Parallel Programming in Computational Engineering & Science

Introduction

Matthias Müller, Dieter an Mey
Center for Computing and Communication

Kindly sponsered by

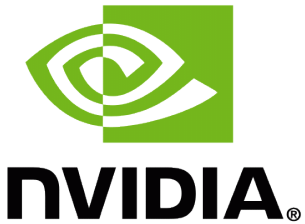

RWTH Systems in Top 500 List

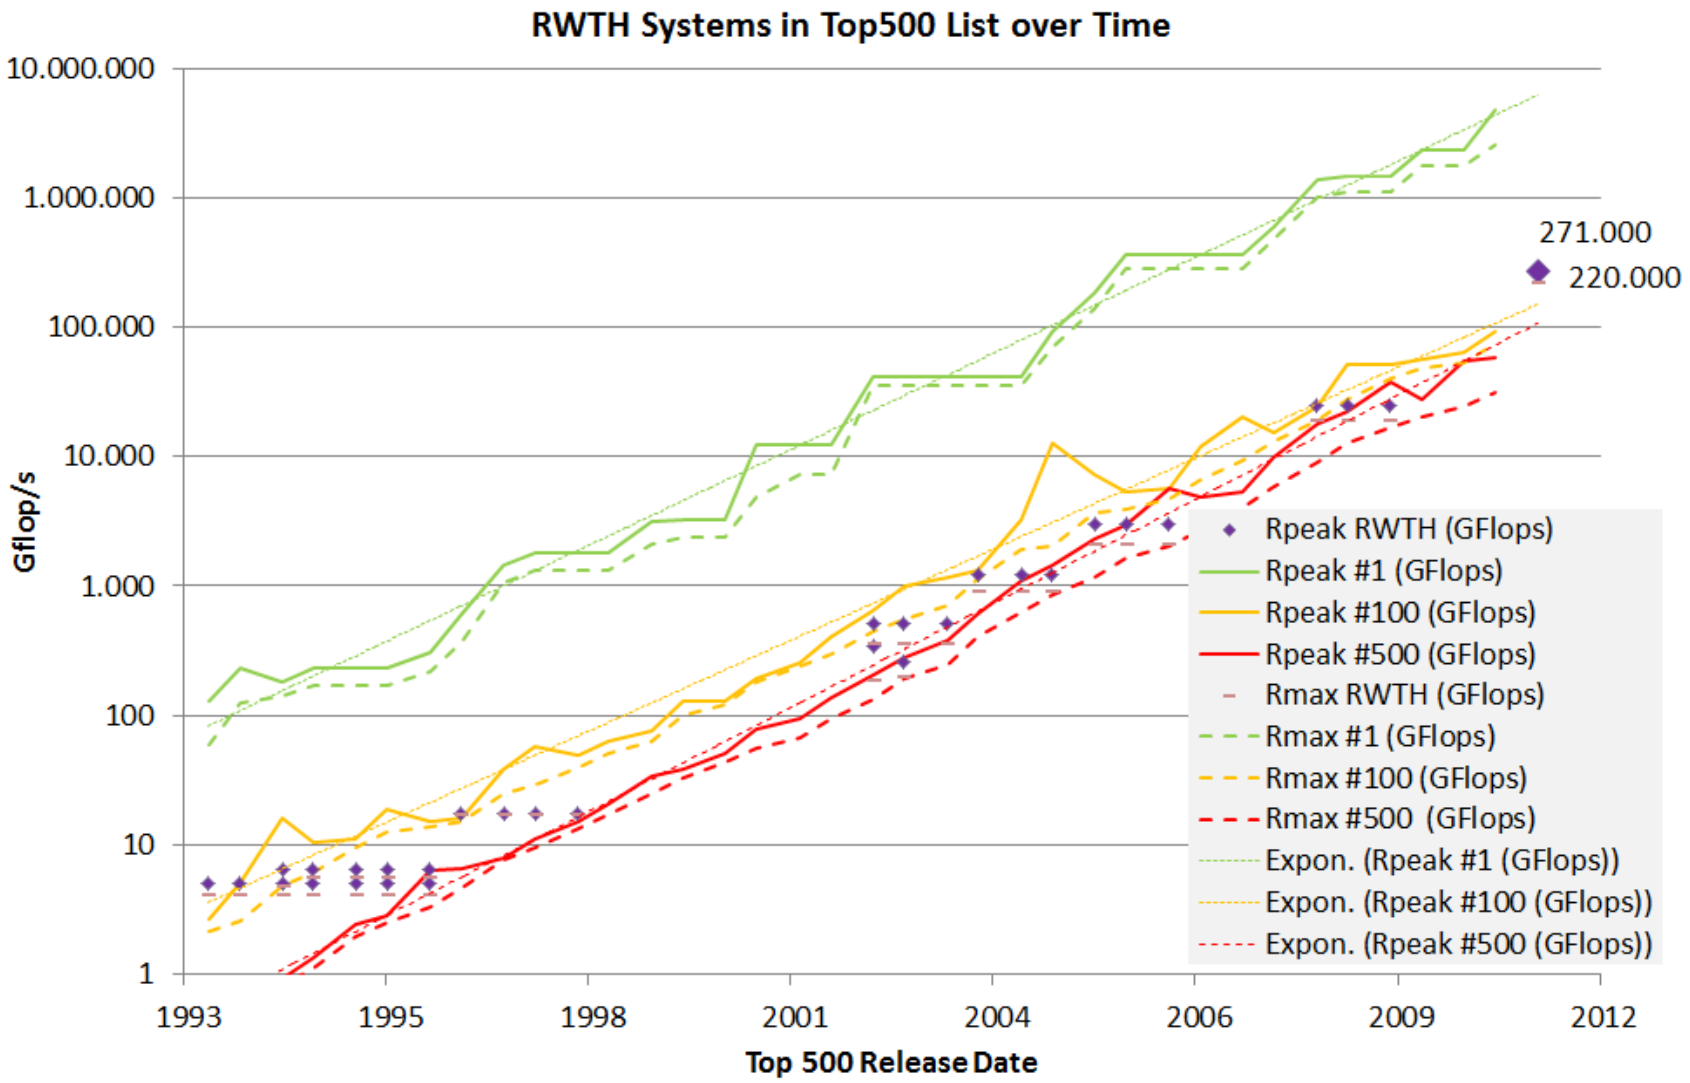

Growth of #OS instances / #processor cores

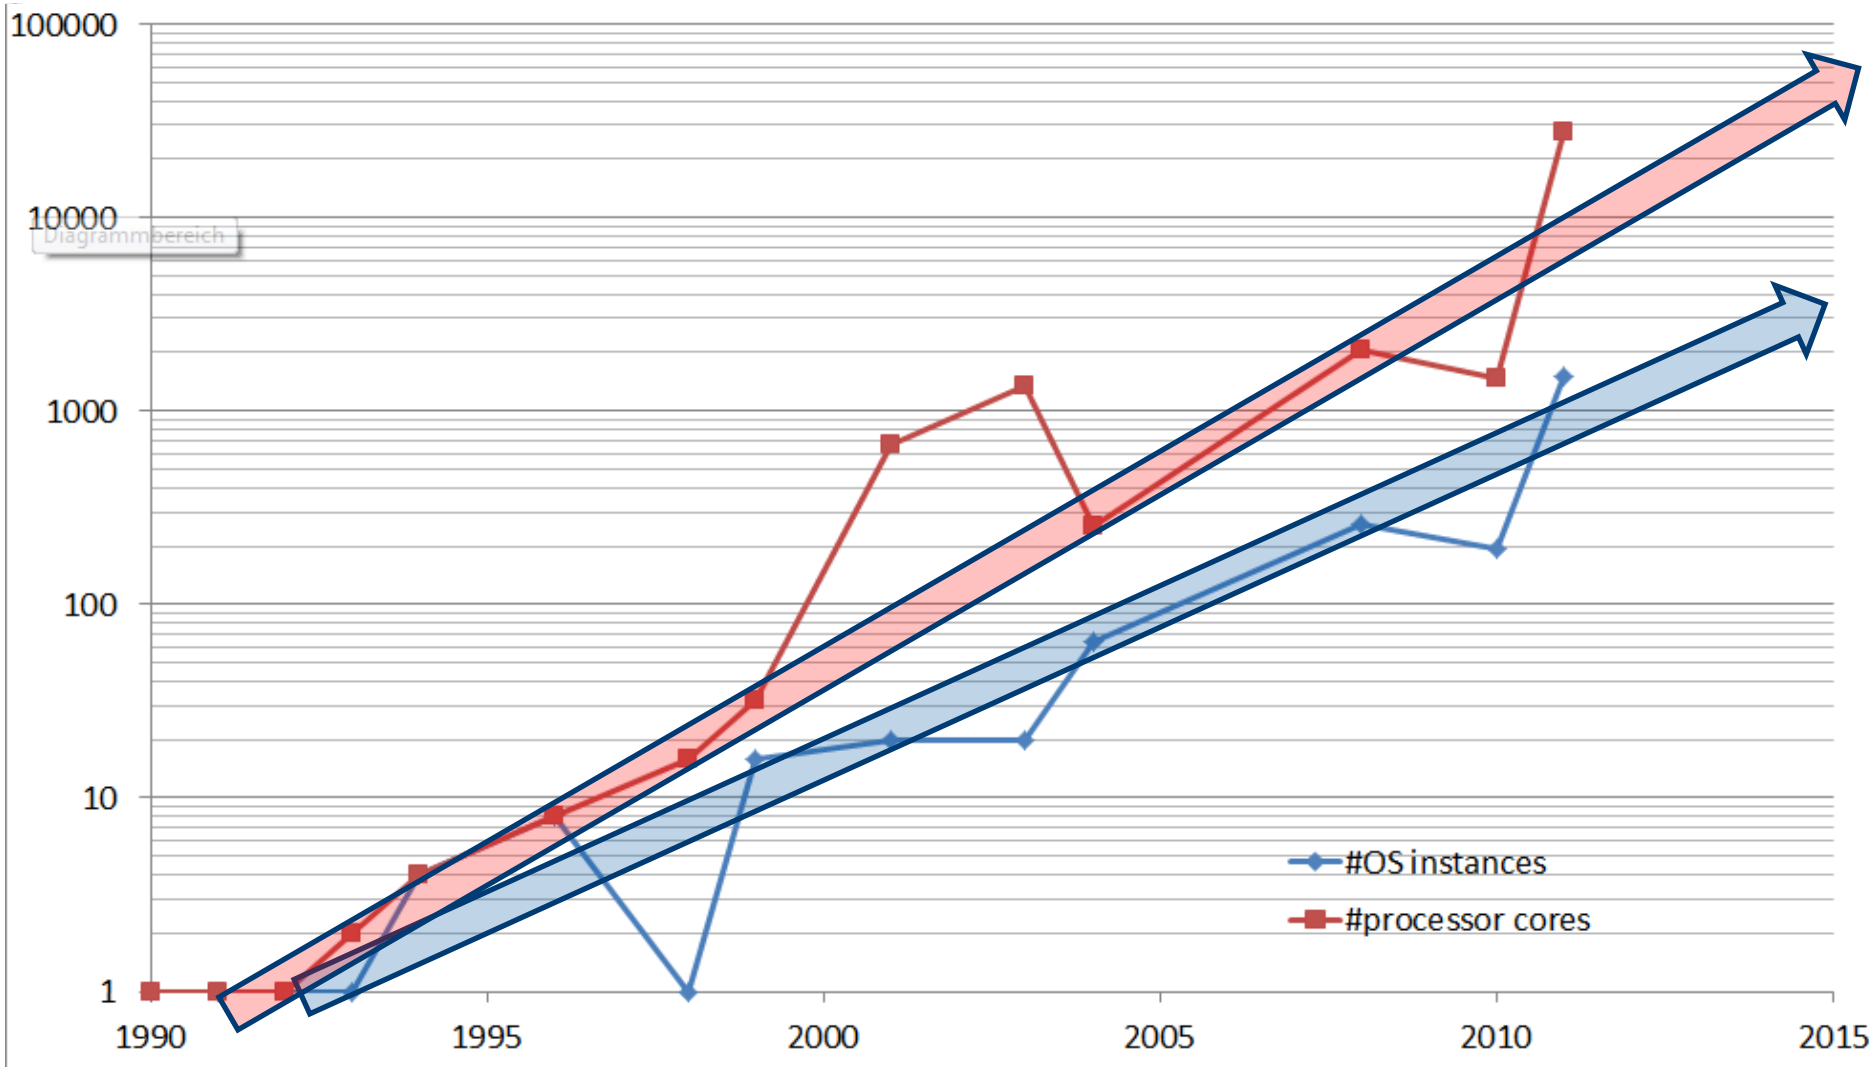

