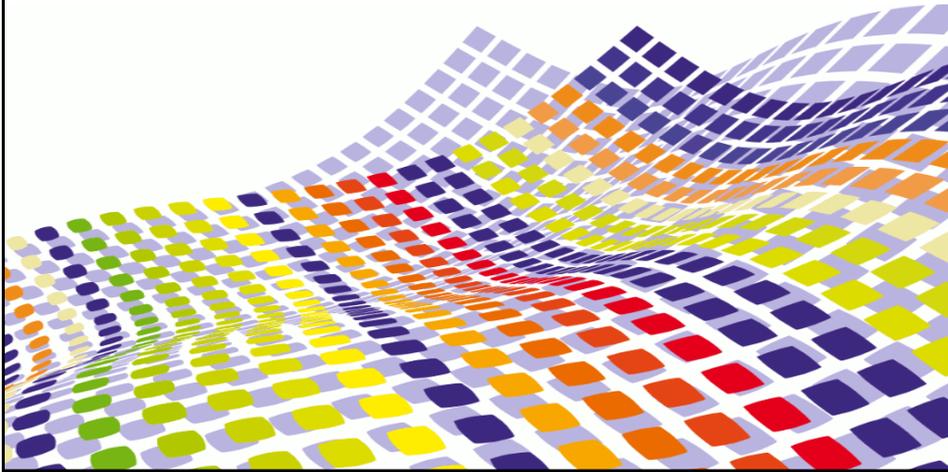


HMPP Origins



Agenda

- GPGPU & Hybrid Parallel computing
- HMPP Concepts and Overview
- Starting with HMPP
- Addressing hardware accelerators with HMPP
- HMPP Toolchain
- HMPP Runtime
- Managing Data with HMPP
- Grouping Codelets
- Sharing Data with HMPP
- HMPP Features & Roadmap



GPGPU & Parallel Hybrid Computing



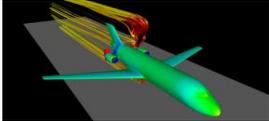
www.caps-entreprise.com

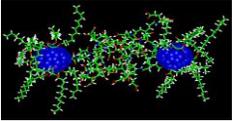
3



Industry and Business Facts

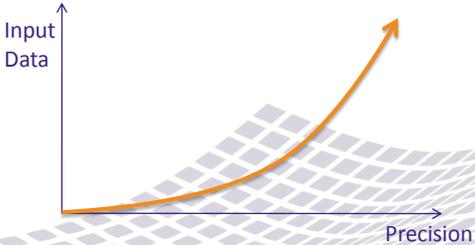
- Modeling & Simulation are pervasive







- Precision is the key to success



www.caps-entreprise.com

4

Why Hybrid Computing?

