

BEYOND EXASCALE: PLAYING THE CMOS ENDGAME

Steve Scott, Cray CTO

NITRD Future Computing Community of Interest Meeting

August 5, 2019




CRAY

Panel: Issues That Will Impact the Computing Field Over the Next 5-10 Years


It Was Good While it Lasted...

CRAY



42 Years of Microprocessor Trend Data

CRAY



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

<https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/>



Cambrian Explosion: Achieving Performance through Specialization

Processors are getting **Heterogeneous** and **Hot**

Serial perf vs. Parallel efficiency

HBM bandwidth & power efficiency
vs. DDR capacity

CPUs, GPUs, FPGAs, DSPs,
DL chips, Blockchain chips,
App-specific processors, ...


Maybe gives us an order of magnitude (no exponential advantage over time)



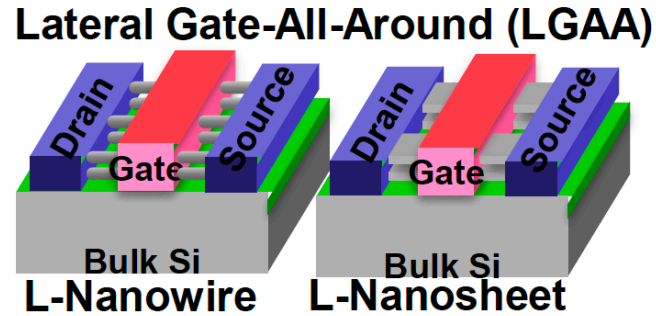
CMOS Device Structure

IRDS – 2017 More Moore Chapter


N7:2017-2019




N7:2019-2021



N5:2021-2024




>N3: >2024



- 3nm is on long-term processor roadmaps.... 1nm is on IRDS roadmap, but not clear we'll get there
 - Very modest improvement per generation
- After we're done with shrinks, will go to **3D stacking**
 - Allows "Moore's Law" to continue almost arbitrarily (transistors/package)
 - But *no* sustained perf/W improvement in logic! (other than reducing interconnect, which does matter)
- Prediction:
 - General purpose CMOS processors get to O(100 TF)
 - Special purpose CMOS processors get to O(1 PF) or clear


GPUs rule for high flop needs

Interconnects



State-of-the-art network

- Low diameter network with high-radix switches, enabled by high serdes rates and cost-effective optics
- > 90% efficiency at scale
- Highly-effective QoS and congestion control
- Uniform low latency (incl. tail latency)
- 80-90% of links use short copper cables



Dragonfly topology with 64-port switches

Over 250K endpoints with 3 switch-switch hops



Evolution over coming decade?

- Incremental improvements to synchronization, manageability, interoperability, etc.
- Next gen (100 Gbps serdes)
 - Electrical/optical crossover in cabinet
- Next-next gen (~200 Gbps serdes?)
 - On-package optics with WDM
- Retain low-diameter topology
 - Dragonfly, flattened butterfly

Software Evolution


- Moves more slowly than hardware...
- High end computing likely to remain MPI + OpenMP through the next 10 years
 - Software like Kokkos and RAJA could also help deal with HW heterogeneity

Workloads also becoming
more heterogeneous




Post-Exascale Asymptote of the CMOS Era

CRAY



O(100 TF) per node
GPUs (wide vectors?)
More for specialized procs




Memory on package
(on proc?)



