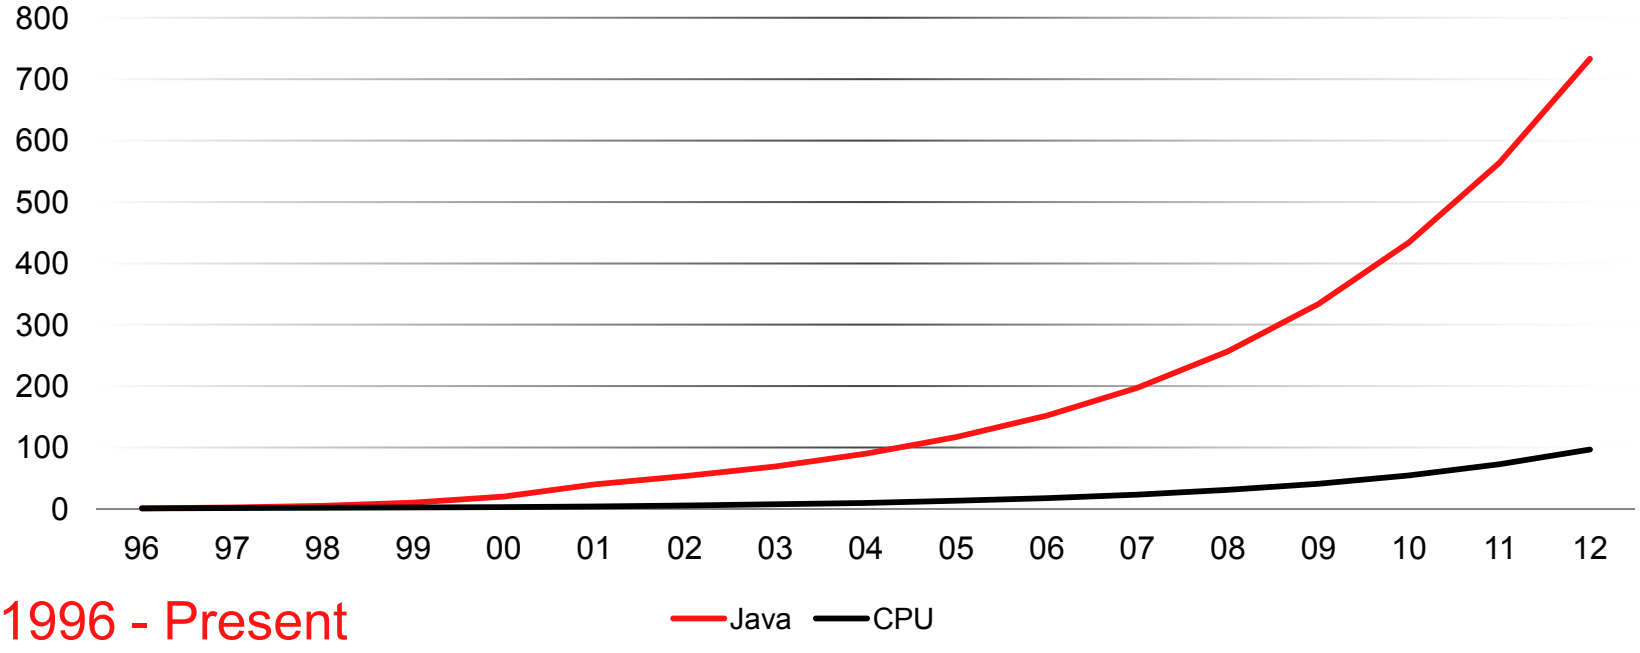

1996 - Present

10x Improvement Beyond Moore

Genius application developers not required!

Von Neumann

Turing

Gödel

10x Improvement Beyond Moore

Superluminary neutrinos not required…

10x Improvement Beyond Moore

Everyone benefits from just a *few* genius developers

A few smart people can massively improve the state-of-the-art in program execution.

10x

Performance Improvements Beyond Moore

There is an opportunity here to recover those CPU cycles…

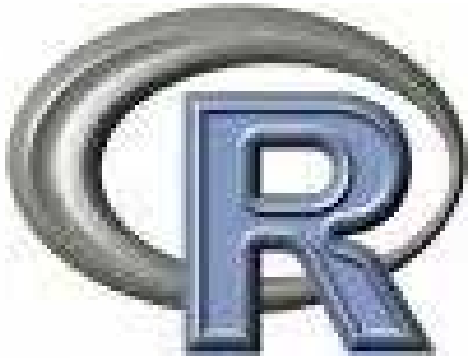

Performance Improvements Beyond Moore

10-100x… 10x is *seven years* of Moore's Law

Improve Compute Cycle Efficiency

Programming Language Development

- 10x Java Runtime Improvement
- Similar opportunities with other languages
- Huge potential gains in aggregate

