RWTH AACHEN UNIVERSITY

Parallel Programming Summer Course 2013

Introduction

Dieter an Mey
Center for Computing and Communication

RWTH Systems in Top 500 List

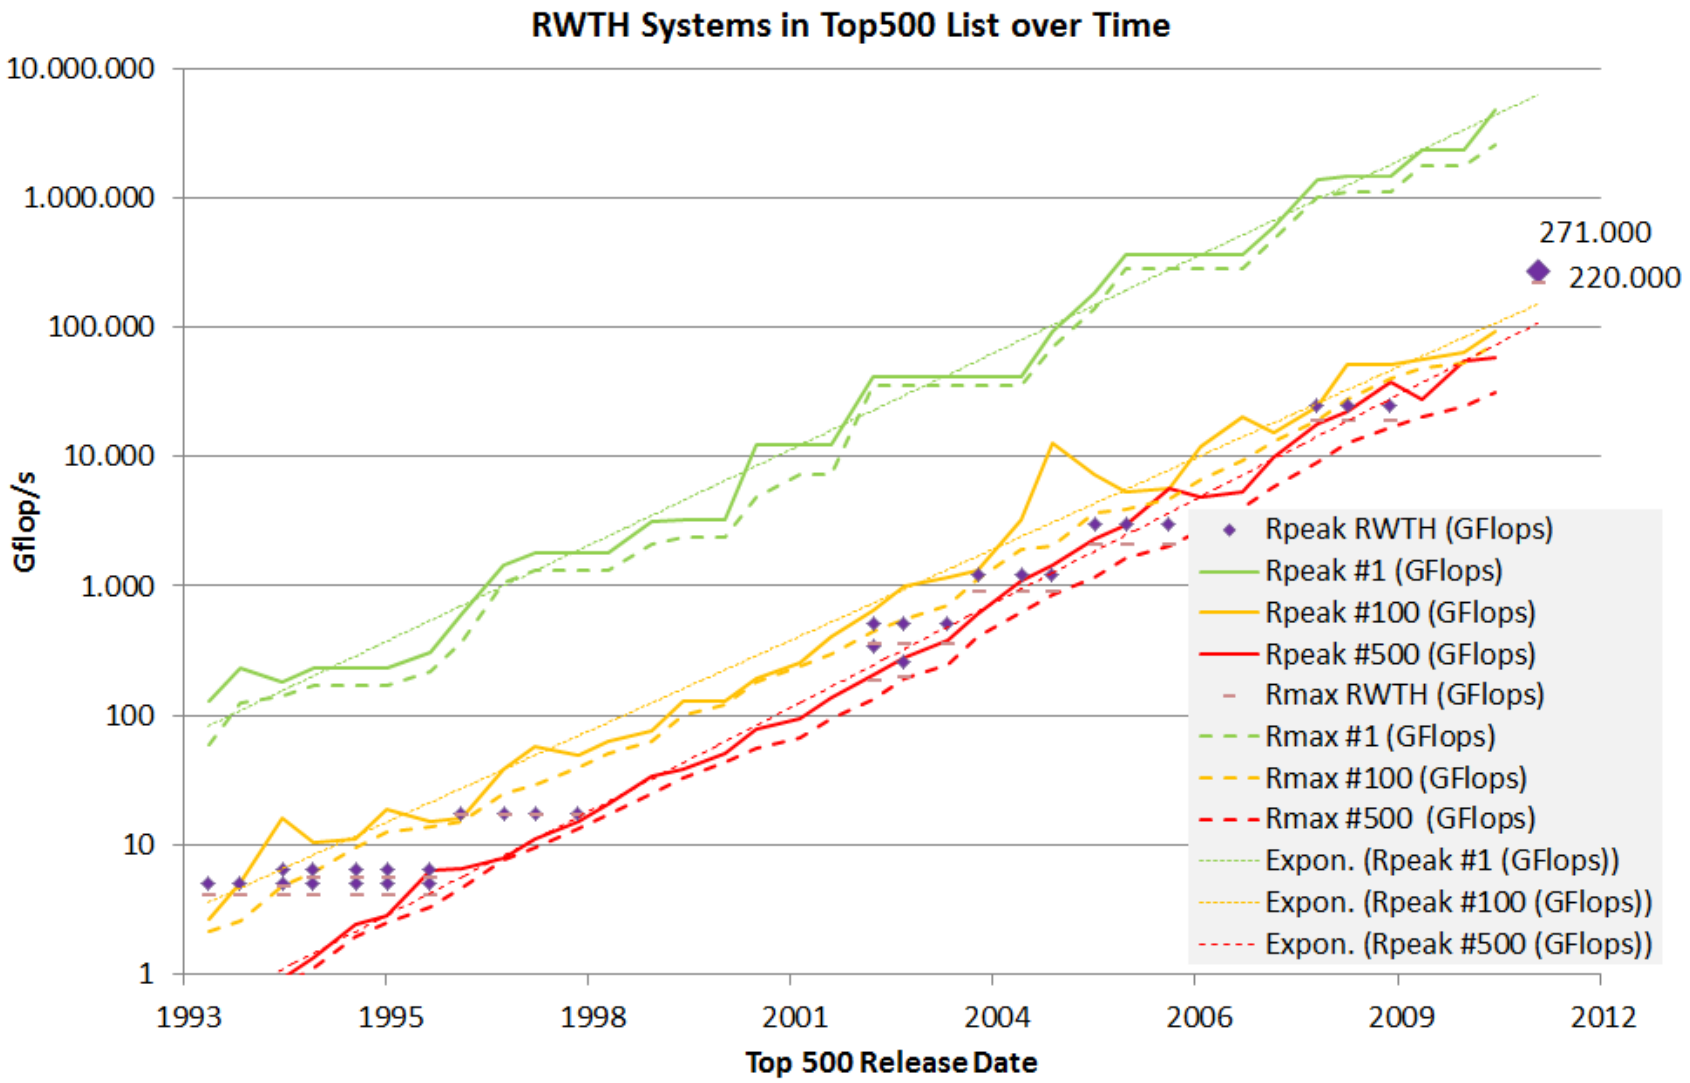