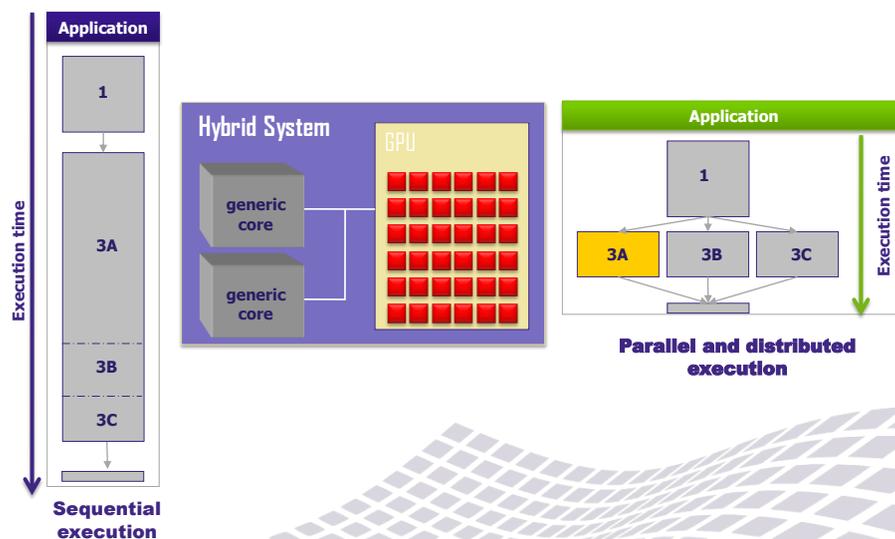


- More precision implies more data to process
 - But still in the same time, so more speed in computations!
- Current technologies reached a limit (in terms of frequency)
- One solution is to use parallelism (increase # of cores)
 - We do more things in parallel instead of doing it quicker
- The efficiency scale evolves according economic issues
 - Performance / Power consumption
 - Performance / Development time
 - ...
- Mainstream applications will rely on these multicore / manycore architectures
- Various heterogeneous hardware
 - General purpose cores
 - Application specific cores (e.g. GPUs)

www.caps-entreprise.com

5

You said Parallel Hybrid Application?



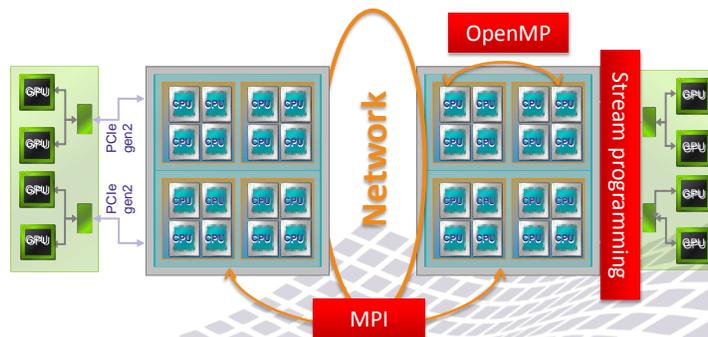
www.caps-entreprise.com

6

Multiple Parallelism Levels



- Adding a new layer of specific hardware is adding a new workload to the developer
- Programming various hardware components of a node cannot be done separately