Computer Architecture: Naive Perspective

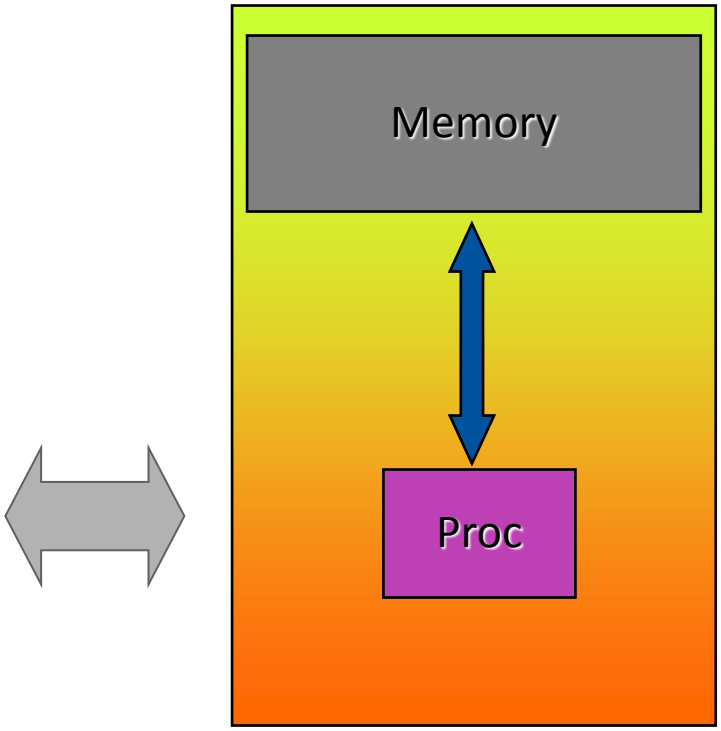

Main memory to store data and programs.

Processor to fetch program from memory, and execute program instructions:
Load data from memory, process data and write results back to memory.

Input/ output
is not covered here.

Accessing memory takes time.
Today memory bandwidth and latency frequently is a severe bottleneck!

