

# Birds of a feather

Trond Runar Hagen, Johan Seland

Geilo Winter School 2008

## Definition

*A BoF session, an informal meet-up at conferences, where the attendees group together based on a shared interest and carry out discussions without any pre-planned agenda.*

Here:

- We wish to invite you all to discuss parallelism and it's impact
- We will have a panel (of experts?) to lead the discussion

Why this time?

- Parallel hardware is now everywhere
- Not only in governmental/department cluster, but on laptops
- Today's computers now have dedicated data-parallel and task-parallel processors, everyone is connected

# Where is the parallelism?

## Message Passing:

- Fast network everywhere
- We could build a 100+ node cluster in this room, **now**
- Industry standard for programming (MPI)
- Driven by user demand

## Shared Memory:

- “All” CPUs now have multiple cores
- Industry standard for programming (OpenMP)
- Forced upon us by the industry?

## Data-parallel:

- All computers now have a GPU
- Often the most costly component in a PC
- No clear industry standard for computations
- Driven by user demand

The Winter School certainly tells us that the age of parallelism has started

- Are university curriculums being adapted to this reality?
- But automatic parallelism still seems very far off

So the big question seems to be:

The Winter School certainly tells us that the age of parallelism has started

- Are university curriculums being adapted to this reality?
- But automatic parallelism still seems very far off

So the big question seems to be:  
Are you willing to retool, rewrite and debug?

- What happens if you don't?

Late 1990s: GPUs became standard on consoles and in gaming PC

## Statement

*The GPU is the only parallel processor that has seen widespread success*

- Gave development completely changed over a few years
- Developers **had** to learn new tools and adapt a new mindset very rapidly
  - Web/mobile games continues the “simpler” tradition
- Curriculums have changed as well
- $\Rightarrow$  is it optimistic to expect this from you as well?

*We currently assume that the kernels of the scalable parallel algorithms/applications/libraries will be built by experts with a broader group of programmers (mere mortals) composing library members into complete applications.*

– Geoffrey Fox



*We currently assume that the kernels of the scalable parallel algorithms/applications/libraries will be built by experts with a broader group of programmers (mere mortals) composing library members into complete applications.*

– Geoffrey Fox

Are you “mere mortals” or experts?

Can we assume that better tools will save us?

- This Winter School does not seem to indicate so
  - And has actually focused on new languages and compilers
- Is it finally time for functional programming?
- Should we (the mere mortals) leave our inner-loops to experts?

- Bandwidth could end up being the real problem
  - Severe enough to use single precision?
- Today's multi-core chips are already semi-NUMA
- Intel expects all processors to be NUMA in some years
- Memory hierarchy is not exposed in today's languages
  - But Operating Systems have (some) support
- Can we expect “mere mortals” to program the memory hierarchy? (CUDA, Cell?)
- Do we need a runtime-system on top of the OS?

Kjør debatt!