Growth of #OS instances / #processor cores

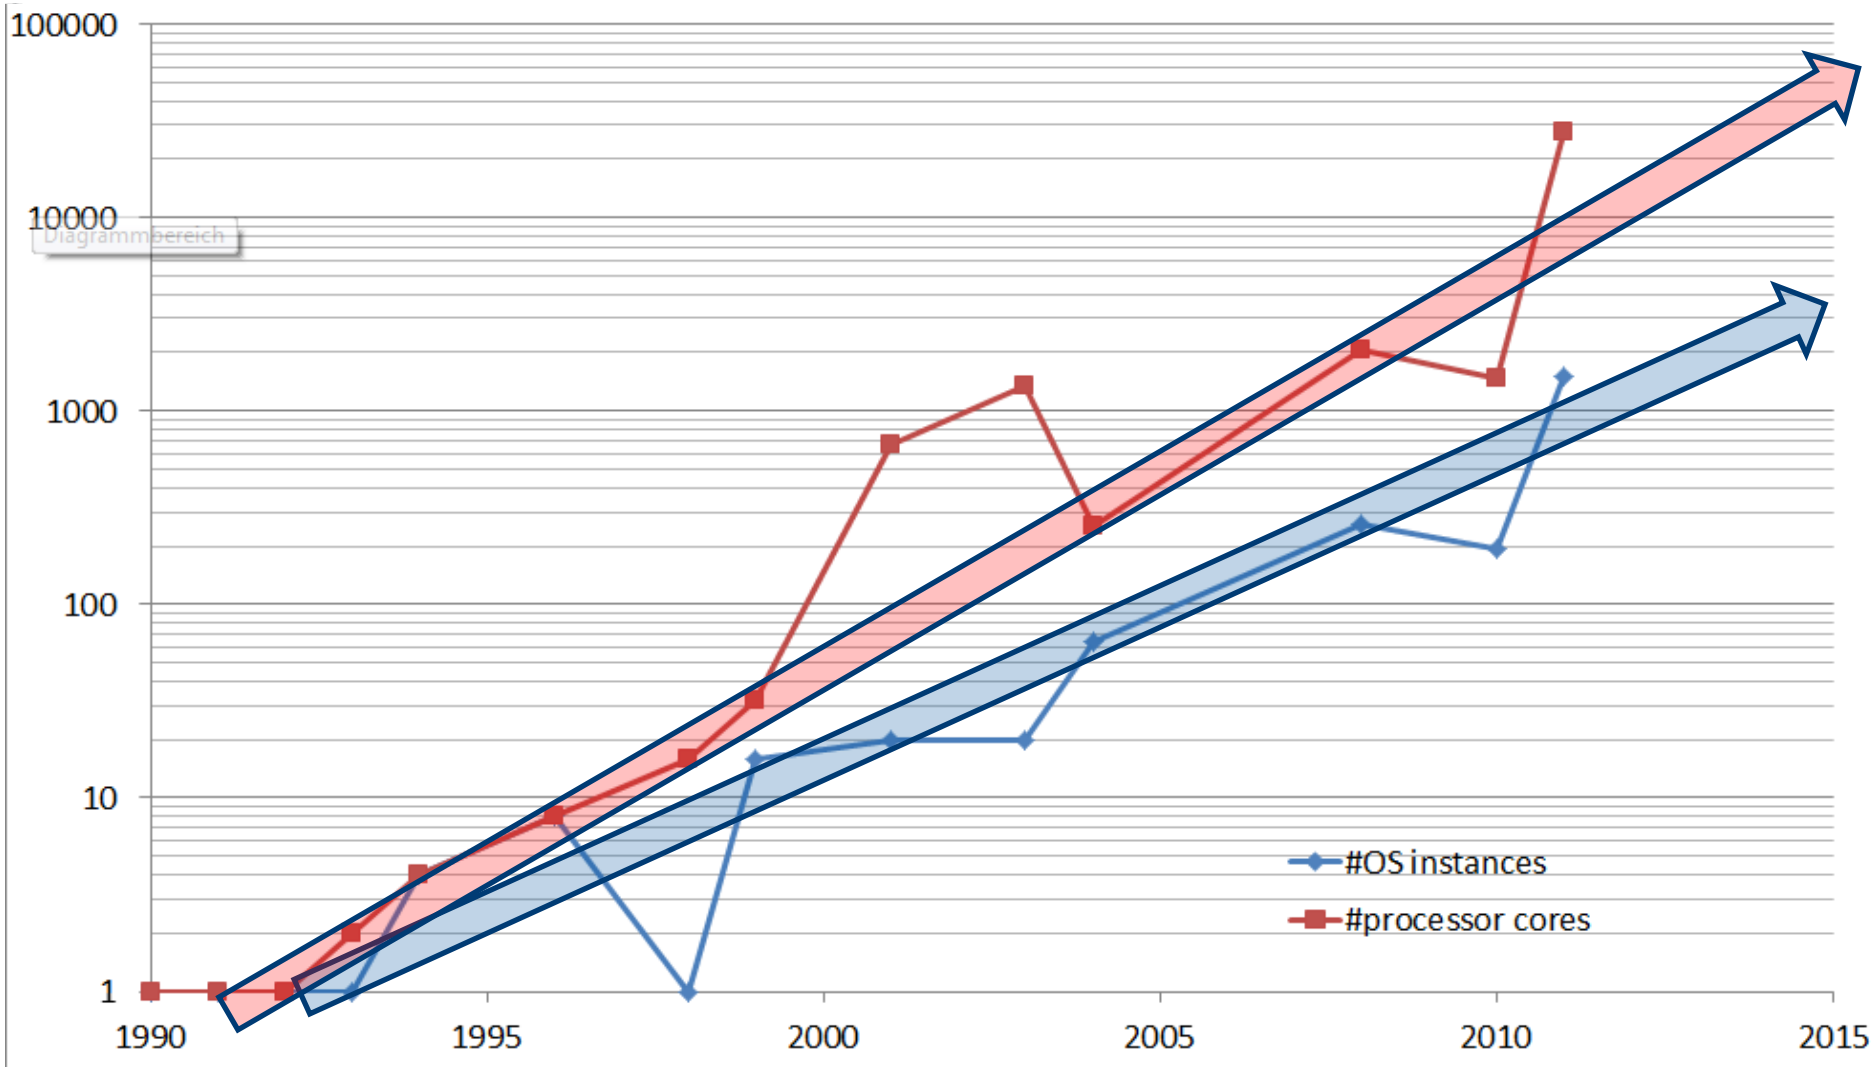

Computer Architecture: Naive Perspective

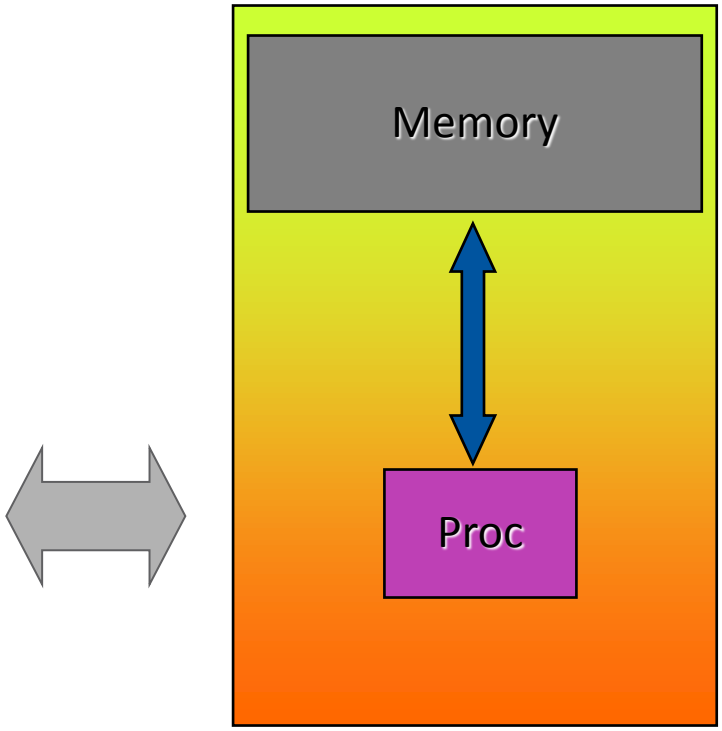

Main memory to store data and programs.