Computer Architecture: Refined View with Multiple Cache Levels

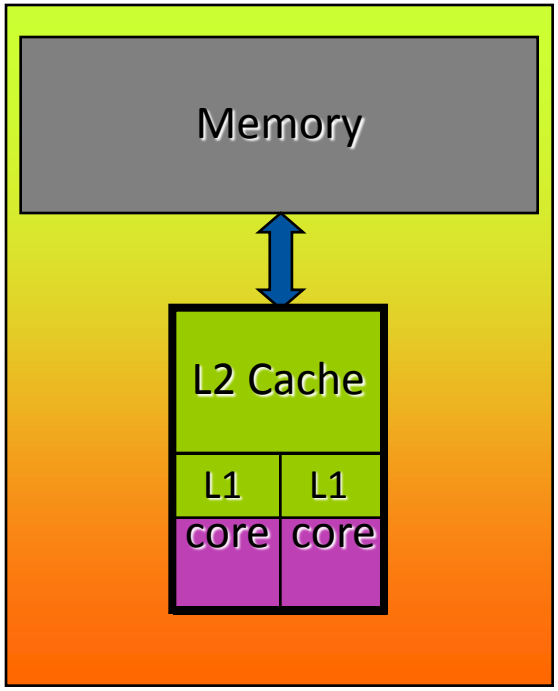

Since 2005/6 Intel and AMD are producing dualcore processors for the mass market.

Today mulitcore processors are ubiquitious.
Currently 4 to 12 cores per chip are quite common.

Caches have been employed since long to remedy the memory bottleneck to a certain degree.
But with a growing number of cores, the memory bottleneck is still growing !

Cluster of Multiprocessor Nodes with Multicore Processors

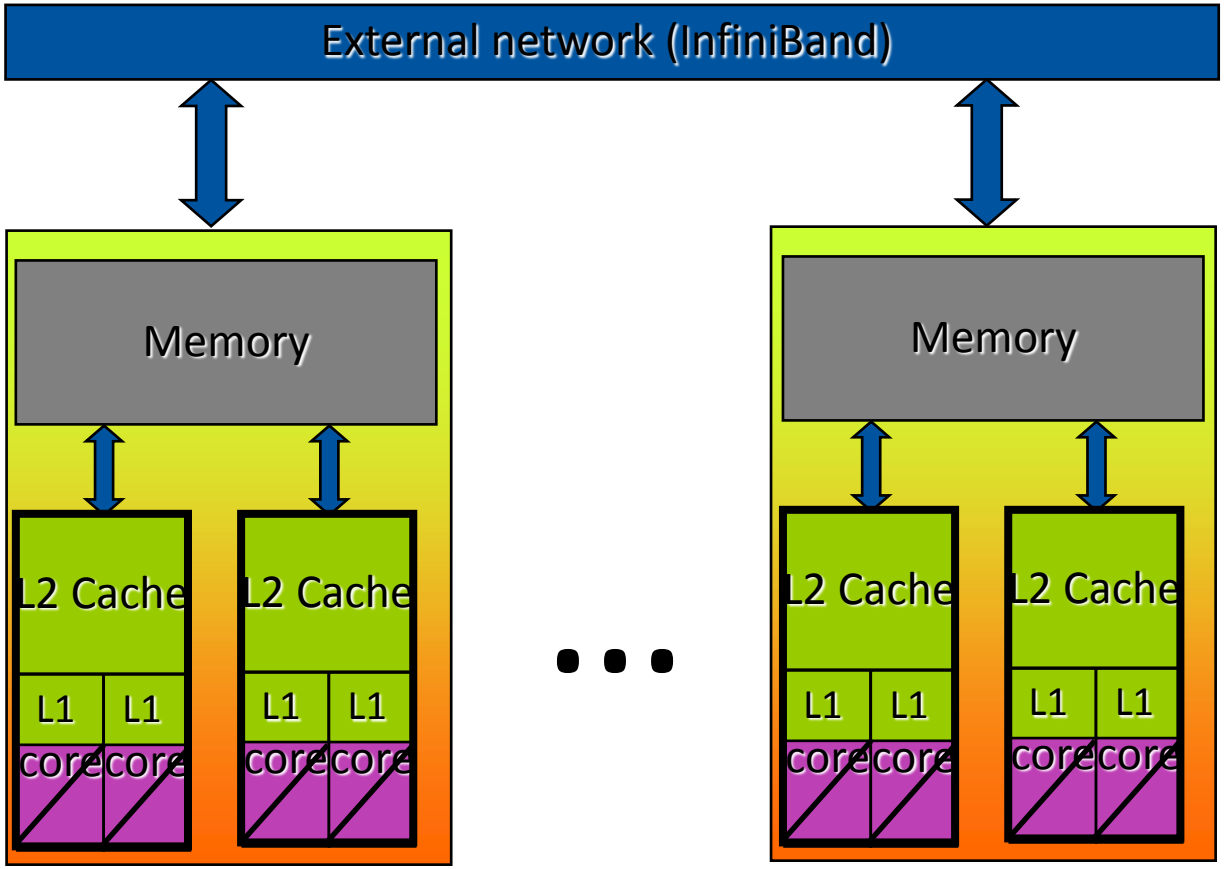

For High Performance Comuting (HPC) large clusters of many nodes with multiple multicore processors are connected by fast networks like InifiniBand.

Each node is a shared memory parallel computer where all cores of all processors have access to one main memory.

