Multicore in Industry

Some random observations

UPMARC Seminar, 2008-10-02

Dr. Jakob Engblom, Virtutech

virtutech

Future Embedded Systems

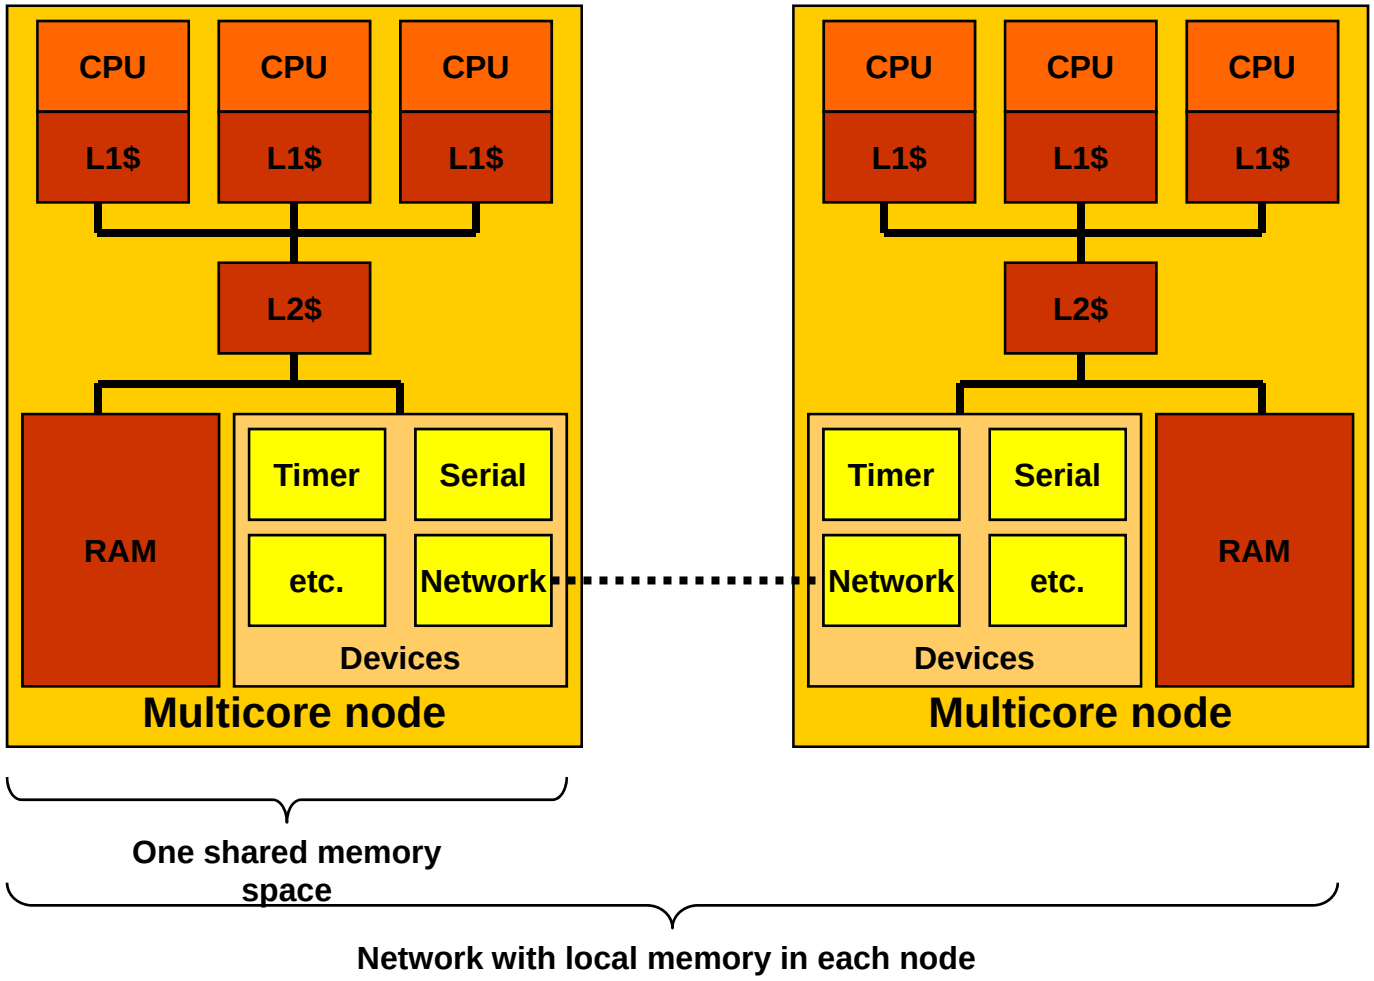

Accelerators

On a modern SoC, the processor cores are just one part of the system

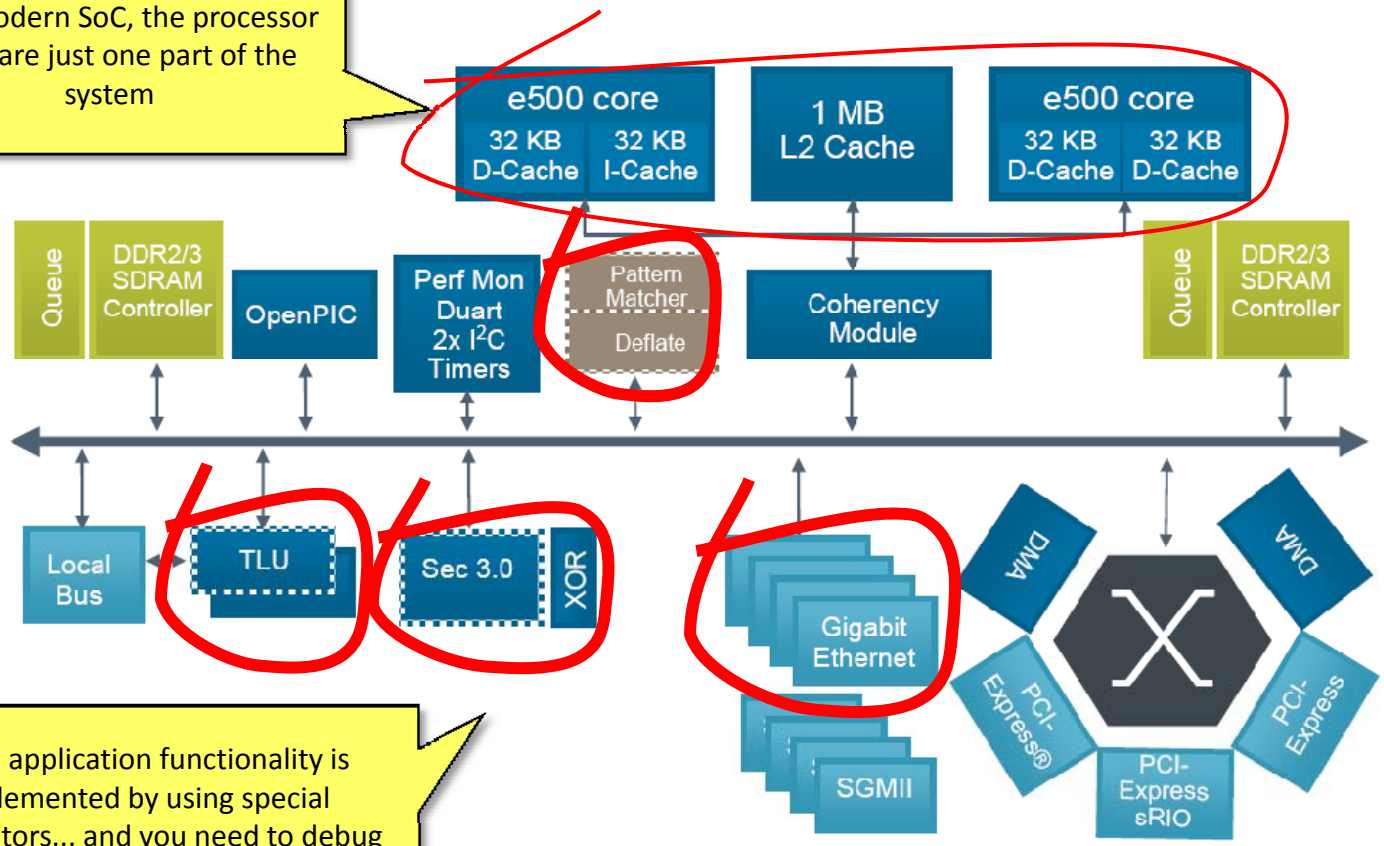

Much application functionality is implemented by using special accelerators... and you need to debug their interaction with the processors & software

virtutech

Accelerators or not?

- Heterogeneous school vs Homogeneous school
- I have some blog posts:
  - http://jakob.engbloms.se/archives/269
  - http://jakob.engbloms.se/archives/283
  - http://jakob.engbloms.se/archives/44
  - http://jakob.engbloms.se/archives/90
- And some opposite views:
  - http://a-vajda.eu/blog/?p=39

Multipro Hardware Debug Support

- Cross-triggering
  - Hardware units listen to events on all cores
    - Breakpoints, raw memory trace, watchpoints, interrupts...
  - Cause action in one core based on events occurring in other cores or elsewhere in the system
    - Stop execution, start tracing, stop tracing, interrupt, ...
    - Requires logic on the multicore chip
    - Basically, it is programmable
- See also:
  - http://www.arm.com/products/solutions/CoreSight.html