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

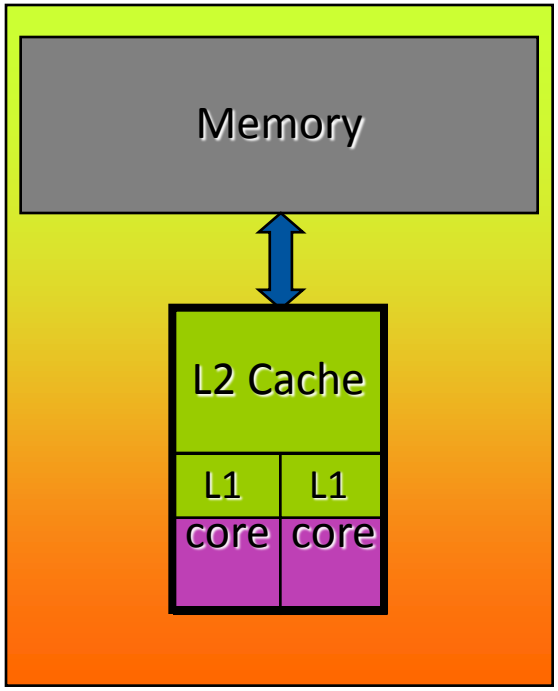

Since 2005/6 Intel and AMD are producing dualcore processors for the mass market.

Today mulitcore processors are ubiquitious.
Currently 4 to 16 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

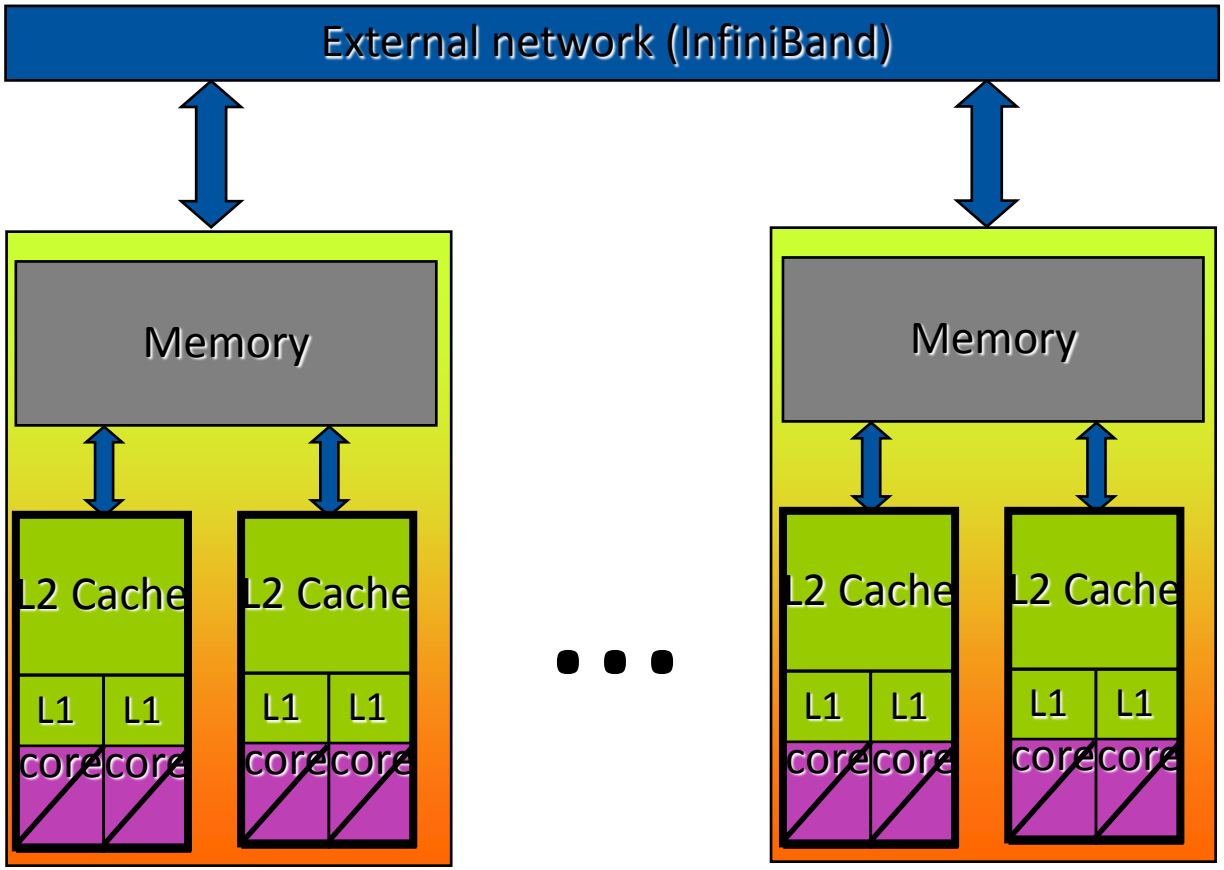

For High Performance Comuting (HPC) large clusters of many nodes with multiple multicore processors are connected by fast networks like InifiniBand.

