SC14

New Orleans, LA | hpc matters.

Highlights from SuperComputing 2014

SFT Group Meeting
Jakob Blomer

Overview

SC14
New Orleans, LA | hpc matters.

Programme:

- Exhibition
- Conference and plenary talks
- Panel discussions
- Workshops
- Posters
- Discussion rounds
- Student showcases

Some Numbers:

- ~10,000 attendees (4+ from CERN)
- ~350 exhibitors
- ~35 workshops
- ~15 panels
- Conference network: 1.5 Tbit/s

http://sc14.supercomputing.org

Plenary Talks

Keynote by string theorist Brian Greene

Charles Leiserson:

- Moore's law ends –
  in its technical formulation
  - Within 5-6 year limit of
    ~5nm structures reached
- Moore's law, in its popular formulation,
  has to continue
  - Scale-out
  - Growth from algorithms, systems
    engineering, compilers, … software
  - Reducing bloat, removing inefficiencies,
    scientifically based performance engineering

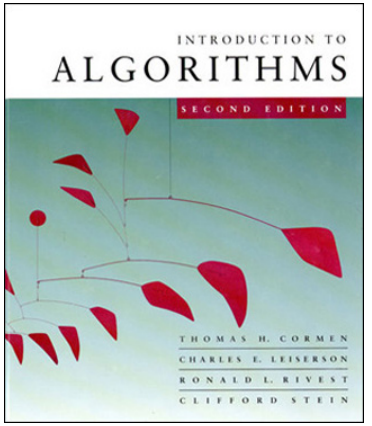

Cilk™ Plus

Plenary Talks (cont.)

Gordon Bell

- "It's a matter of programming"
- 2 eras of supercomputing
  - Mono computers,
    40% faster per year, 100x (sic!) per decade
  - Distributed clusters,
    2x faster per year, 1000x per decade
    Every generation 10x more expensive

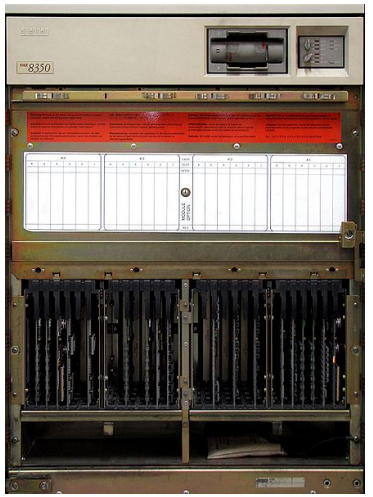

Computer History Museum

Today's Architecture

- Cluster's of CPU + Accelerator
- Non volatile memory per box
  deep memory architecture
- Fast interconnects (MPI)
- Sum TOP500: > 100PFLOPS

Energy efficiency becomes important

- E.g. "Green 500"
- Better hardware packaging
  (like mobile devices)

How and when can HPC reach "exascale"?

Memory for the Exascale (Panel)

Participation from Intel, ARM, Nvidia, Micron, IBM, AMD

- No disruptive technology on the horizon to replace SDRAM on a larger scale
- DDR4 standard for the next 5+ years
  work on DDR5 has not yet begun
- Deep memory hierarchy expected with
  - new RAM technologies (e.g. 3D-DRAM) close to the chip
  - SDRAM and NVRAM on higher levels
- Virtual memory programming model considered problematic
- Error correction capabilities have to increase

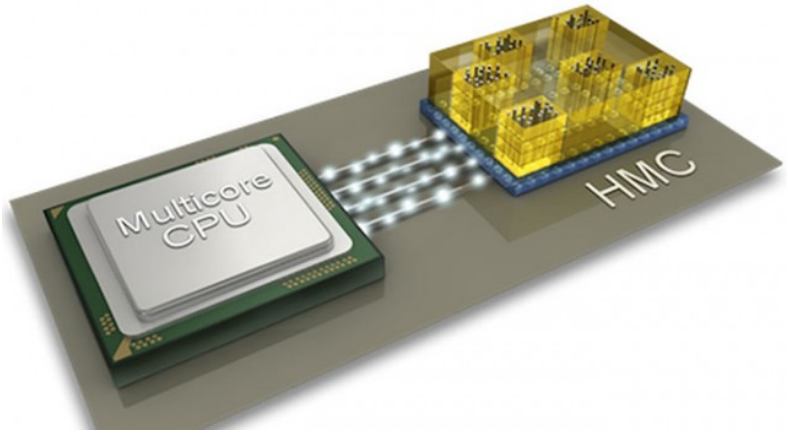

Products: Processors

IBM Power8

- 6-12 cores, 8 threads per core, up to 4.5GHz
- CAPI: port to accelerators or flash, unified memory with CPU
- Adopted ARM's licensing model