Disk for
archive



Flash for
main storage




NVM for
storage cache

*Movement towards high-perf object store
or KVS on storage-class memory*



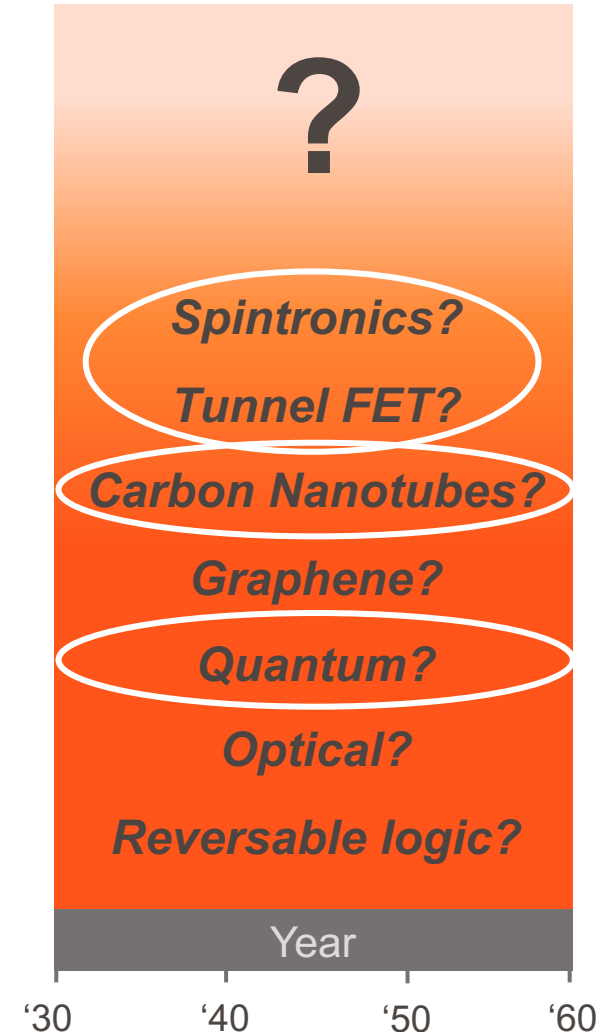
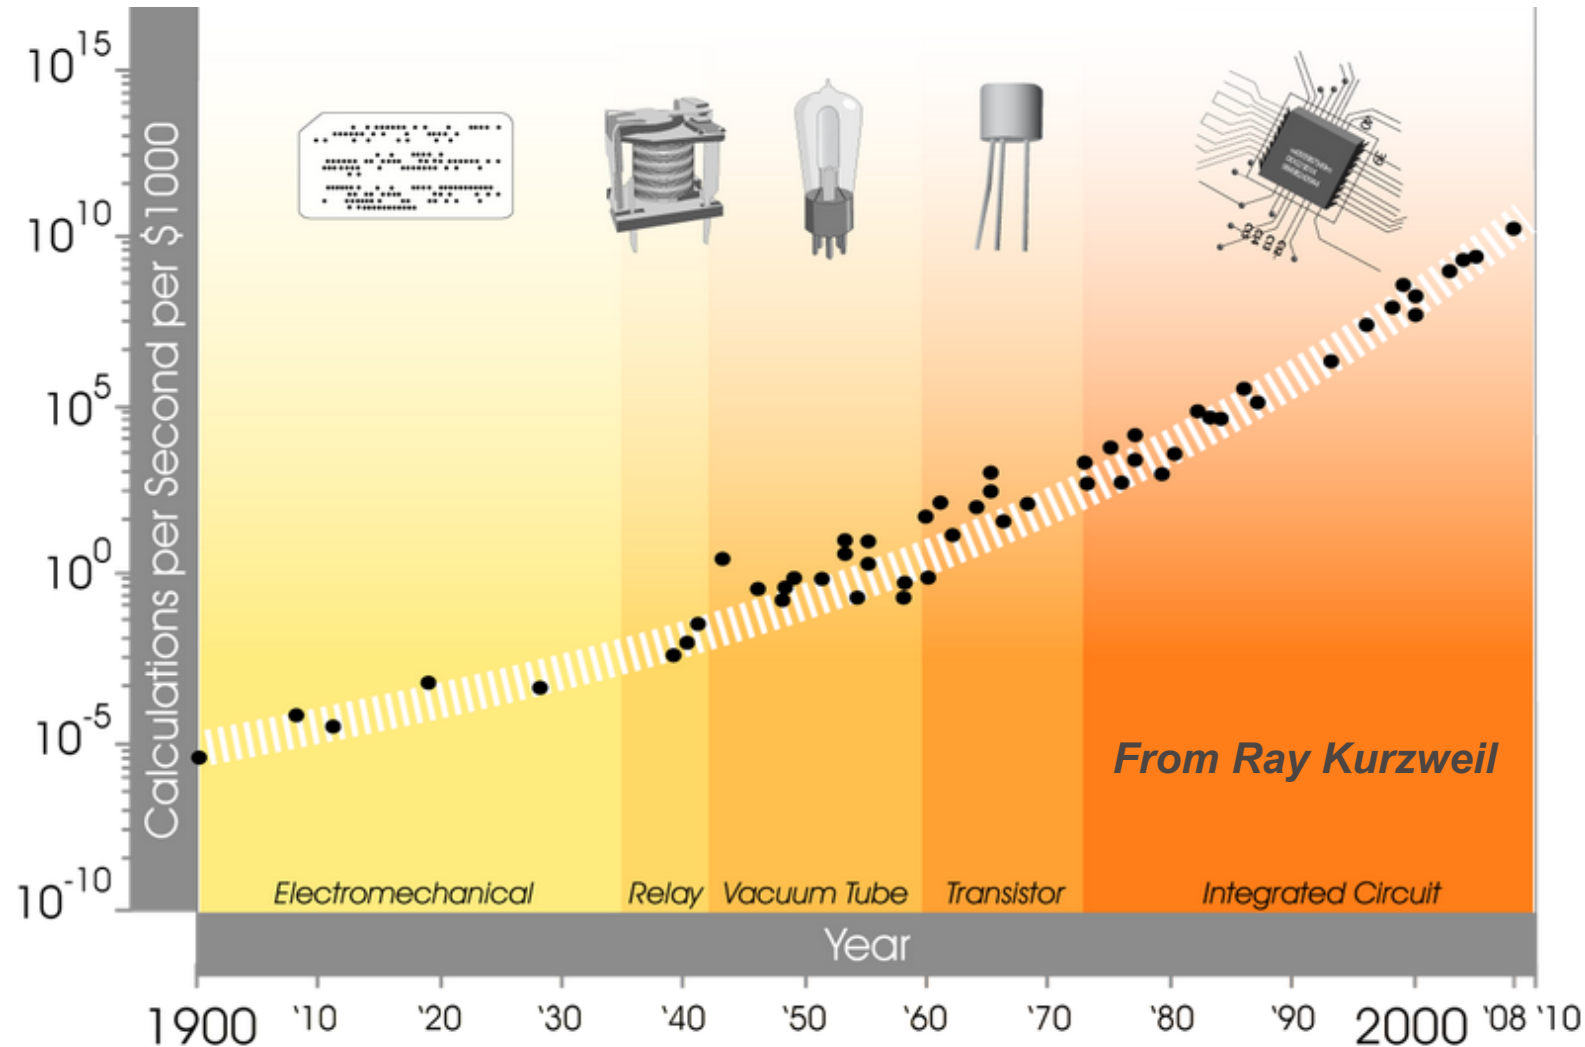
O(100k) nodes, ~30-50 MW
O(10 Exaflops)