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

Bull BCS System

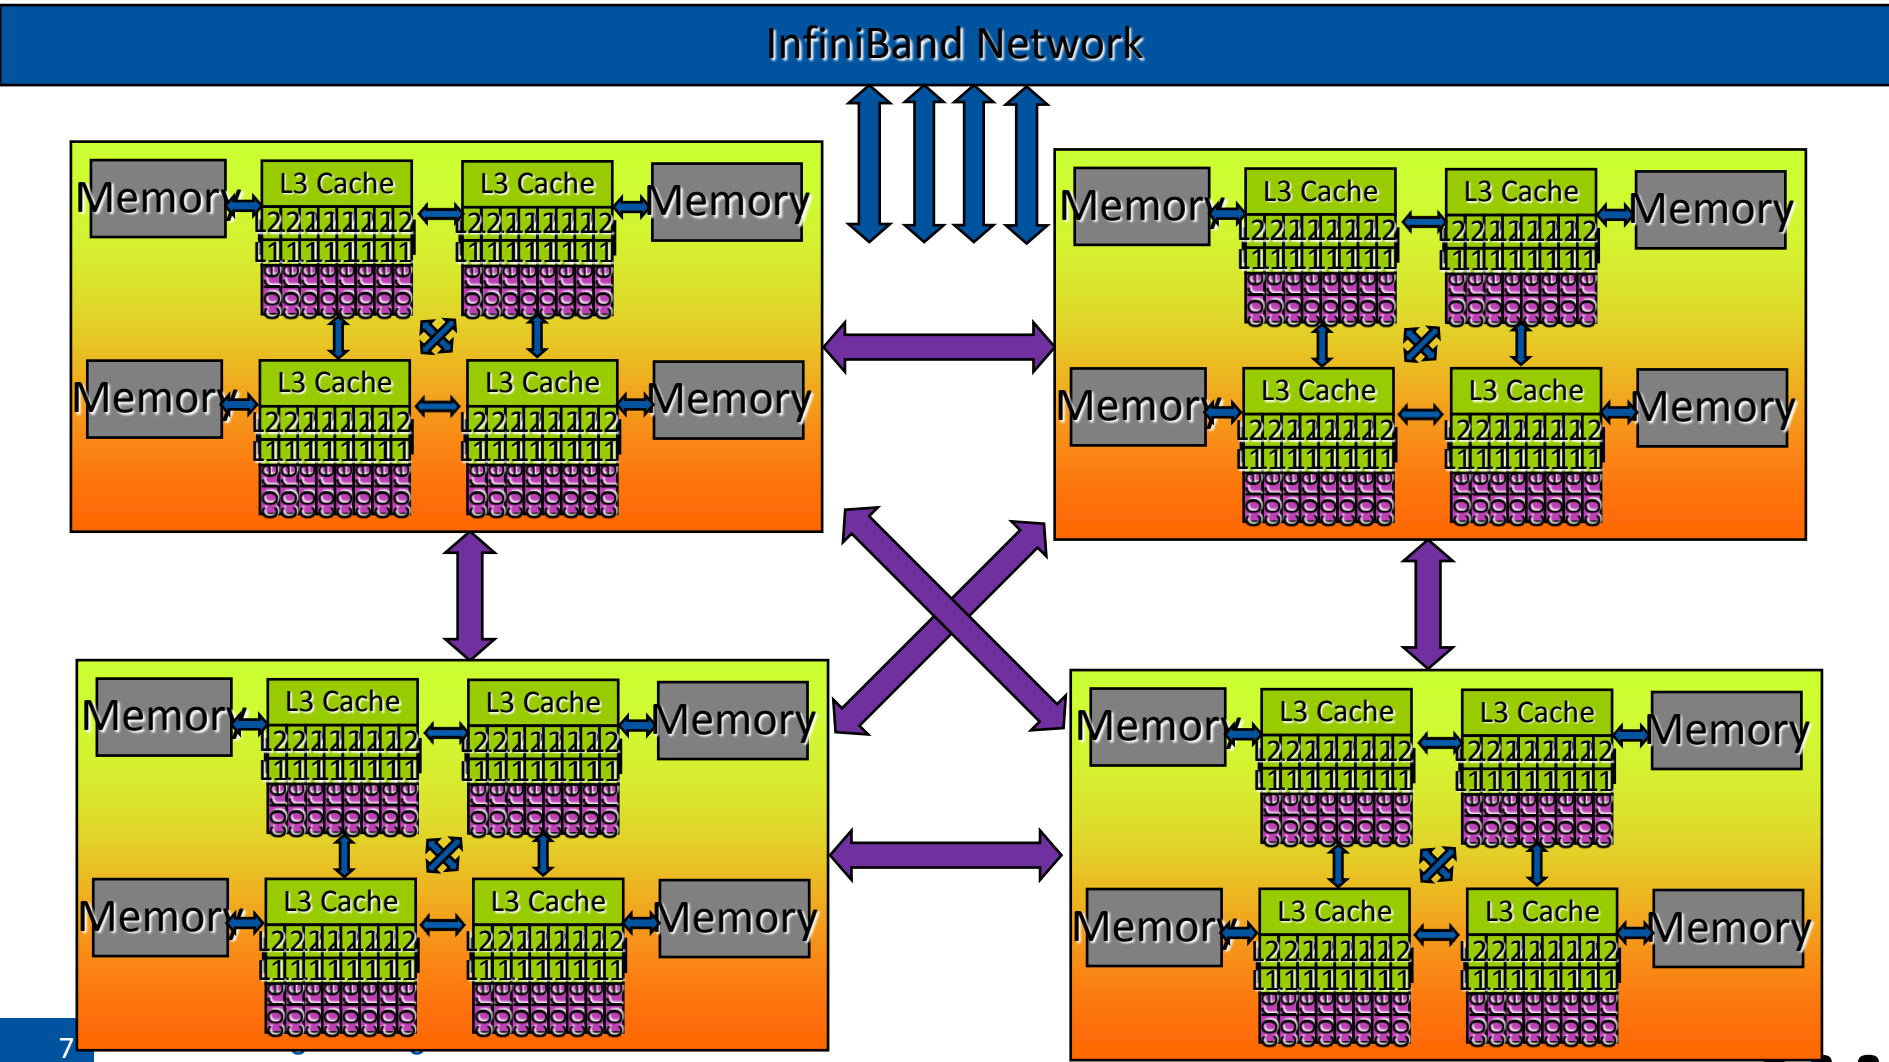

Message Passing with MPI
On Distributed Memory Parallel Computers

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

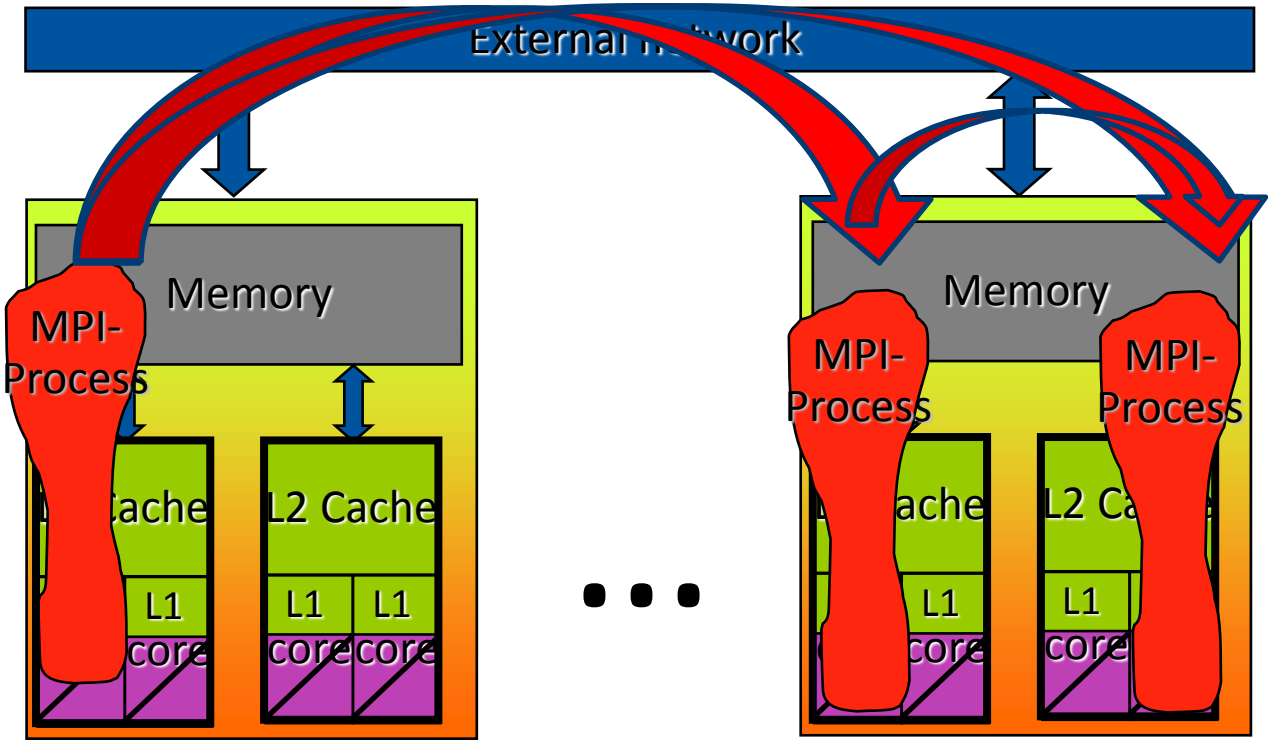

MPI is the de-facto standard for message passing.