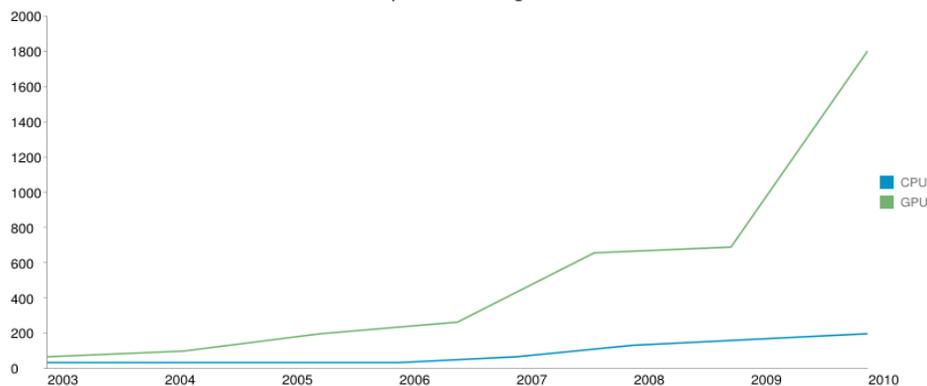
www.caps-entreprise.com

7

Parallel Processor Architectures

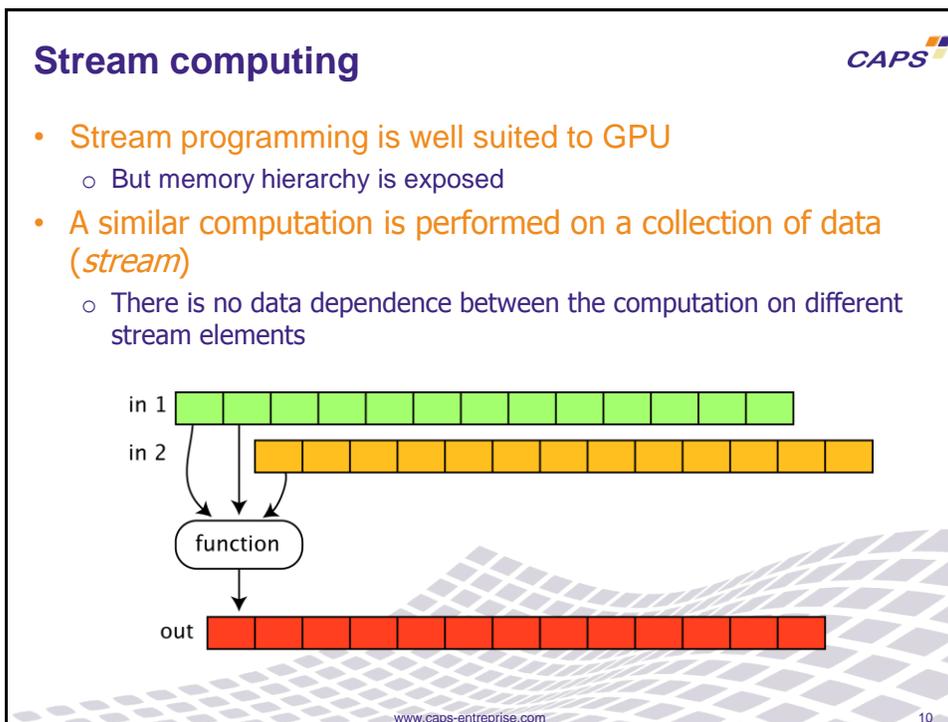
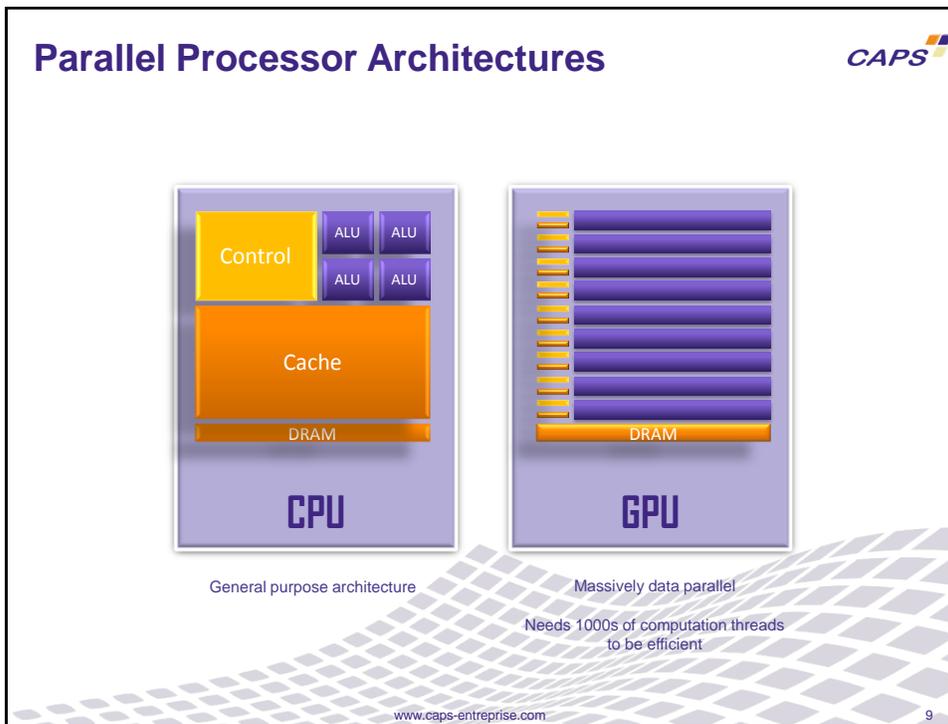