Bull BCS System

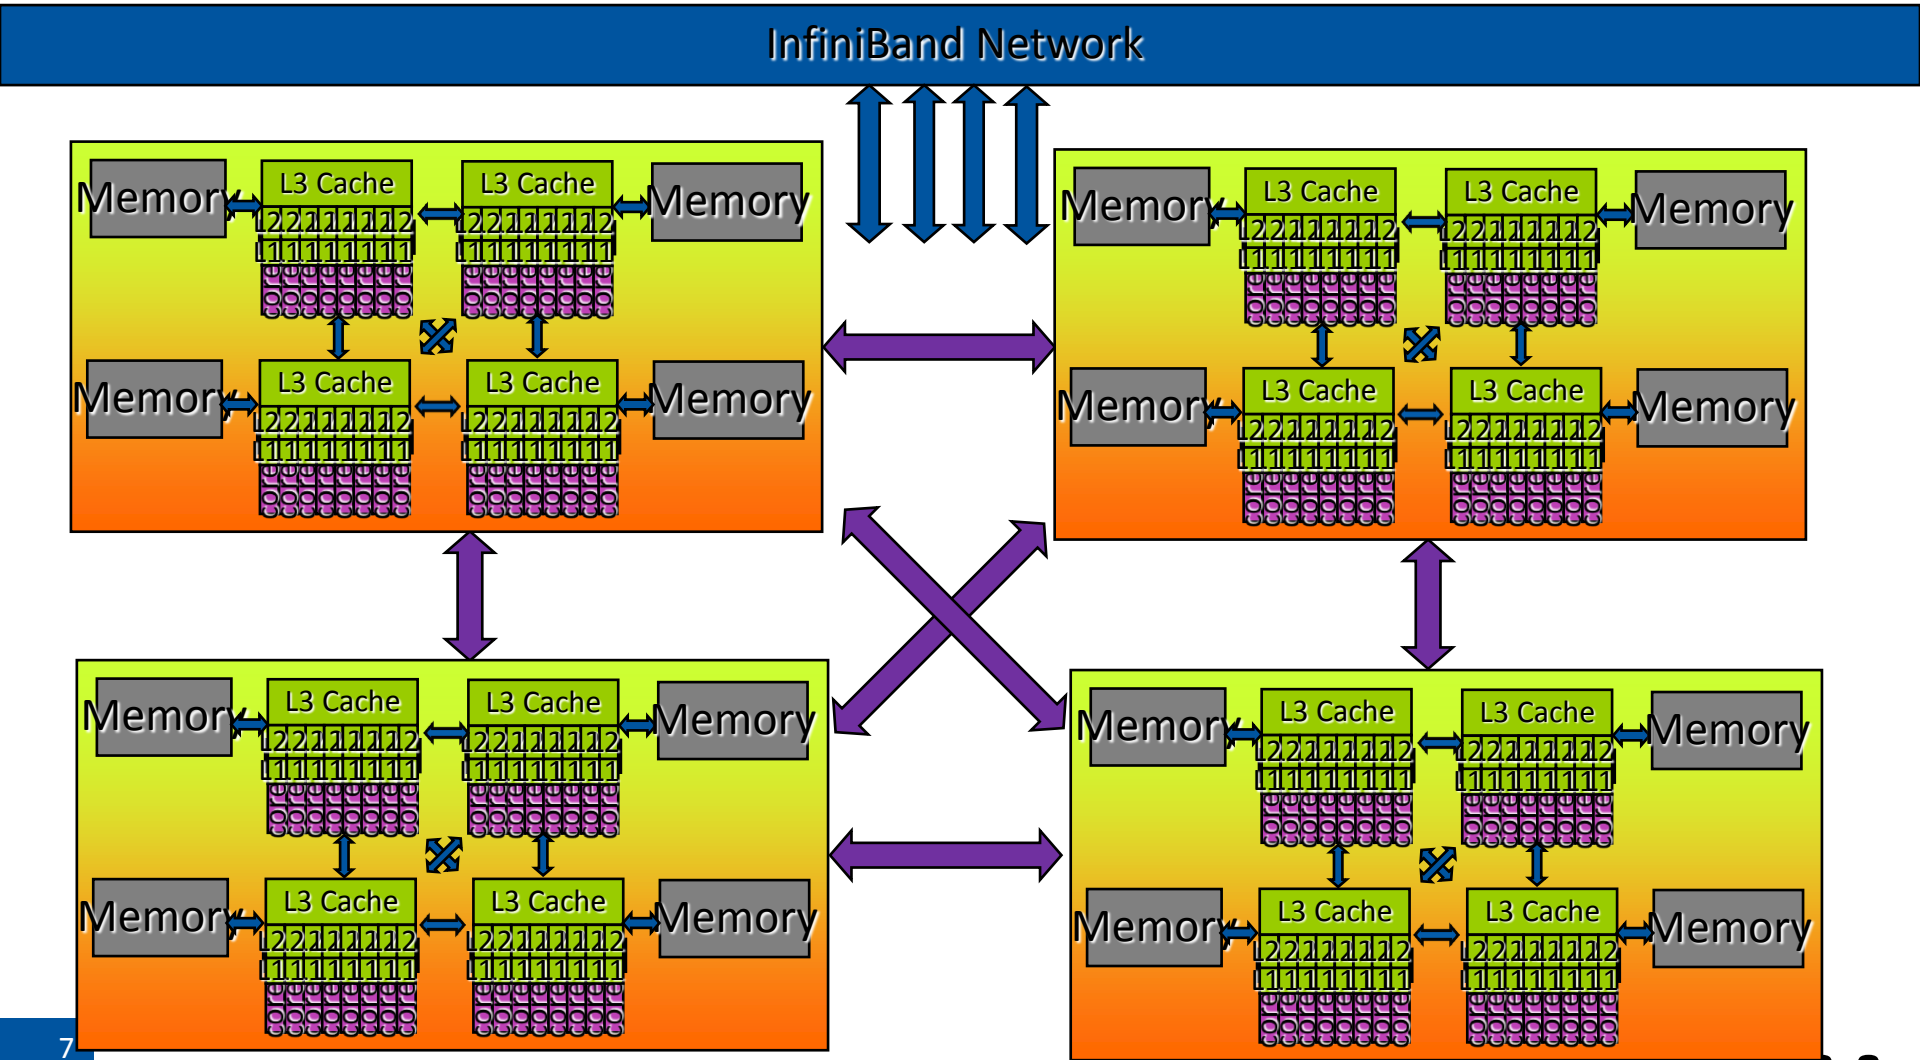

Message Passing with MPI
On Distributed Memory Parallel Computers

Typically, when using Message Passing with MPI, one MPI process runs on each processor core

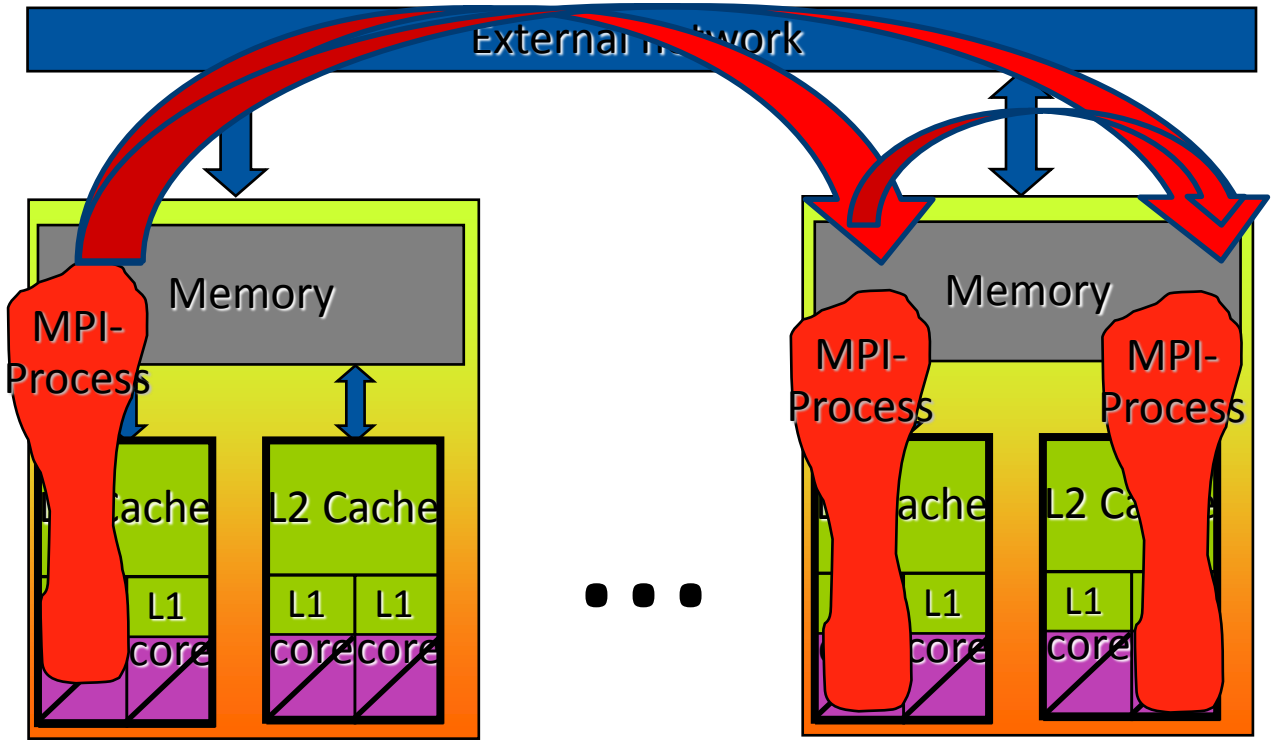

MPI is the de-facto standard for message passing.

MPI is a program library plus a mechanism to launch multiple cooperating executable progams.

Typically it is the same binary, which is started on multiple processors.