Low diameter networks with optics
Dragonfly, Flattened Butterfly

Over the next decade, computers won't change that much from the current model.

What Comes After CMOS?



Post-Exascale Processor Advancement



Specialized CMOS architectures

- DNN accelerators
- Neuromorphic
- Configurable/spatial
- PIM
- Application-specific
- ...

Low hanging fruit

Definitely pursue

Post-CMOS general switching technology

- Tunnel FET
- Spintronics
- Graphene
- Superconducting/cryo
- Photonic computing
- ...

No clear candidate

Needs research! (please fund!)

Meanwhile... 3D stacking

Specialized Non-CMOS, Non-Von Neuman

- Quantum
- Analog
- RNA/DNA
- ...

Worth pursuing (long term)

***Not** a real replacement for general purpose computing*

Expect a couple generations beyond Exascale without too much change. But we won't get to zettascale without some breakthroughs.

Quantum Computing



Quantum is a revolutionary technology with the potential to radically change computing as we know it!

?

Quantum is an esoteric technology that is far from productization and will not significantly change the HPC landscape.


Both. Simultaneously.

Fors CRYO

Tempering Expectations for Quantum Computing

CRAY

(universal gate model)



- Practical issues: cooling, isolation, coherence, control, I/O, etc.
- High error rates: large overhead to achieve reliable logical qubits; limited circuit depth
- Small number of core algorithms
 - Near term algorithms are mostly heuristic
 - Algorithms with strong quantum advantage require very large machines
- Very limited state for problem input & output
- Predictions:
 - 5+ years to contrived example of quantum beating classical
 - 10+ years for any *practical* quantum advantage
 - 15-20 years for large machines with strong quantum advantage
- And quantum will still address a small subset of what we use classical HPC for

Courtesy Fred Chong, University of Chicago
Lead PI for EPIQC (*Enabling Practical-scale Quantum Computation*)

THANK YOU

QUESTIONS?



"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development
(NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674,
Fax: 202-459-9673, Email: nco@nitrd.gov, Website: <https://www.nitrd.gov>