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

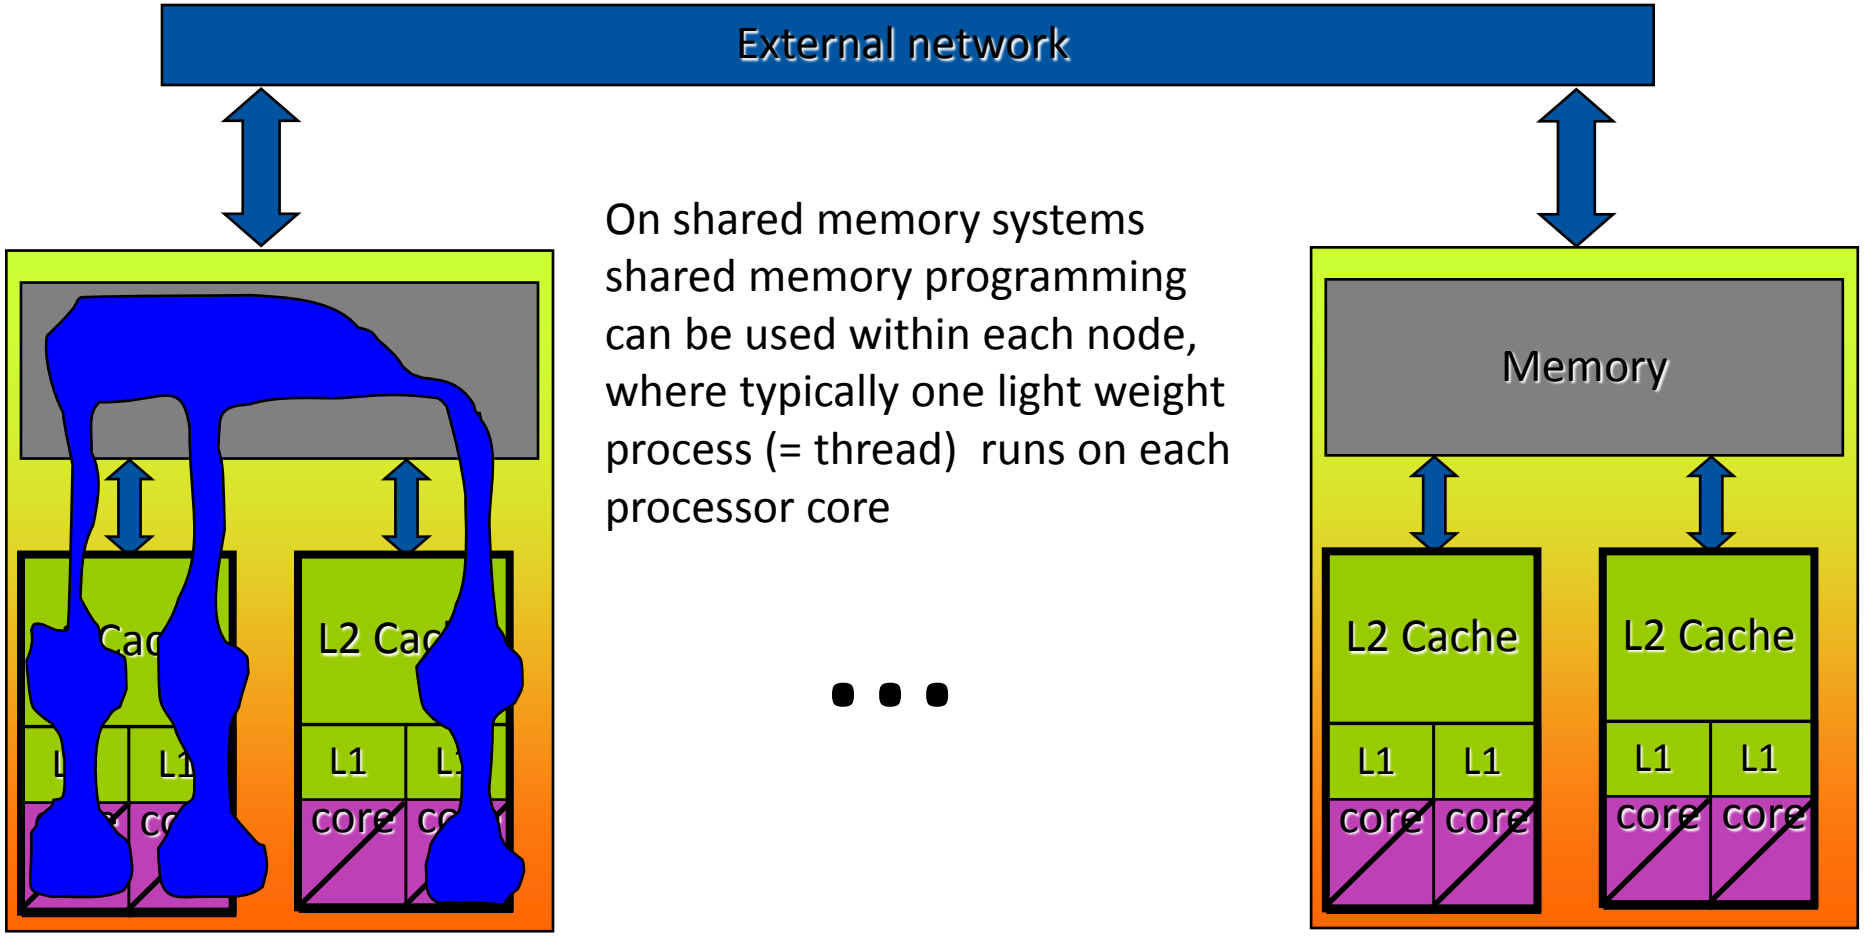

On shared memory systems shared memory programming can be used within each node, where typically one light weight process (= thread) runs on each processor core