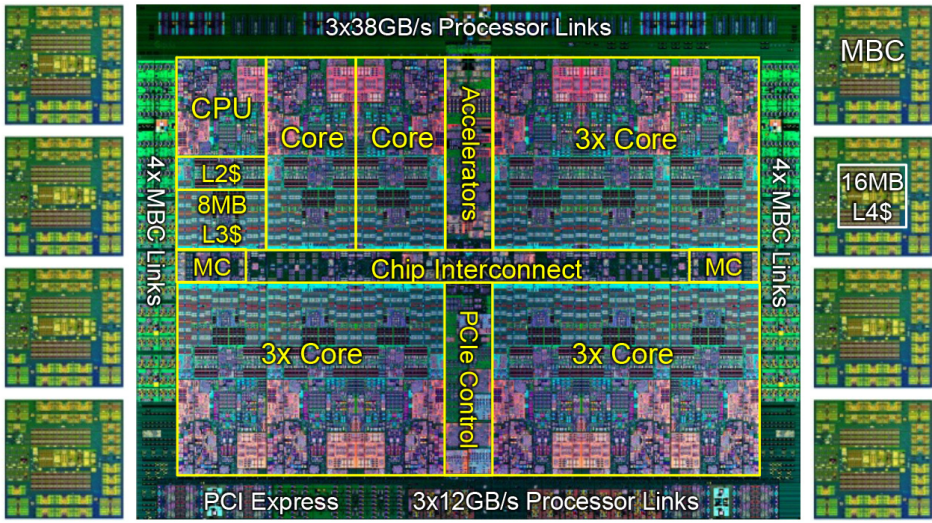

ARM-64

- Not yet in mass production

Xeon Phi ("Knights Landing", no new public announcements)

- 72 Atom cores, 4 threads / core
  (Only slight increase in core count,
  full x86 binary compatibility except TSX)
- Self-hosted mode (act as a node, not as a co-processor)
- 8-16 GB in "on-package" 3D DRAM (~5x bandwidth of DDR4) either as a cache or explicitly programmed
- Deep memory hierarchy, more programming models

Products: Network and Storage

Ethernet

- 2 x 40GbE card available from Mellanox, Solarflar, Cisco
  Can be split into many 1GbE PCI devices, e.g. for virtualization
- Mellanox: 100Gb InfiniBand/Ethernet adapter
- 1GbE per core is feasible

Flash (SSDs)

- Flash memory: $1 – $2 per gigabyte
  100,000 – 200,000 IOPS/s, 500MB/s
- IBM FlashSystem: 1 PB,
  21 x 1.2MOps / second

Integrated Storage

- BLOB store (S3-like interface), called
  - "Software-defined storage" by Scality
  - "Object store" by Caringo
- Erasure coding
- Tens to hundreds of Petabytes deployed at customers

Products: FPGAs

Hardware

- Many board vendors
- PCIe connected
- Typically order of 10 gigabytes of RAM
- Low energy consumption (order of 10W)
- Price range from a few hundred dollars to a few thousand dollars
  Number of logic gates are important factor for the price
  → Competing with GPUs

Software

- Traditionally VHDL, Verilog
- Altera and Xilinx offer OpenCL compilers for FPGAs
- Example:
  image processing algorithm, 6 months in VHDL vs. 4 weeks in OpenCL
- Lots of open questions, potential student project?

Products: ArrayFire

- Formerly commercial accelerator library
- C, C++, Fortran interfaces with
  CPU, OpenCL, CUDA backends
- BSD licensed, https://github.com/arrayfire/arrayfire
- Arithmetic, matrix operations, sorting, histograms, fft,
  image processing

2nd Workshop on Sustainable Software in Science

- ~ 50 participants from different sciences
- Key note, discussion rounds, and lightning talks
- Broad range of topics:
  - Software engineering, tools and packages, long-term preservation of software, community building, policy and funding discussions
- Community still in very early phase
  - E.g. “sustainable” not yet well-defined (~ reproducible)
  - Unfortunately little participation from funding agencies
- Discussions tightly coupled to the scientific publishing process

- Interesting pointers for us:
  - Journal of open research software
  - The Software Sustainability Institute (U.K.)
    looks very similar to our HEP Software Foundation efforts

Overall impression

- Unique preview of latest products and upcoming hardware landscape, all major vendors present
- Technical program partly off-topic for our field, but
  - Excellent plenary talks and
  - Excellent panels
- Convergence of HPC and BigData claimed, but
  - No presence of Google, Facebook, Amazon, …
  - Little coverage in the technical programme
- The “exascale supercomputer” until 2018 is very ambitious
- Given the growing share we get on HPC,
  this conference can be an important point of exchange

Increasing chip density will stop very soon!
Improvements must come from software.