Peak performance growth



www.caps-entreprise.com

8



Offloading computations

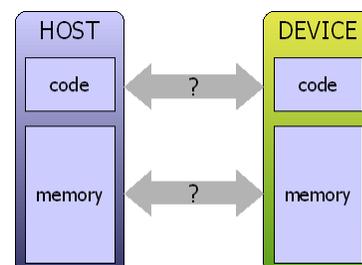


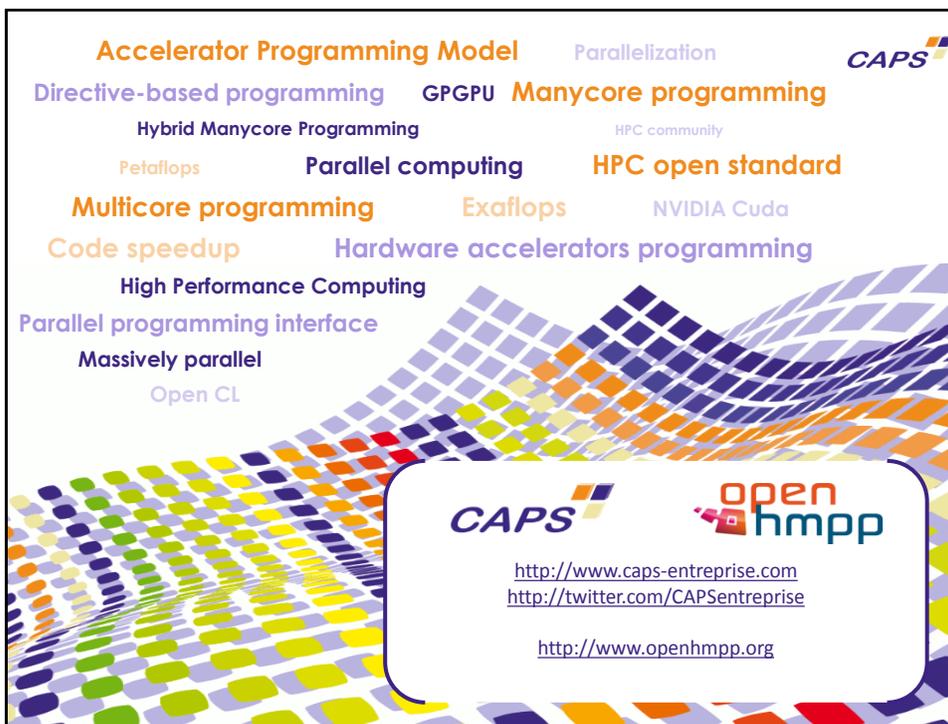
- **Host: General purpose cores**
 - Share a main memory
 - Core ISA provides fast SIMD instructions
- **Device: Streaming cores**
 - GPU, DSP, FPGA... (vector, SIMD)
 - Application specific architectures
 - Can be extremely fast
- **Hundreds of GigaOps**
 - But not easy to leverage
 - Restriction to one platform is not acceptable

Key Issues



- Using accelerators makes you stick to a proprietary language/environment/technology
- Huge potential performance but accelerators are far from host memory
 - Data must be copied on the remote device
 - Due to narrowband links between CPU/HWA, data transfers are critical
- Therefore rethink the computation organization/algorithm





Accelerator Programming Model Parallelization **CAPS**

Directive-based programming GPGPU **Manycore programming**

Hybrid Manycore Programming HPC community

Petaflops **Parallel computing** **HPC open standard**

Multicore programming Exaflops NVIDIA Cuda

Code speedup **Hardware accelerators programming**

High Performance Computing

Parallel programming interface

Massively parallel

Open CL

CAPS **open hmpp**

<http://www.caps-entreprise.com>
<http://twitter.com/CAPSentreprise>

<http://www.openhmpp.org>