• • •

MPI + OpenMP = Hybrid Parallelization

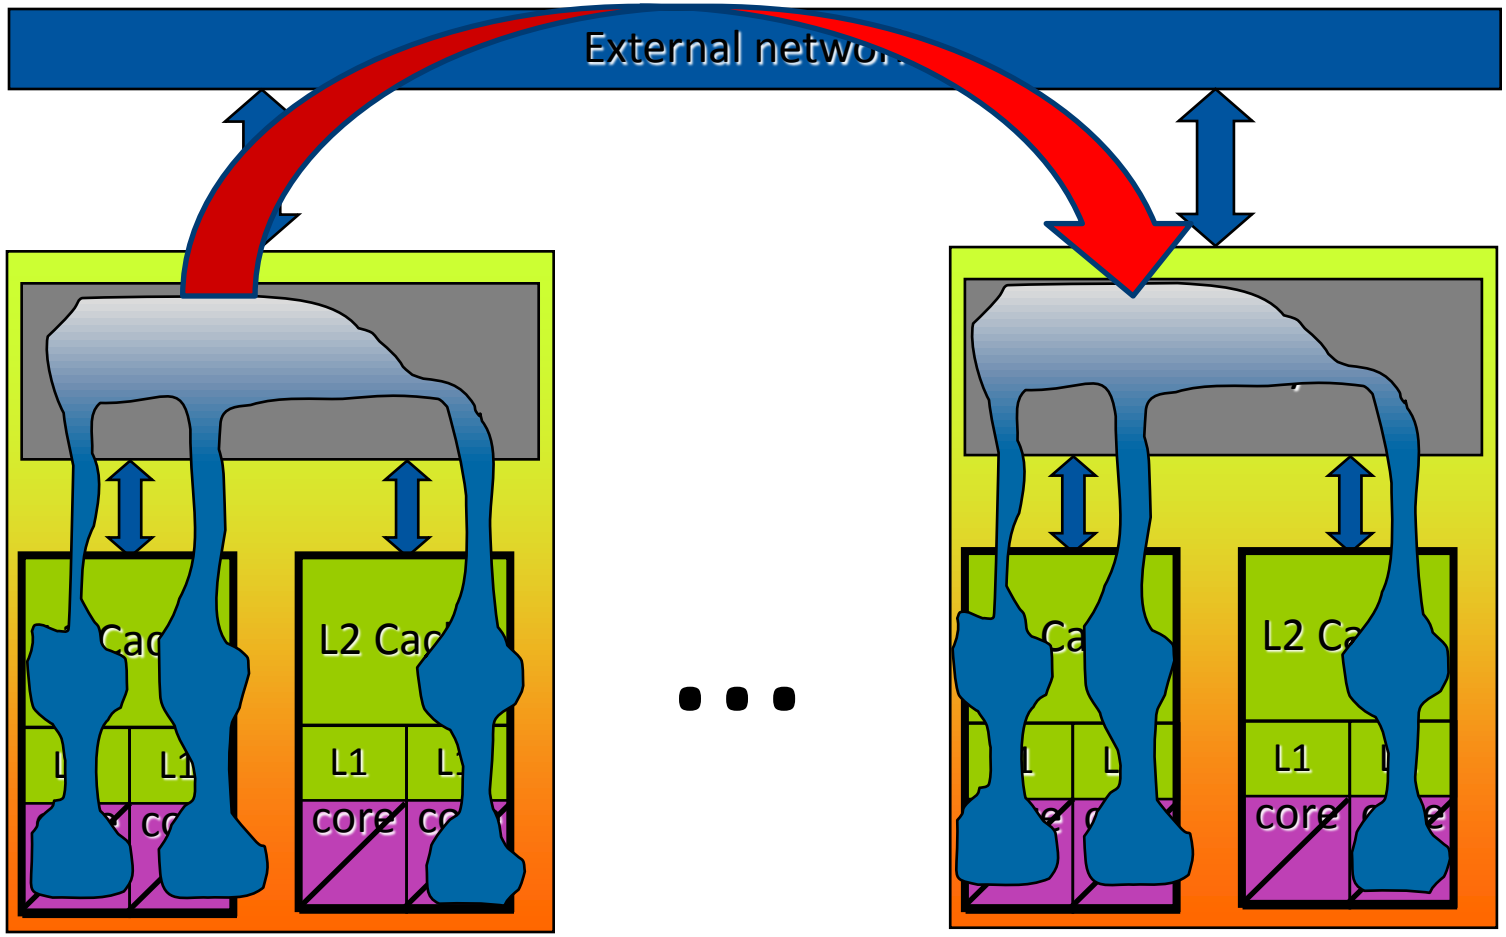

Virtual Shared Memory Programming with OpenMP (using Software from ScaleMP)