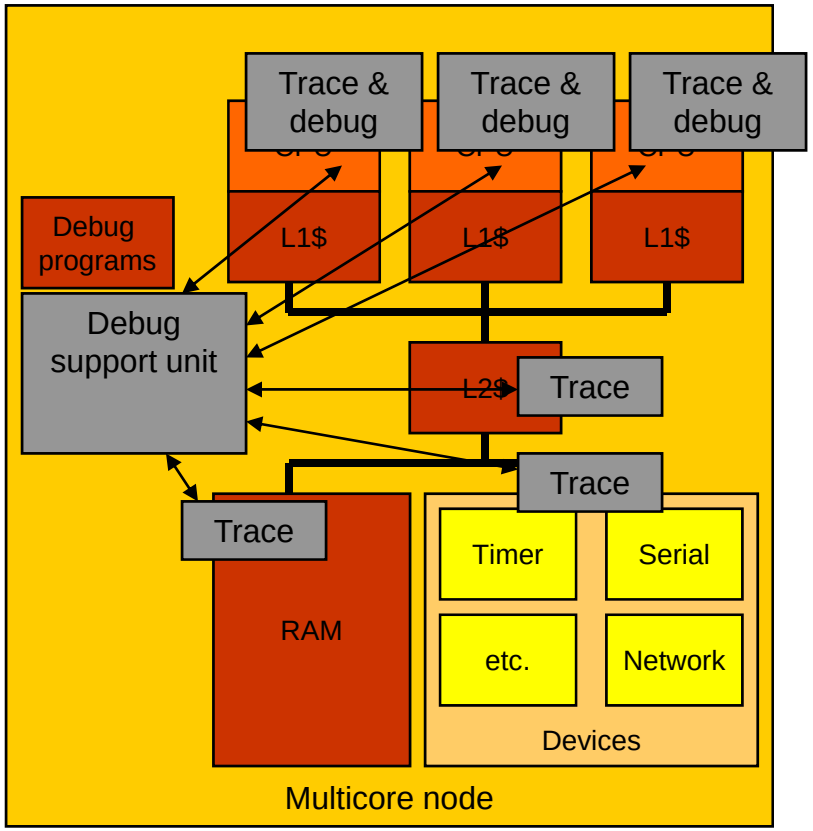

virtutech

AMP Use on a Single Chip

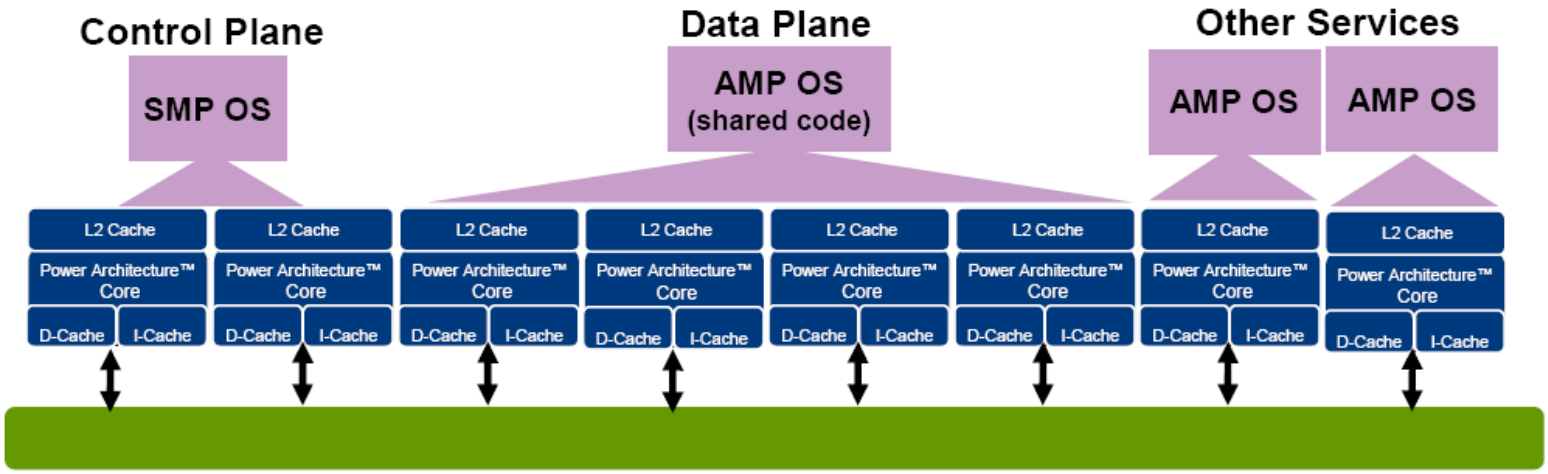

virtutech

Hypervisors

- Hypervisors will multiplex operating systems on cores
- IOMMU to gate access to hardware, just like regular MMUs gate access to memory
- IOMMUs are coming in hardware, real soon!

- See
  - http://jakob.engbloms.se/archives/37
  - http://jakob.engbloms.se/archives/22

Domain-Specific Languages

- Create easy models for each programmer
  - Should be mostly isolated local sequential
  - Do not expect all programmers to be parallel masters
  - Execute many of these in parallel
- Tailored to domain
  - My example:
  - DML for device modeling in Simics
    - Automatically thread safe
    - Local-memory
    - Five times less code than C
    - Appropriate primitives for domain
- See also:
  - http://jakob.engbloms.se/archives/209
  - http://www.scdsource.com/article.php?id=166

virtutech

Cores Become Free

- We used to have that CPUs were expensive and had to be shared efficiently
  - That is not true anymore
  - Cores are “free”
  - IO to and from a chips is “expensive”
- How can you leverage many cores to simplify software and make it more robust and simple to create?
  - I have a series of blog posts on this:
    - http://jakob.engbloms.se/archives/58
    - http://jakob.engbloms.se/archives/123