Make Every Transistor Count

Specialization

Making every transistor count

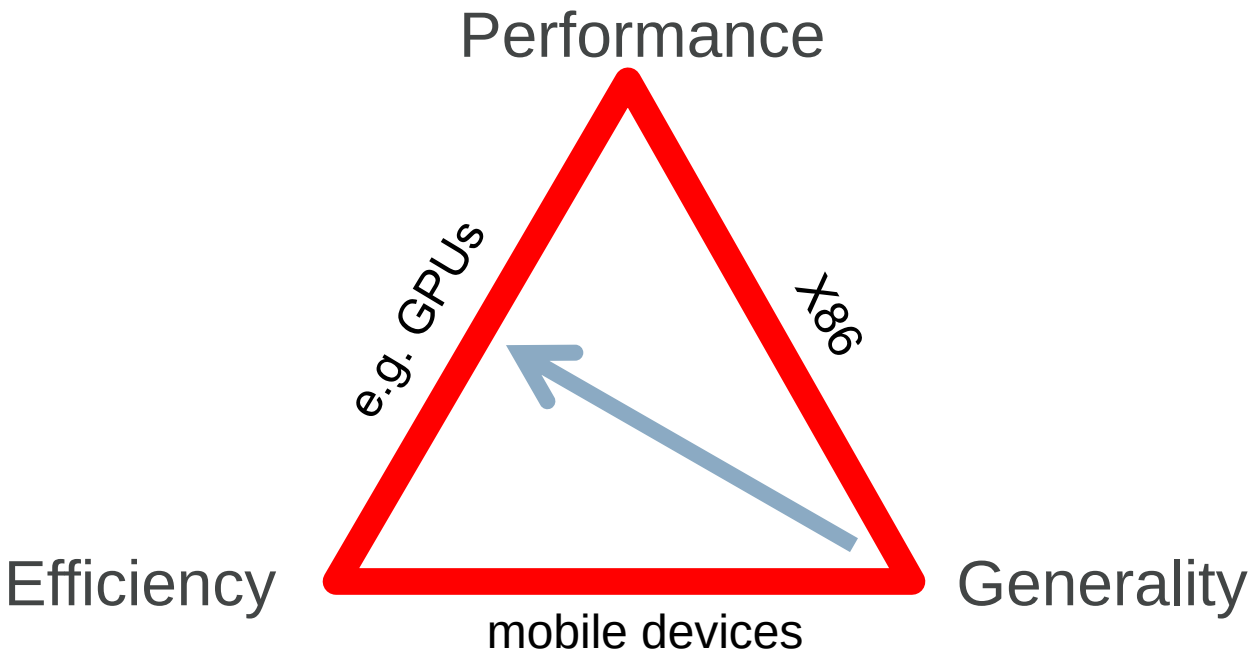

Specialization

Making every transistor count

- “GPGPU” trend – general computation using GPU shader pipeline
- Optimized for computationally-intense workloads
- Enterprise workloads – RDBMS – can be memory-bandwidth intense