(SPMD=single program mutliple data paradigm)

Shared Memory Programming with OpenMP

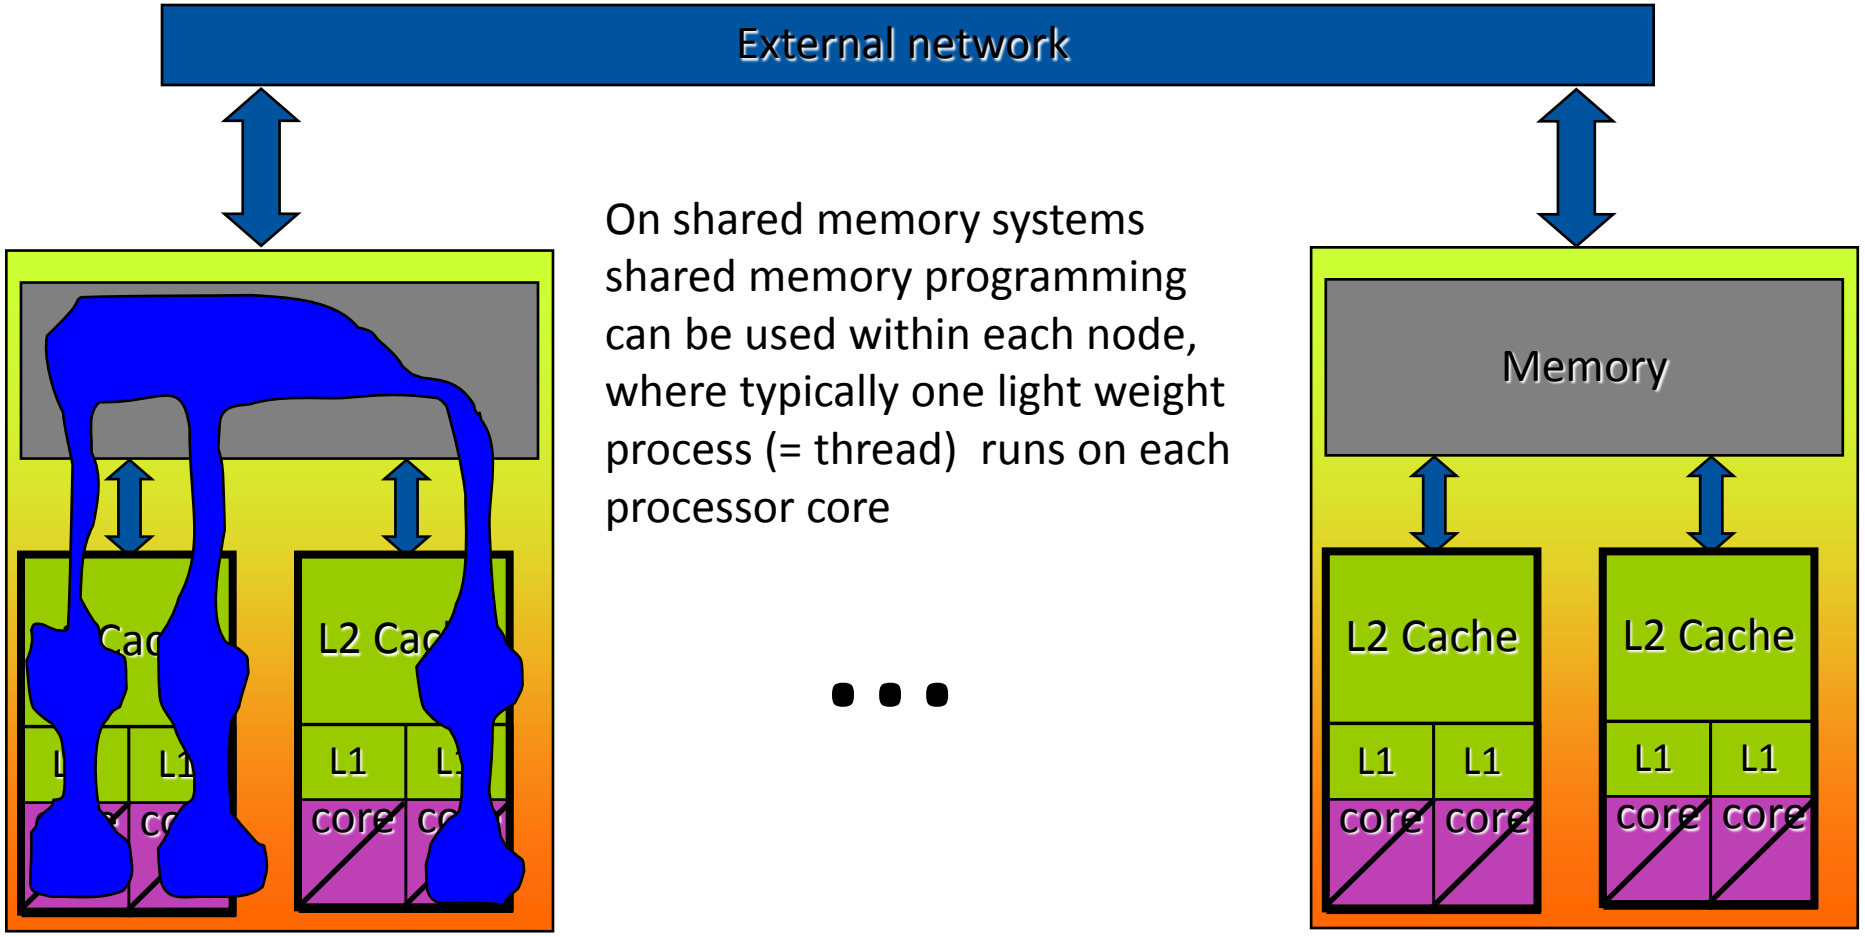

On shared memory systems shared memory programming can be used within each node, where typically one light weight process (= thread) runs on each processor core