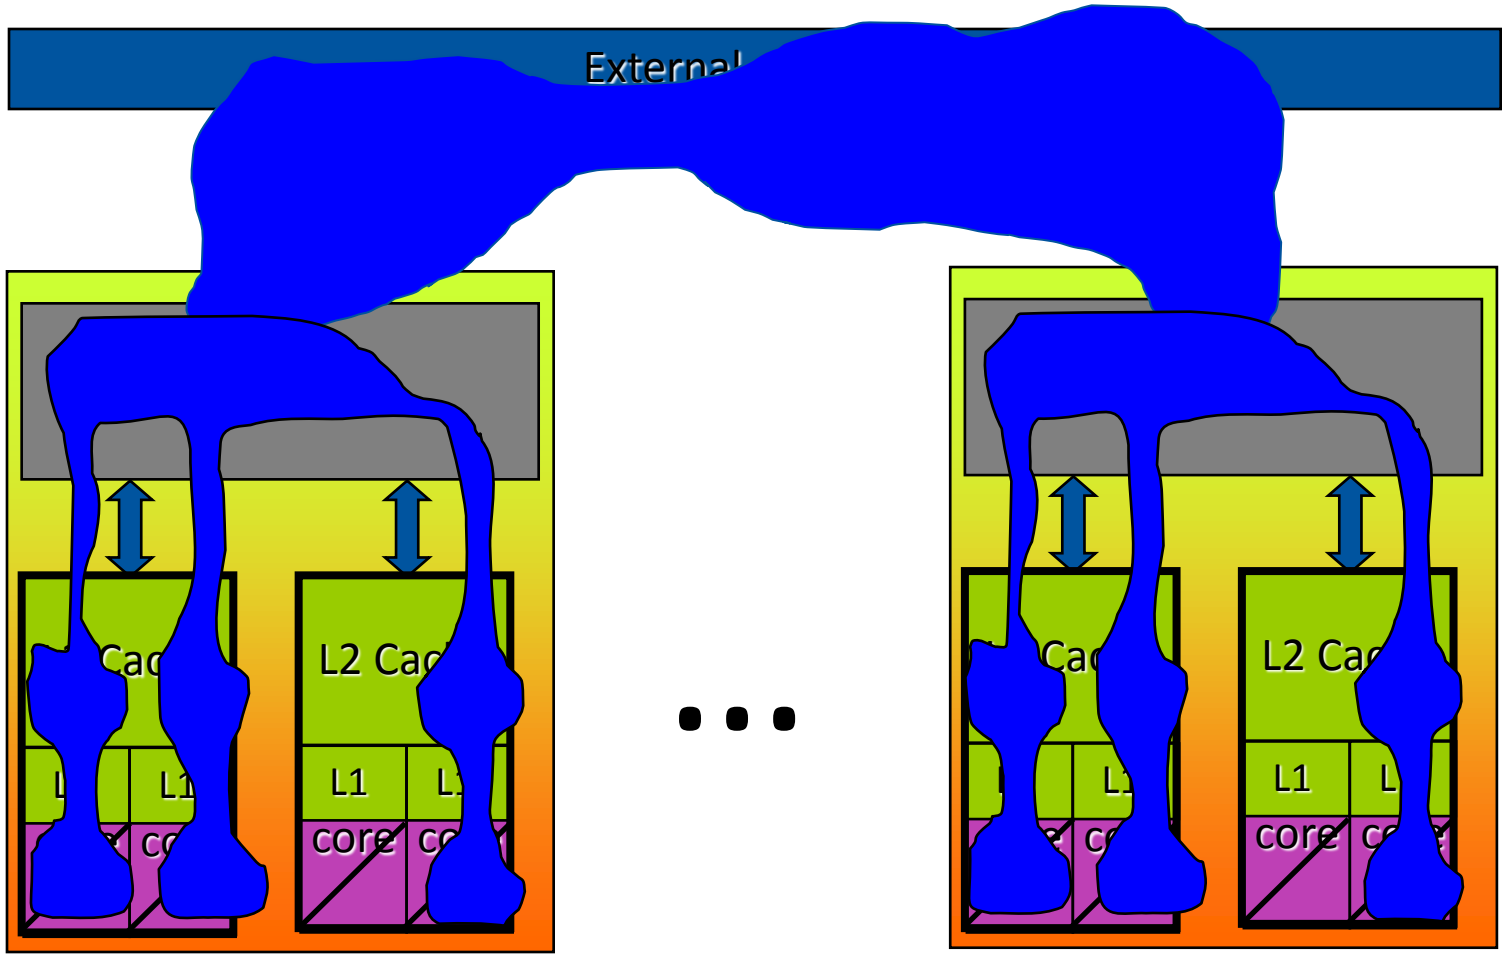

and then there are Accelerators / Coprocessors (NVIDIA GPGPUs, Intel Xeon Phi) …

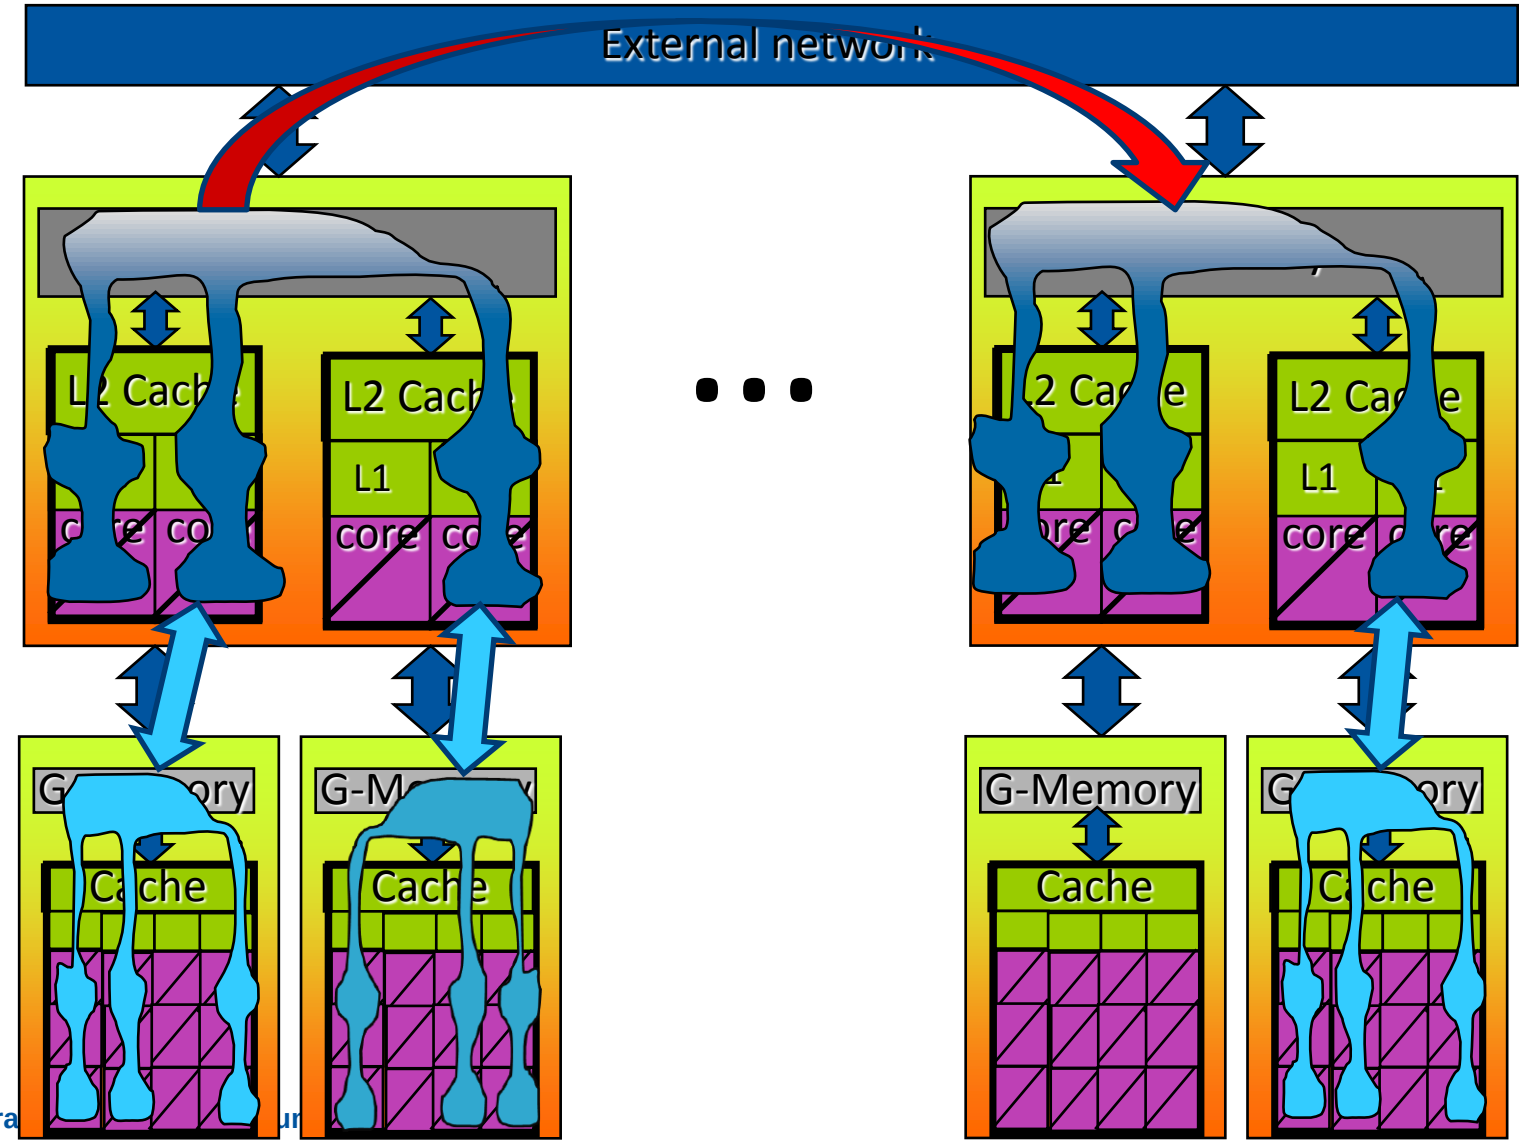

How to deal with this Complexity?