- So we prefer… a screwdriver

Driver: fcb981@en.wikipedia
Screw: haragayato

Engineered Systems

Making every transistor count

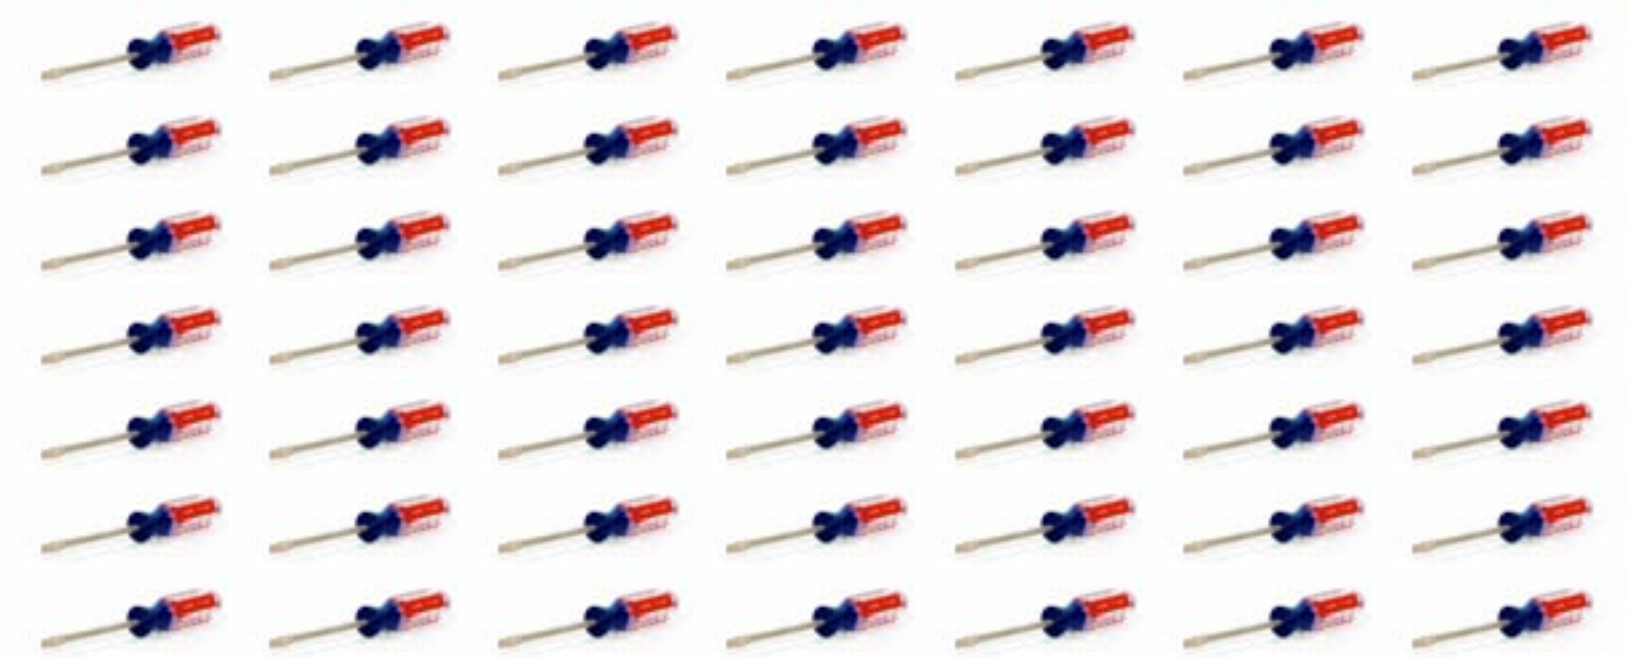

Making every transistor count for the Enterprise

Hardware Specialization

- Hardware / Software co-design
- Oracle uniquely positioned

Build Massive Scaleout Compute

Build Massive Scaleout Compute

Recall…

Performance = **parallelism** * frequency

We need to look here!

Build Massive Scaleout Compute

SPARC T4

- 3GHz
- Eight Cores
- Eight threads / core
- Out-of-order execution
- Single-pipeline, dual-issue

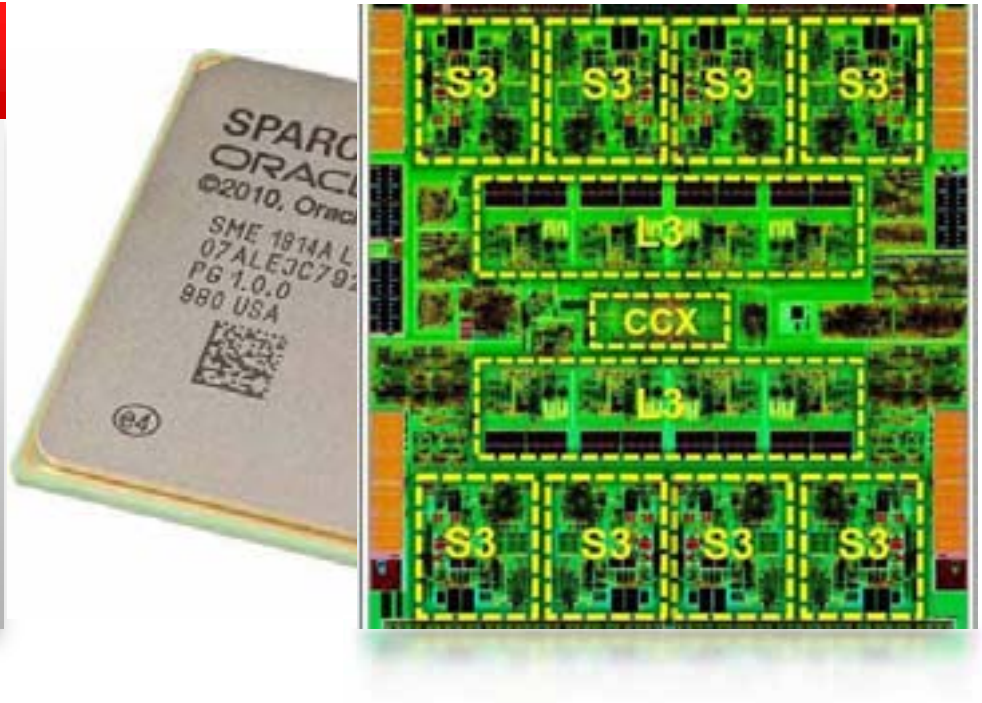

Build Massive Scaleout Compute

- Recall… there are die-size limits

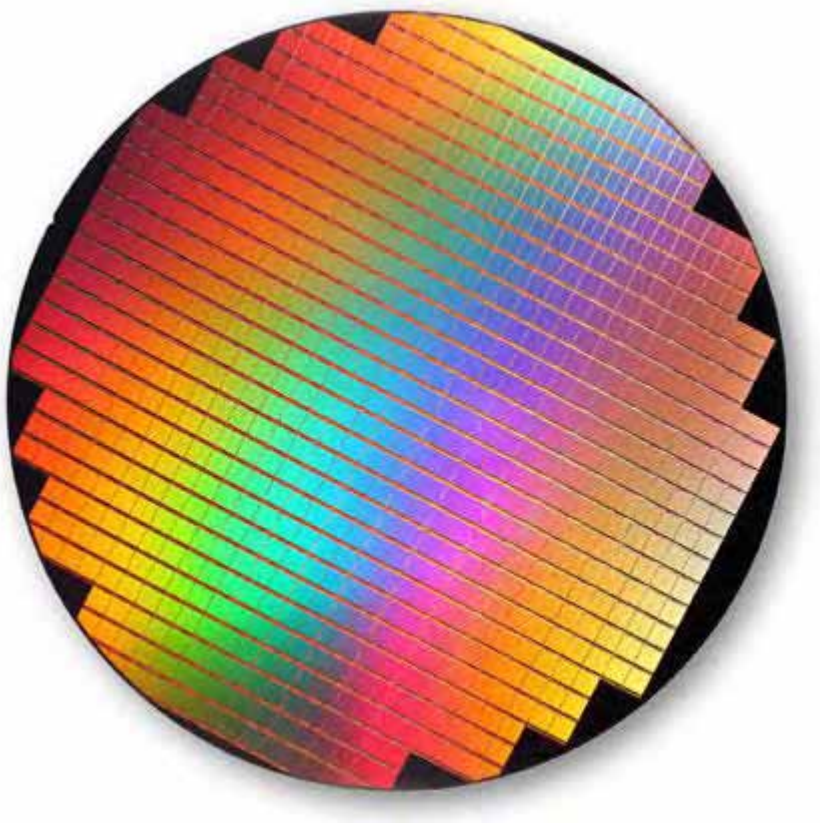

How Do We Continue to Advance Integration?