• • •

MPI + OpenMP = Hybrid Parallelization

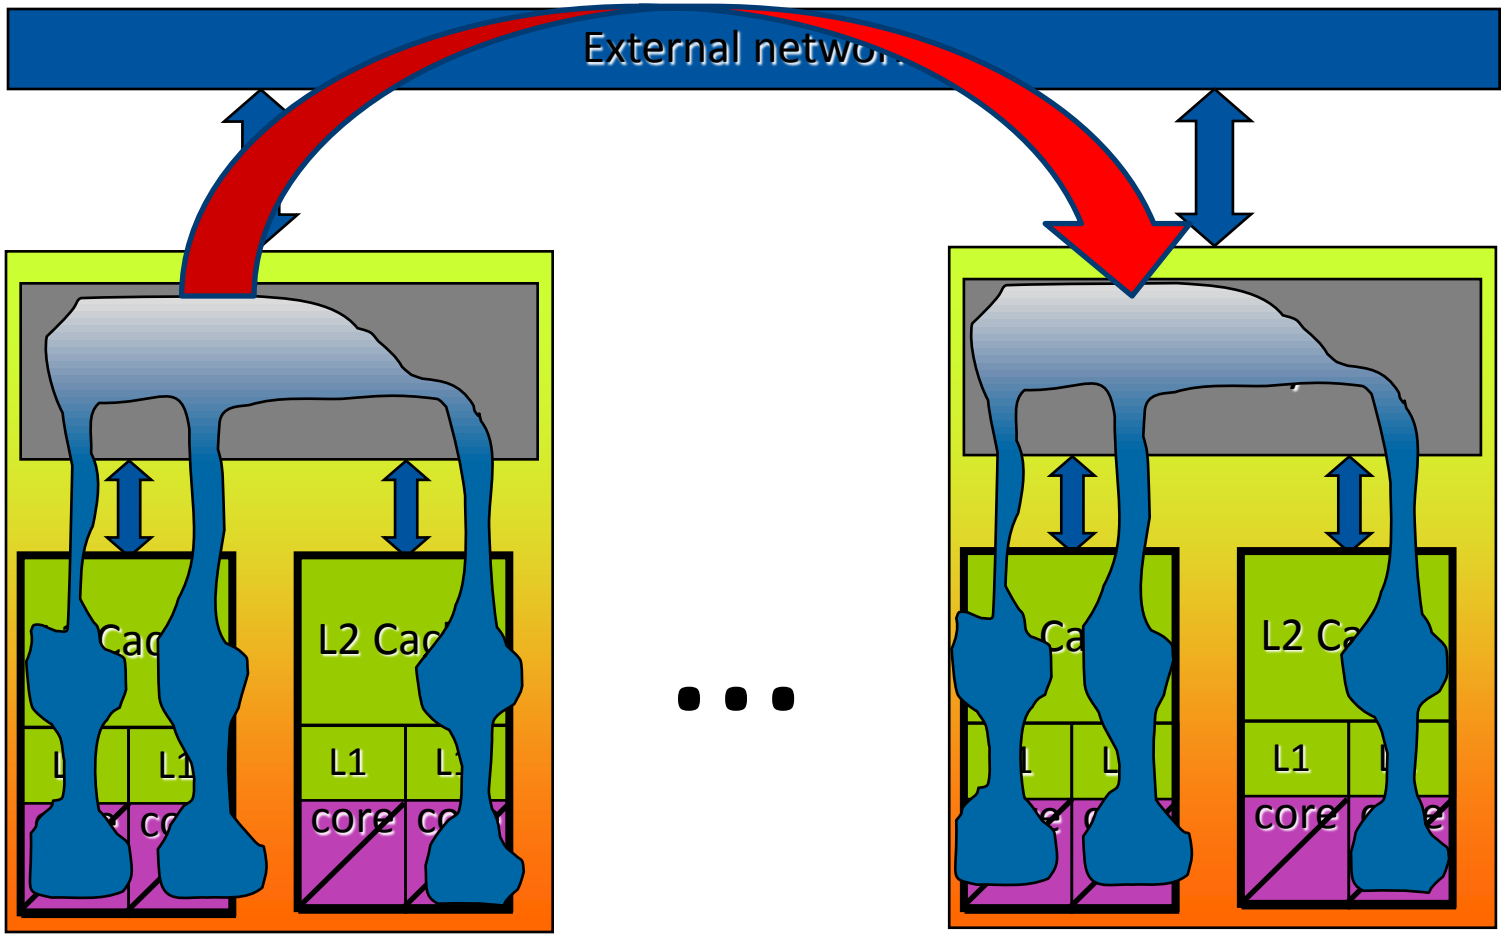

Virtual Shared Memory Programming with OpenMP (using Software from ScaleMP)

