



**KTH Information and
Communication Technology**

Scalability and Programmability in the Manycore era

Mats Brorsson

Royal Institute of Technology (KTH)

New research group together with

SICS (Swedish Institute of Computer Science), and

Uppsala university (Erik Hagersten)



By 2018 the hardware platform...



KTH Information and
Communication Technology

- ***Cores will be many***
ITRS roadmap predicts 100-1000+ cores (x1.4/year)
- ***There will be non-uniformity***
NUMA, NUPerf (Non-uniform performance)
- ***On-chip storage till be abundant***
Gbyte of cache, shared address space
- ***Off-chip bandwidth is limited***
<200 Gbyte/s data rate, 3D-stacking may give more
- ***Point-to-point networks will be used***
Mesh networks likely used because of wire delays
- ***Locality is a major issue***
Computations need to be placed where data is (or vice versa)
- ***Cores are unreliable***
Reliability is realized through software

SKOLAN FÖR INFORMATIONSS-
OCH KOMMUNIKATIONSTEKNIK

Software development

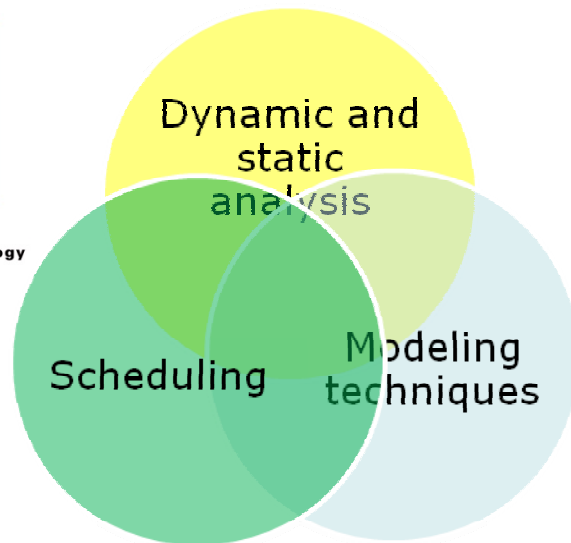


KTH Information and
Communication Technology

- ***Safe parallelism should be adopted***
 - Parallel programming should have a sequential reading
 - Debugging should be done in the sequential domain
- ***Established programming languages will be used***
 - The world will not convert into new "strange" languages
- ***The important parallel construct is the task***
 - The thread model becomes too difficult and coarse grain to reason about
 - Potential models: OpenMP tasks, cilk, cilk++
- ***Hardware structure should not be exposed to programmers***
 - In order to be future-proof, software should never assume a particular hardware structure

SKOLAN FÖR INFORMATIONSS-
OCH KOMMUNIKATIONSTEKNIK

Research areas



- **Dynamic and Static Analysis** to find and expose safe parallelism
 - Finding data dependences
 - Using type and effect systems
- **Modeling techniques**
 - Understanding and predicting performance (and power consumption) on future systems
- **Scheduling**
 - Scheduling tasks on worker threads taking heterogeneity and locality into account
 - Scheduling hardware resources (cores) to adapt to performance needs.
- Architectural support for the above:
 - Data dependence analysis
 - Modeling
 - Scheduling of tasks

SKOLAN FÖR INFORMATIONSS-
OCH KOMMUNIKATIONSTEKNIK

Interest in research collaborations



KTH Information and
Communication Technology

- Host for short visits on specific topics (paper writing)
- Partner(s) in EU FP7 STREP proposals
- Complementing competences
 - Network systems and architectures (NoC)
 - Memory architectures (e.g. Coherence & TM)
 - Application providers
 - Performance tools (task-oriented)

Contact:

- Mats Brorsson
matsbror@kth.se



Current and past achievements



KTH Information and
Communication Technology

- Operating system algorithms for scheduling hardware resources (processors) for power saving applications (KTH)
- World-leading analysis and modeling tool for performance estimation (Uppsala university)
- Dynamic data-dependence analysis for simplified and correct parallel programming (SICS)
- Development of the first open source OpenMP implementation and run-time system for SMP:s and clusters (KTH)
- World-leading implementation of parallel Prolog (SICS)
- High-performance software distributed shared memory systems (Uppsala university, KTH)

SKOLAN FÖR INFORMATIONSS-
OCH KOMMUNIKATIONSTEKNIK