This is probably not the way to do it

- Cable latencies
- Signaling power
- Management

Image from Al Davis, Utah

Build Massive Scale Compute

Communication costs are significant

- Data centers are power limited
  - Much of that power goes to I/O
  - At least 10 mW per Gb/s of data comms
  - 1 b/s of external traffic spawns 0.1 Gb/s (1 mW) of internal traffic
  - So 10 XB/month is 80 PW!
- Costs are real
  - 2 MW costs $1M in OpEx
  - Need dedicated substations

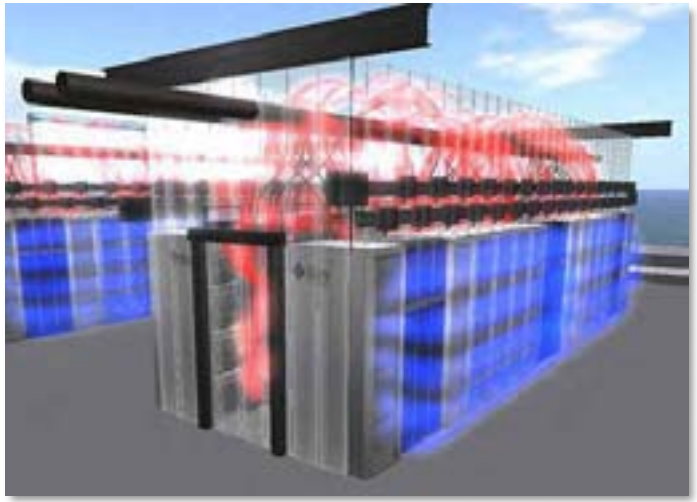

ORACLE

Build Massive Scale Compute

Towards extreme integration through chip aggregation

- Build upwards (“3-D”)
- Build outwards

- Requires low-energy, high-bandwidth interchip I/O

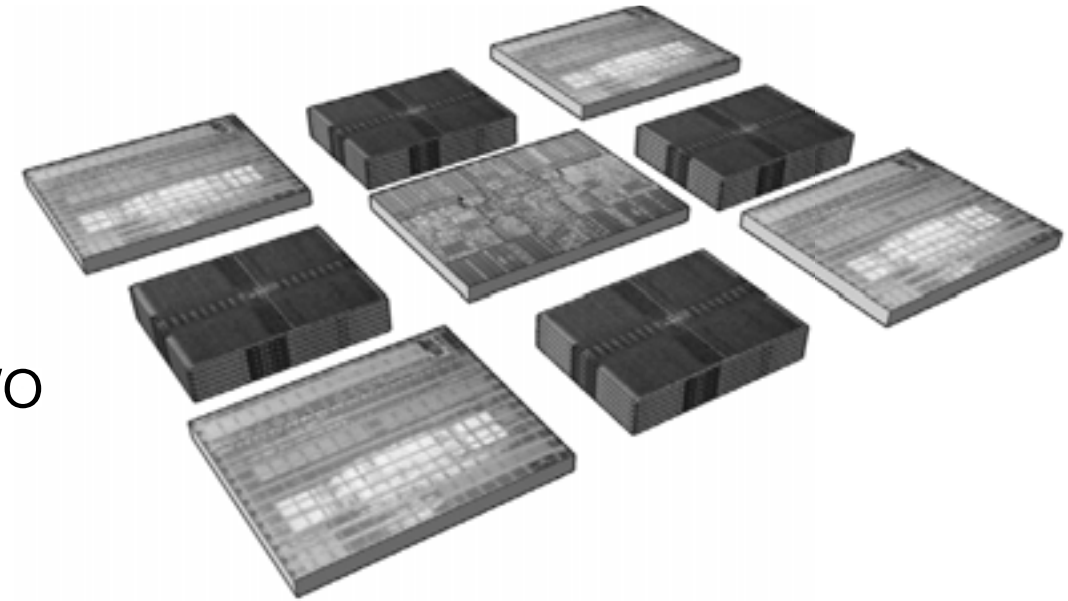

Build Massive Scale Compute

Low-energy, high-bandwidth interchip I/O

- Electrical connectivity increasingly challenged
- Optics reigns where distance * BW exceeds 100 Gb*m