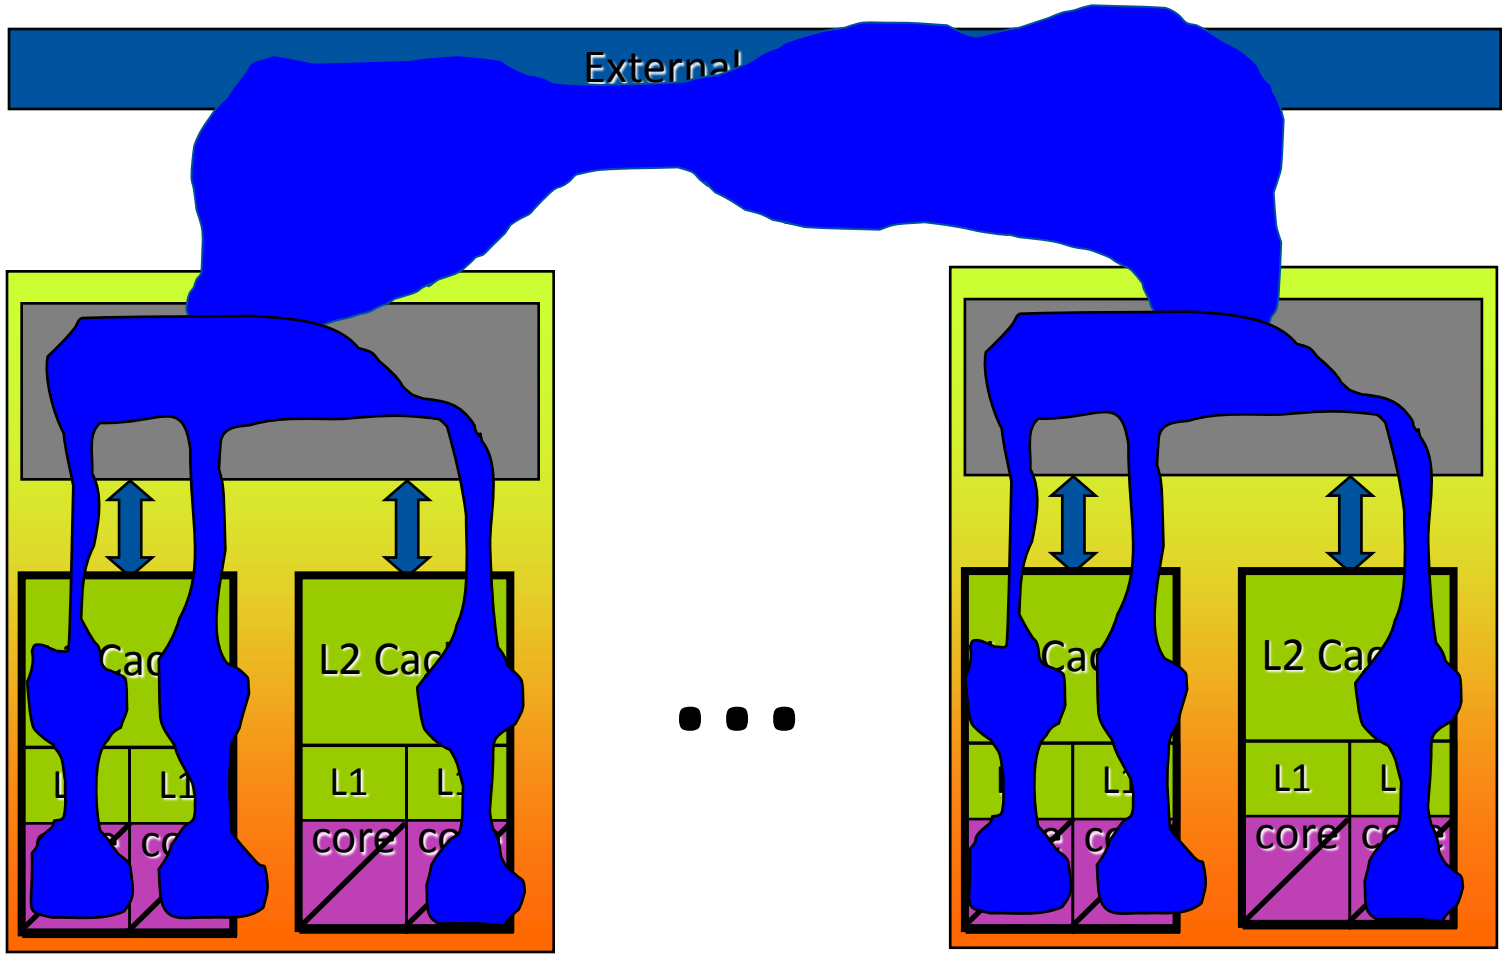

and then there are Accelerators / Coprocessors (NVIDIA GPGPUs, Intel Xeon Phi) …

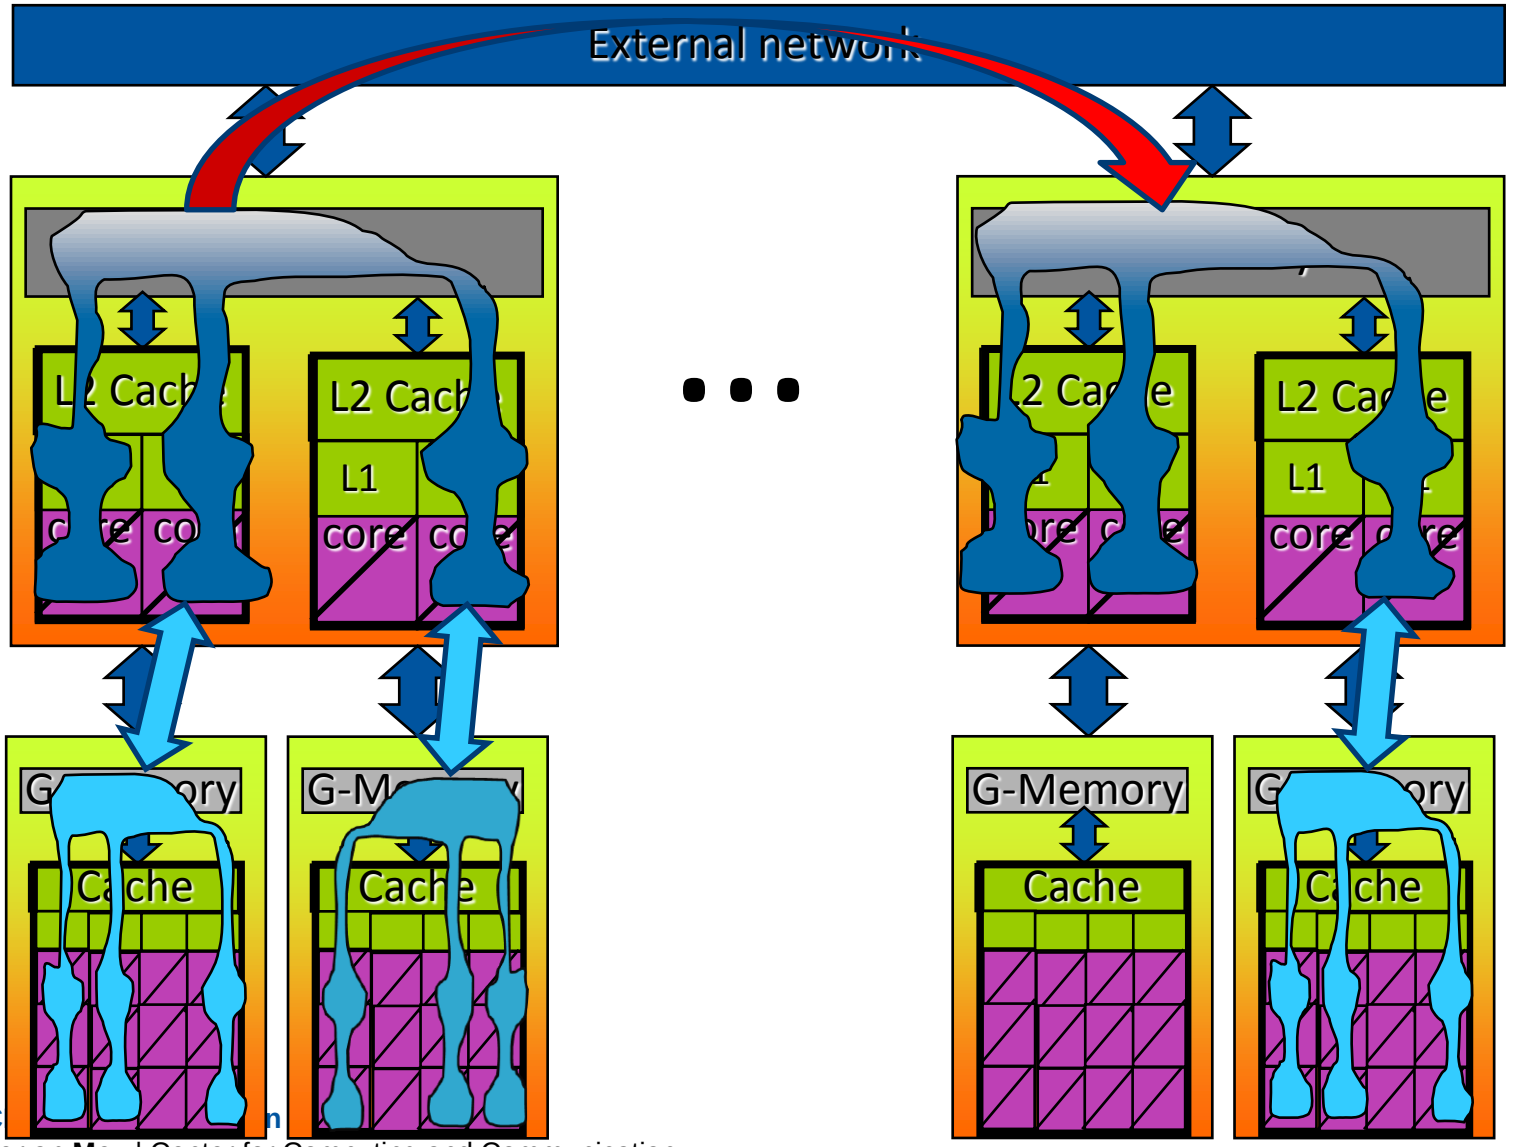

How to deal with this Complexity?