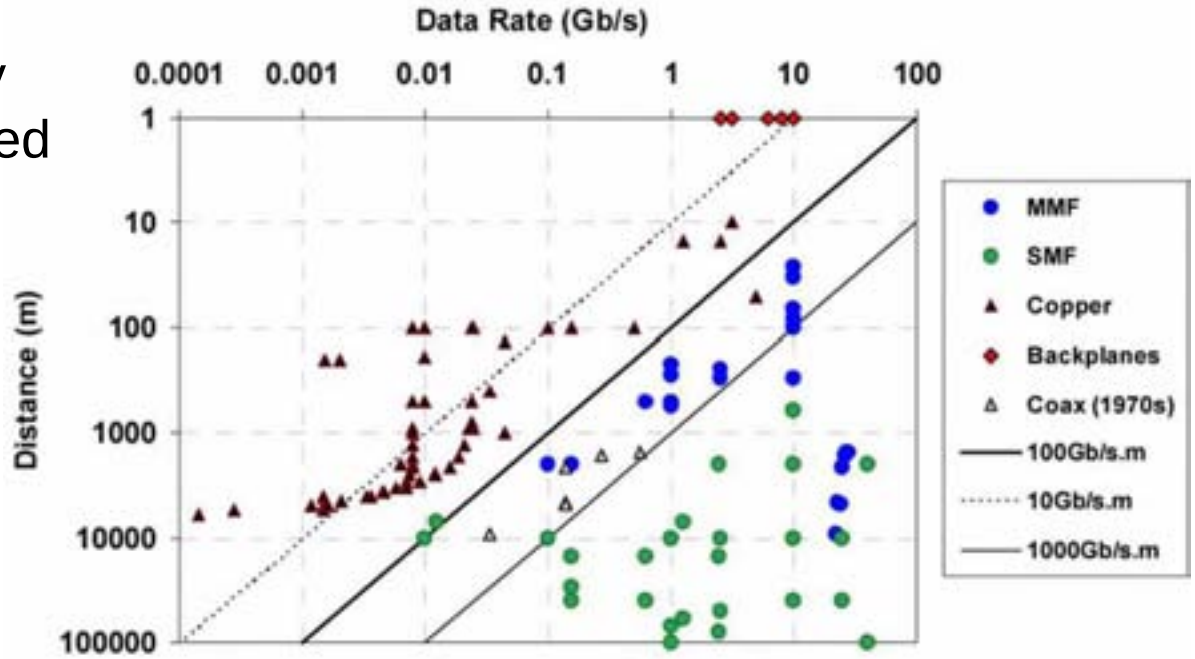

Build Massive Scale Compute

Optical inter-chip communication via silicon photonics

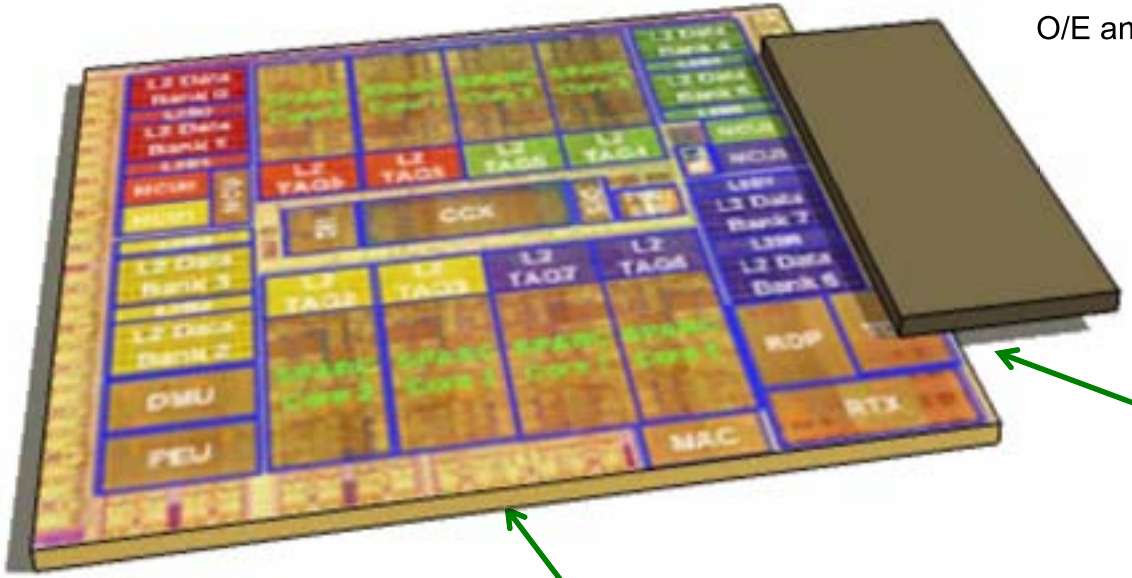

O/E and E/O “bridge” chip

Fine-pitch solder
(25 μm pads demonstrated)

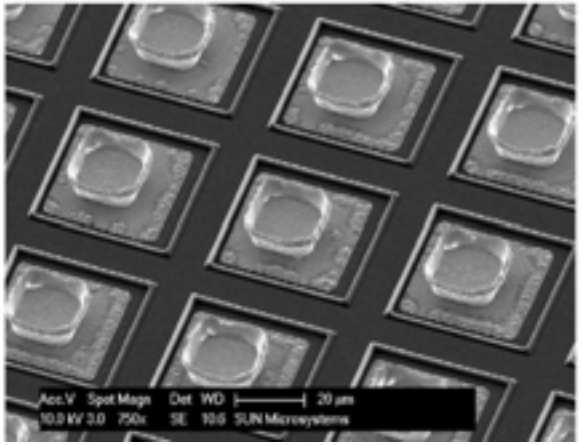

CMOS chip (no optics)

Higher reliability, low power, better integration than current optical technology

You will deploy massive-scale compute…

…Based on silicon-photonic-enabled massive-scale integration

- A server-in-a-package based on CMOS and optics
- More than 5 Tbps IO BW per chip

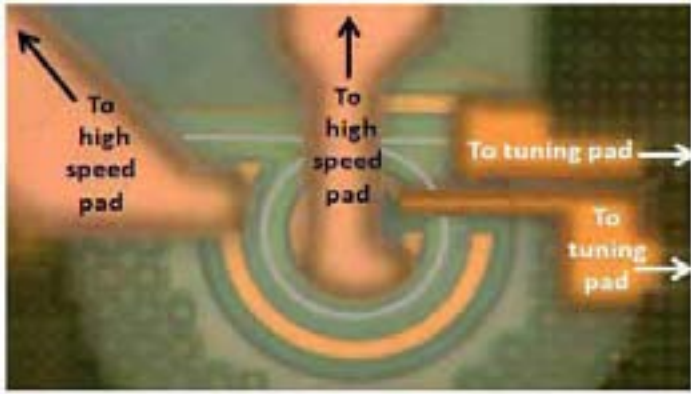

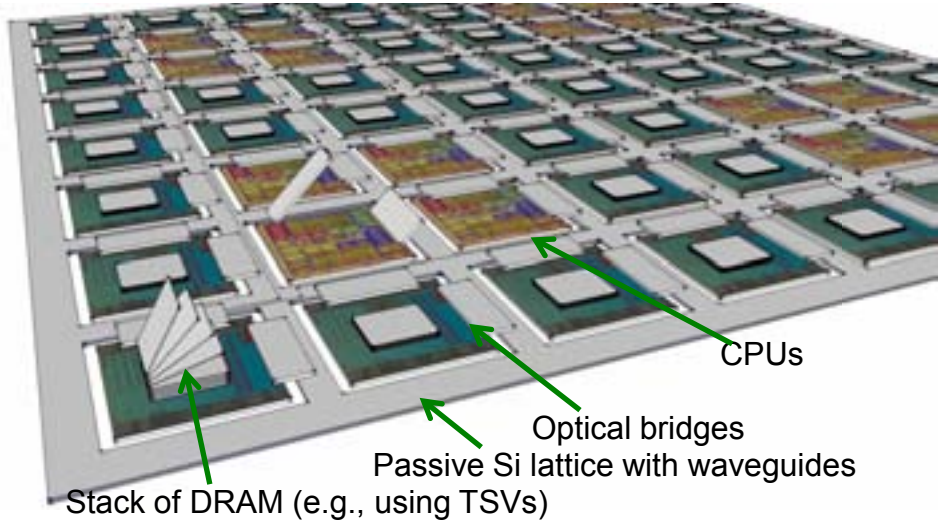

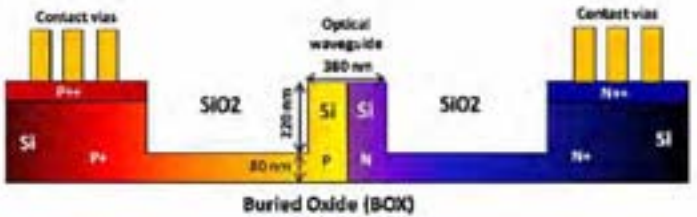

Massive scale computing

Next Generation Integration

- Chip integration
- High bandwidth I/O
- Optical inter-chip communication – Si Photonics

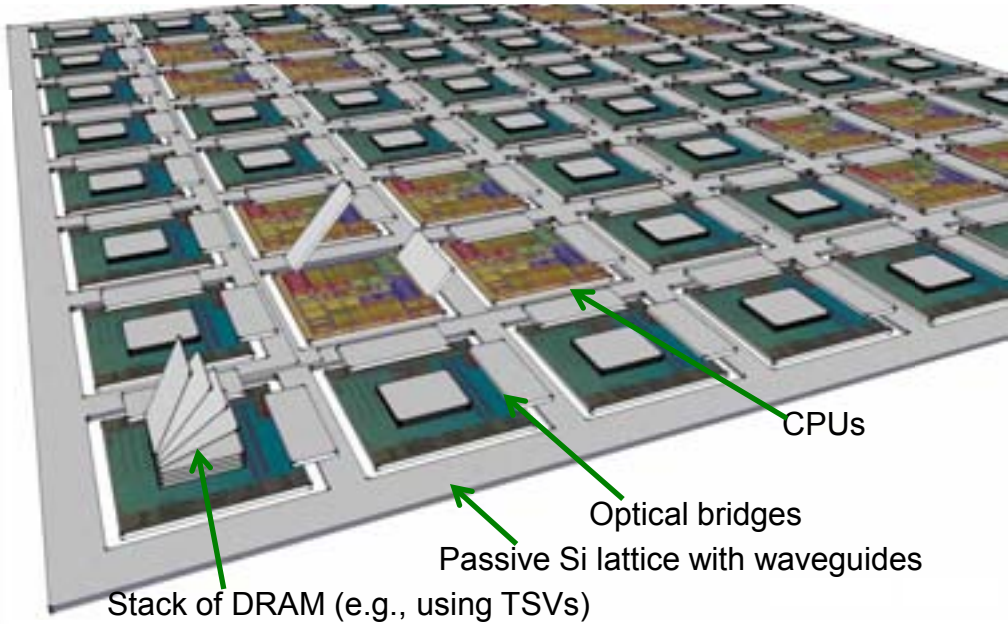

Harness Massive Scaleout Compute

Harnessing Massive Scaleout Compute

Programming massive numbers of cores – and specialized hardware

- Programmers have a hard time handling the concurrency implicit in massive parallelism
  - Exception - “Embarrasingly parallel” problems – web servers
- Tools are largely inadequate to the task