RWTH AACHEN UNIVERSITY

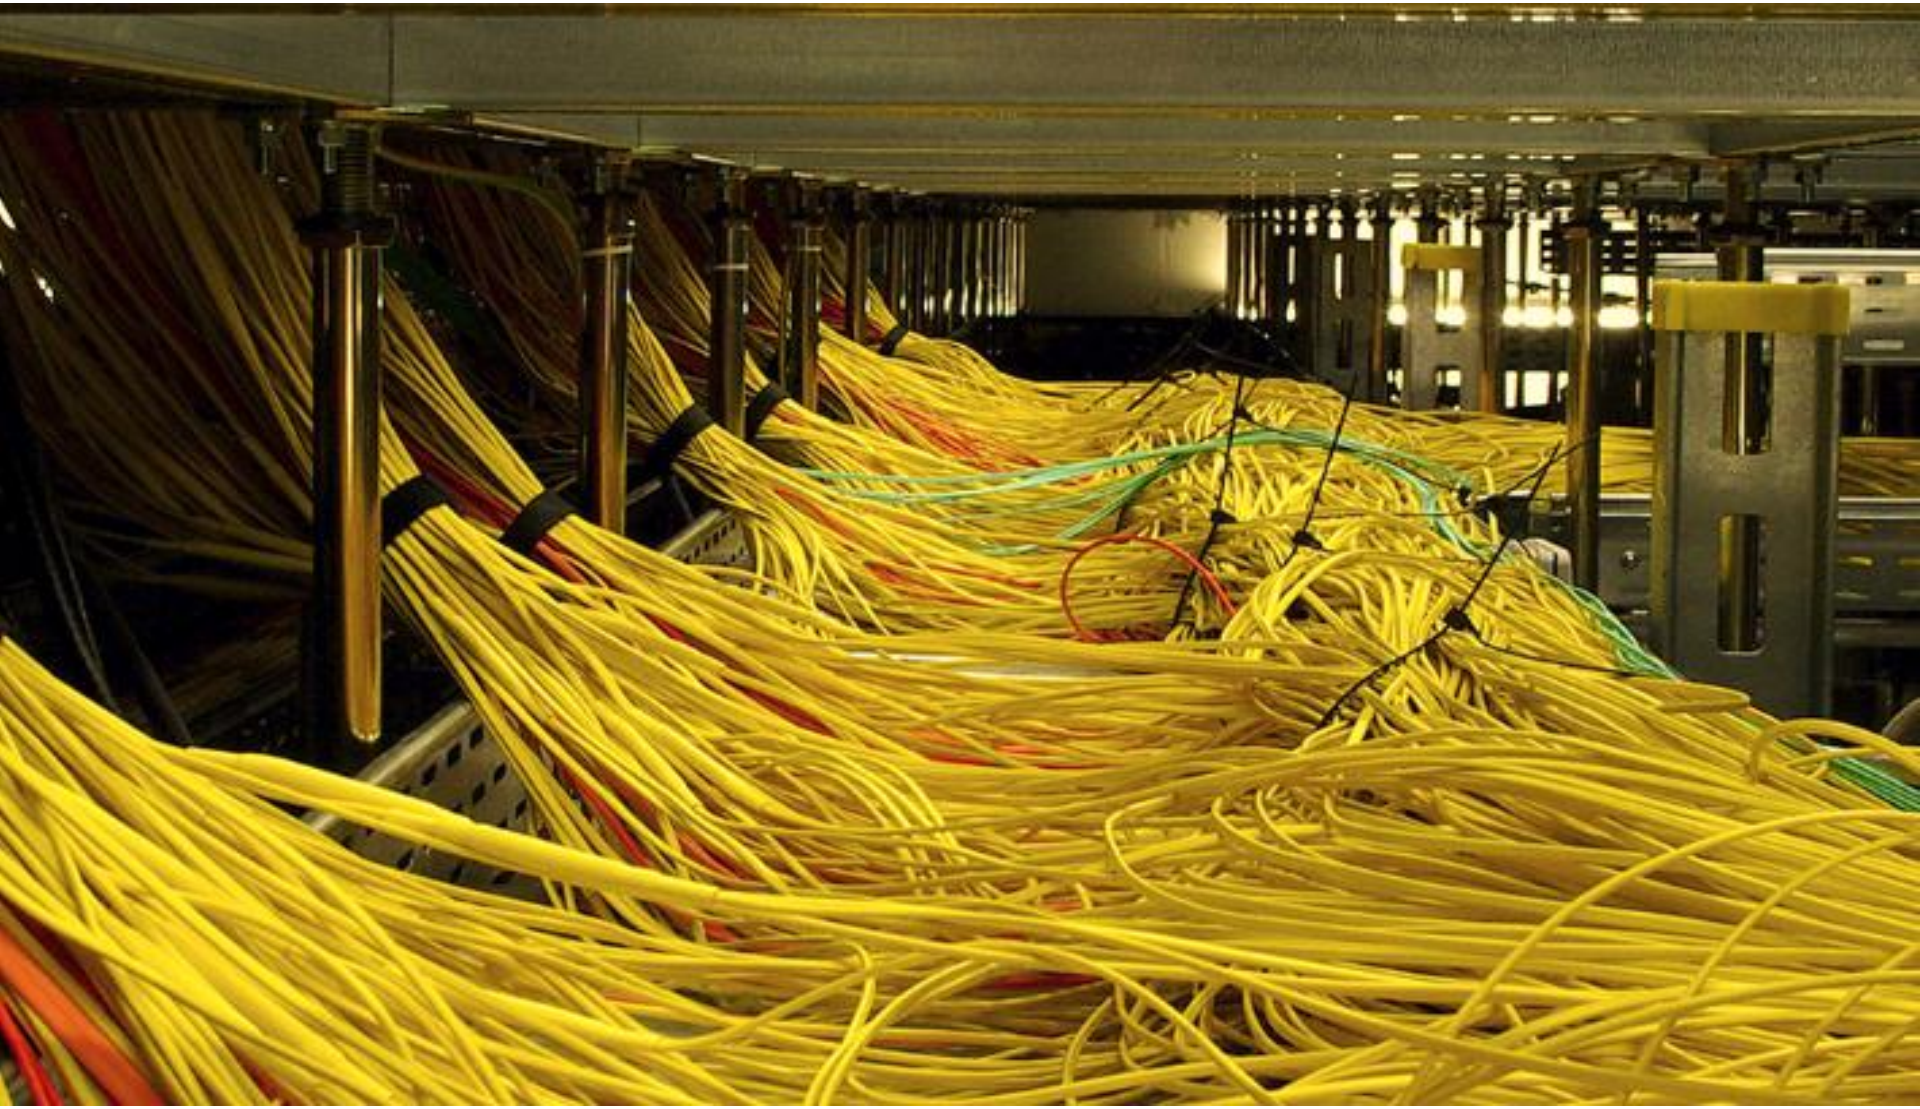

PPCES 2013 - Overview

Part I **Monday, March 11,** **14-17:30**	**Processor Architectures and Serial Programming** **incl Labs**
Part II **Tuesday, March 12,** **9-17:30**	**Message Passing with MPI** **incl Labs**
Part III **Wednesday, March 13,** **9-17:30**	**Shared Memory Programming with OpenMP** **incl Labs**
Part IV **Thursday, March 14,** **9-17:30**	**GPGPU Programming** **incl Labs**
Part V **Friday, March 15,** **9-15:30**	**Programming Intel Xeon Phi** **incl Labs + Code Tuning**

Parallel Programming for Computational Engineering & Science

HPC Seminar

March 11-15, 2013

RWTH Aachen University, Center for Computing and Communication

	Introduction Serial Programming, Tuning, Processor Architectures	Message Passing with MPI	Shared Memory Programming with OpenMP	GPGPU Programming with OpenACC	Intel Xeon Phi, Code tuning
	Mo, March 11	**Tu, March 12**	**We, March 13**	**Th, March 14**	**Fr, March 15**
9:00 - 10:30		**MPI Basics: Part I** • General Concepts • Point-to-Point communication *H. Iliev, RWTH*	**Introduction to Parallel Computing with OpenMP** *C. Terboven, RWTH*	**GPU Introduction, OpenACC Basics** *S. Wienke, RWTH*	**Programming the Intel® Xeon Phi™ Coprocessor** *T. Cramer, RWTH*
10:30 - 11:00		*coffee break*	*coffee break*	*coffee break*	*coffee break*
11:00 - 12:30		Lab: MPI	Lab: OpenMP Part I	Lab: OpenACC	Lab: Intel Xeon Phi, Bring your code tuning (not only for Xeon Phi)
12:30 - 14:00		*lunch break*	*lunch break*	*lunch break*	*lunch break*
14:00 - 15:30	**Welcome and Introduction** *15 min, Prof. M. Müller, RWTH* **Parallel Computing Architectures** *75 min, C. Terboven, RWTH*	**MPI Basics: Part II** *H. Iliev, RWTH*	**Getting OpenMP up to Speed** *R. v. d. Pas, Oracle*	**OpenACC Advanced** *70 min, S. Wienke, RWTH* **Case Study** *20 min, P. Springer, RWTH*	Lab: Intel Xeon Phi, Bring your code tuning (not only for Xeon Phi)
15:30 - 16:00	*coffee break*		*coffee break*	*coffee break*	
16:00 - 17:30	**The RWTH Compute Cluster** *25 min, T. Warschko, Bull* **RZ Environment** *20 min, T. Cramer, RWTH* 45 min, Lab: Performance Tuning for Cached-based Systems	**Vampir** *15 min, H. Iliev (RWTH)* **TotalView** *30 min, T. Cramer, RWTH* *45 min*, Lab: MPI and Tools	**OpenMP Programming on ScaleMP** *25 min, Dirk Schmidl, RWTH* **Misc. Advanced OpenMP Programming** *20 min, C.Terboven, RWTH* *45 min*, Lab: OpenMP Part II	**Outlook on OpenMP for Accelerators** *30 min*, C. Terboven, RWTH *60 min*, Lab: OpenACC	
19:00 –			***Social event in Palladion*** Schmiedstraße 3		

Welcome to PPCES 2013