Harnessing Massive Scaleout Compute

Programming massive numbers of cores – and specialized hardware

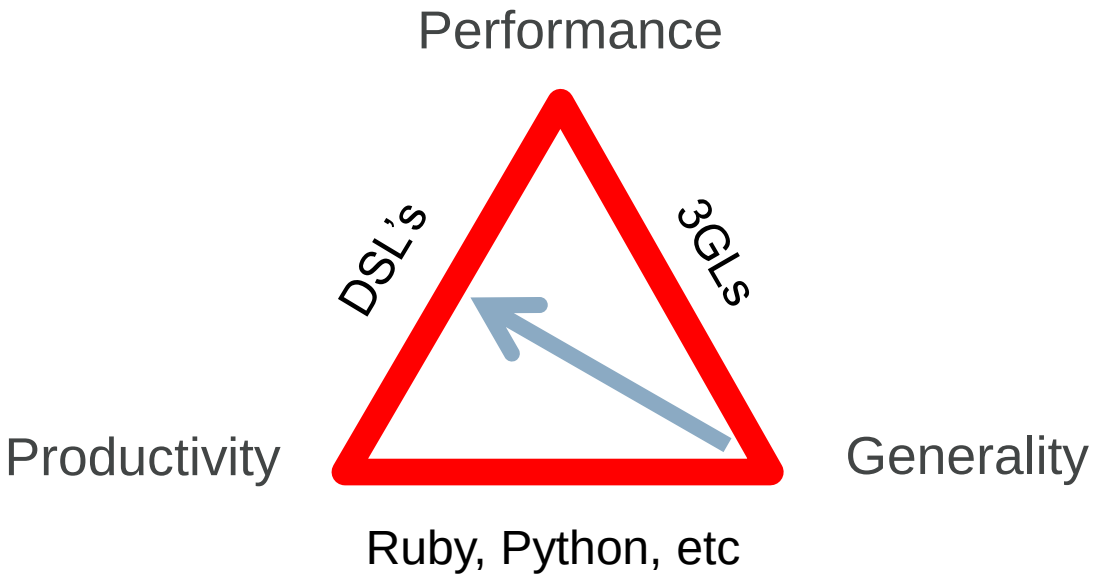

Harnessing Massive Scaleout Compute

Domain Specific Languages

- Why would you specialize a programming language?
- To capture abstractions at a high level

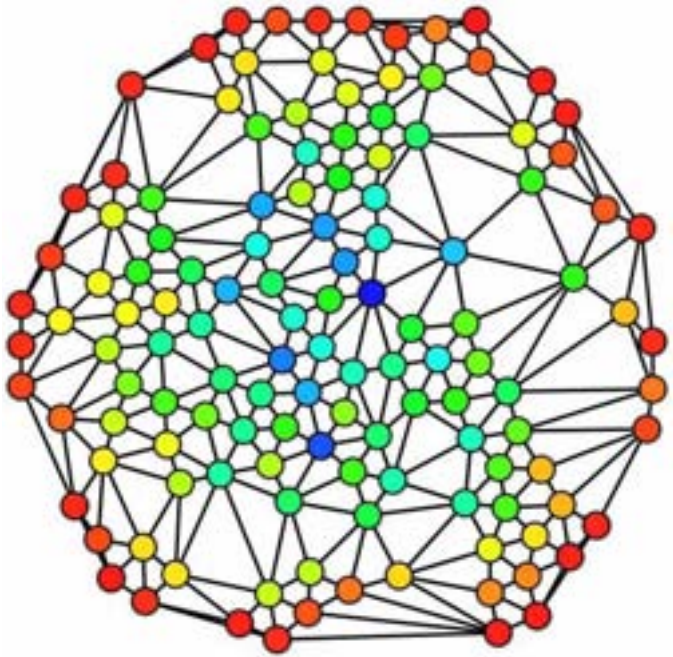

Claudio Rocchini

Harnessing Massive Scaleout Compute

Domain Specific Languages

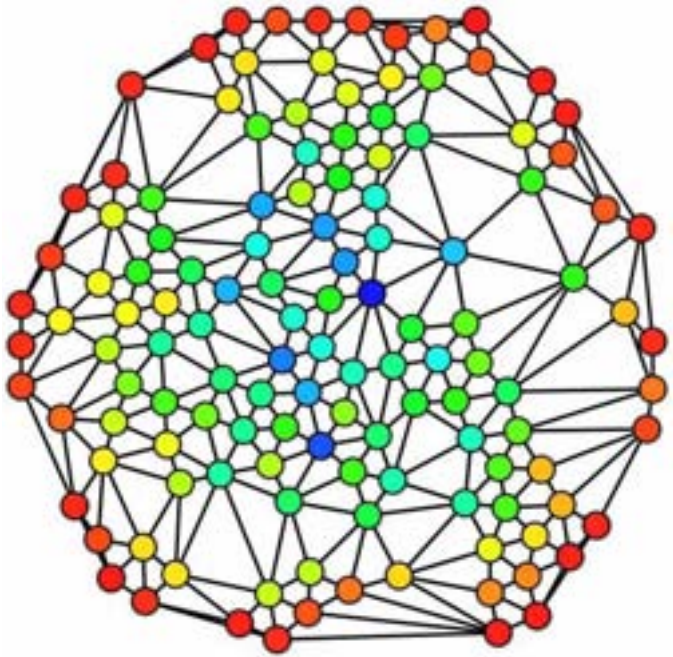

Brandes 2011, Hong et al 2012

Claudio Rocchini

Harnessing Massive-Scale Compute

Domain-Specific Languages

- Isn’t it wasteful to create new languages for every problem ?
- We’re creating a DSL framework for new domain-specific languages
  - And you can create your own
  - So you’re not hiring high-end SQL or map/reduce programmers
- What are some interesting domains other than graphs?
- Statistics… Finance…. You name it.

You Will Harness Massive-Scale Compute

..through Domain-Specific Languages

- In fact, you already do!
- The world's most popular domain-specific language is
  - Capable of expressing problems succinctly and clearly
  - Highly optimized for execution on heterogenous hardware
  - Scalable to massive numbers of nodes
- And it's called...

SQL

Exploiting Massive Scale Computers

Specialized Software

- Highly performant and productive languages
- Targeting specific domains

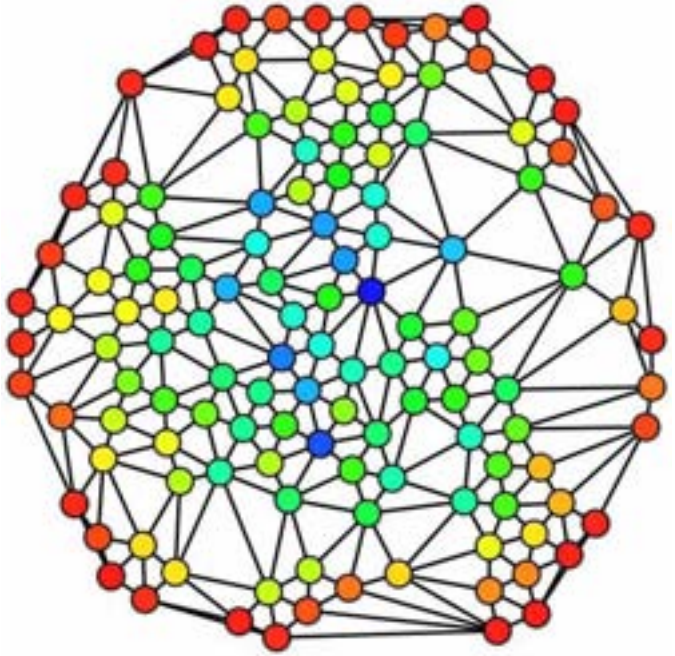

Agenda

- Research and Innovation at Oracle
- The End of Moore's Law?
- Challenges Beyond Moore's Law
- Opportunities Beyond Moore's Law
- **Your Future Computing Infrastructure**

Your Future Computing Infrastructure

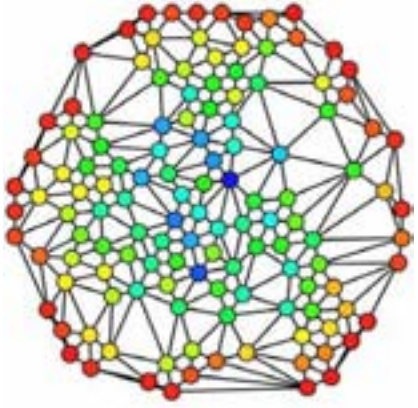

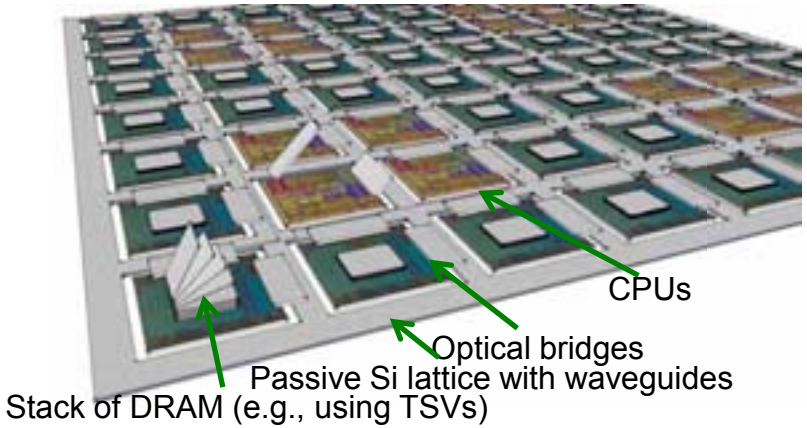

SQL

Hardware and Software
ORACLE
Engineered to Work Together

Hardware and Software

Engineered to Work Together


ORACLE®