RWTH AACHEN UNIVERSITY

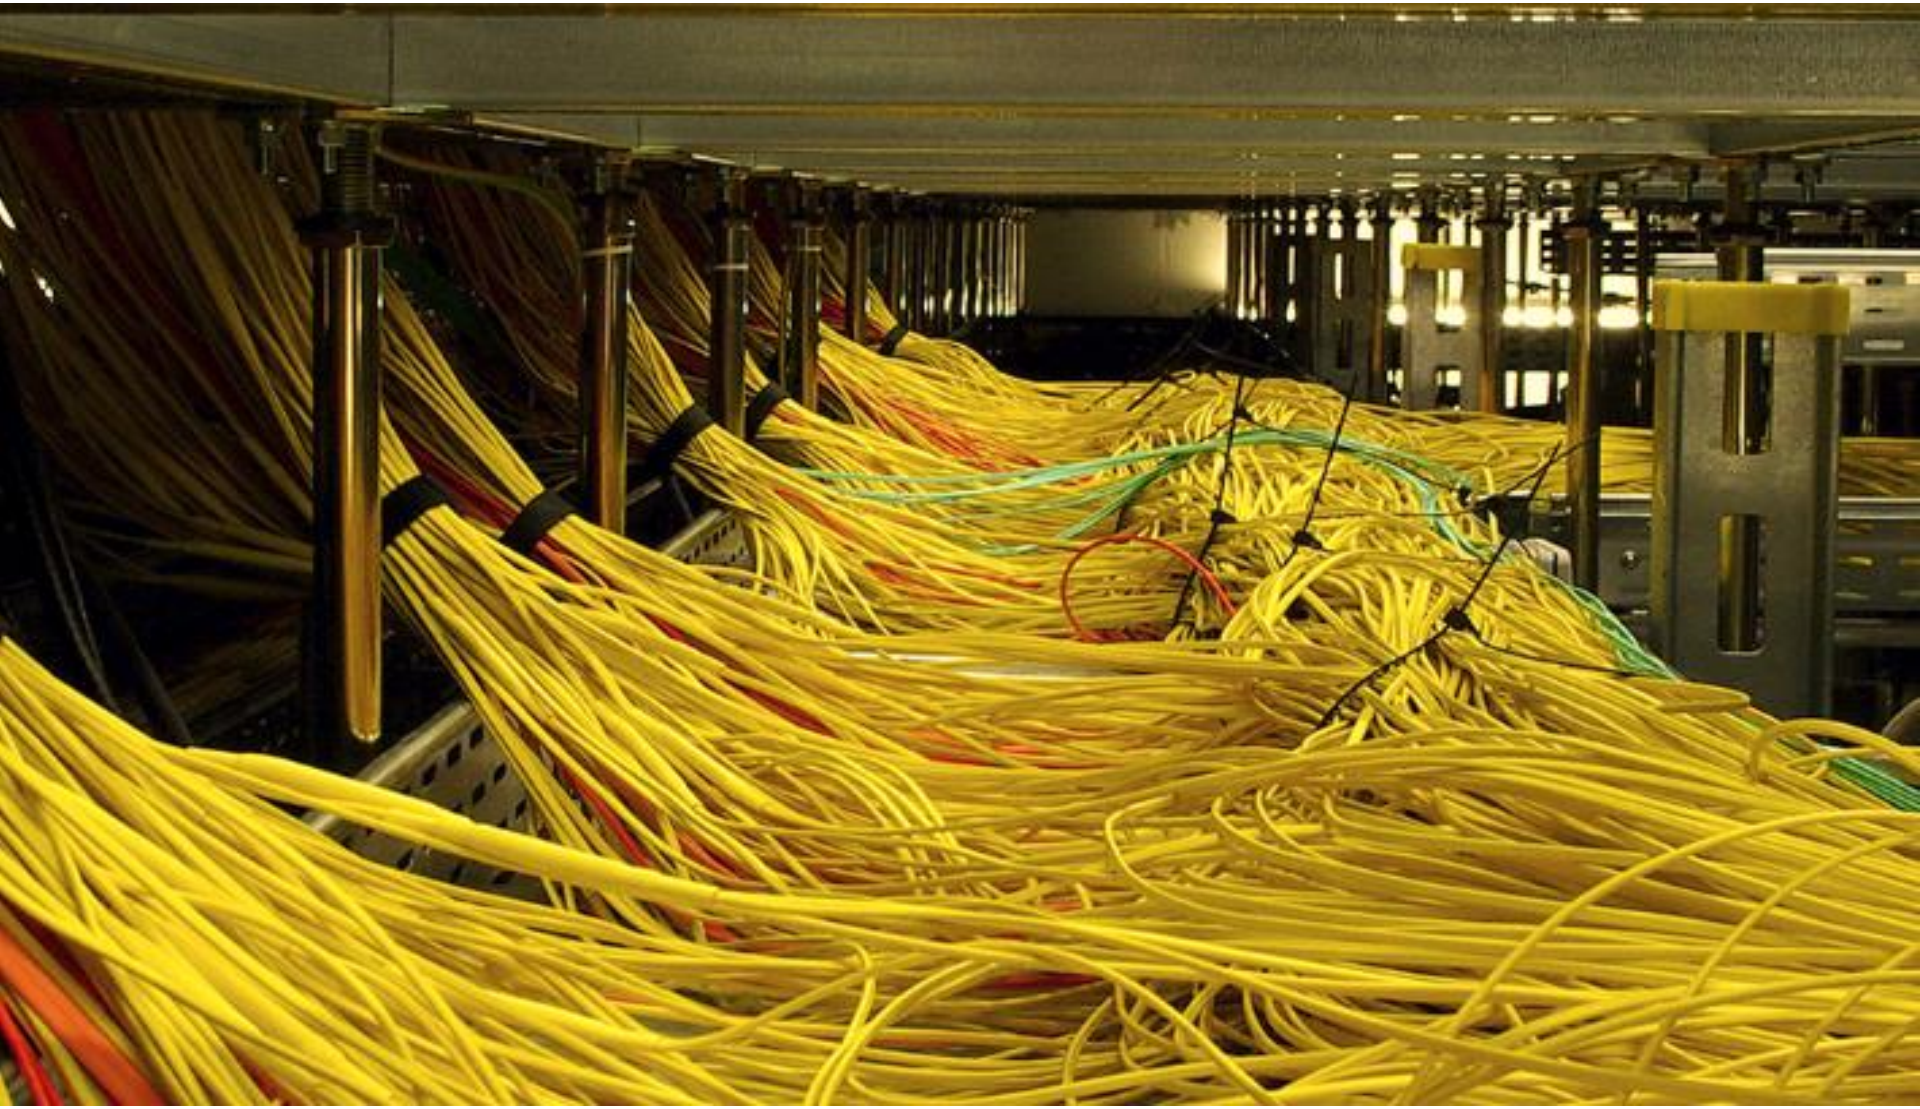

Parallel Programming Summer Course 2013 Overview

Part I **Monday, July 29,** **14:00-17:00**	**Parallel Computing Architectures &** **The RWTH Aachen Cluster Environment**
Part II **Tuesday, July 30,** **9:00-17:00**	**Introduction to OpenMP Programming** **incl Labs**
Part III **Wednesday, July 31,** **9:00-17:00**	**Advanced OpenMP Programming** **incl Labs**
Part IV **Thursday, August 1,** **9:00-17:00**	**Basic message passing with MPI** **incl Labs**
Part V **Friday, August 2,** **9:00-15:30**	**Advanced MPI, profiling and debugging of MPI applications** **incl Labs**

Parallel Programming Summer Course 2013
Agenda

RWTH AACHEN UNIVERSITY

	Monday, July 29	Tuesday, July 30	Wednesday, July 31	Thursday, August 01	Friday, August 02
09:00 - 10:30		Introduction to OpenMP Programming I (Christian Terboven)	Advanced OpenMP: Programming for Performance (Christian Terboven)	MPI: Parallel programming with message passing. Basic communication operations in MPI. (Hristo Iliev)	MPI: Hybrid programming. Often used parallel programming patterns. (Hristo Iliev)
10:30 - 11:00	break	break	break	break	break
11:00- 12:30		Lab Time	Lab Time	Lab Time	Lab Time
12:30 - 14:00	lunch break	lunch break	lunch break	lunch break	lunch break
14:00 - 15:30	Welcome (Dieter an Mey) Parallel Computing Architectures (Christian Terboven)	Introduction to OpenMP Programming II (Christian Terboven)	Advanced OpenMP: Tools (Dirk Schmidl)	MPI: Collective communication. User-defined datatypes. Virtual topologies. (Hristo Iliev)	MPI: Debugging and profiling of MPI applications. (Hristo Iliev) Lab Time
15:30 - 16:00	break	break	break	break	break
16:00 - 17:00	Parallel Computing Architectures: Outlook (Christian Terboven) The RWTH Aachen Cluster Environment (Tim Cramer)	Lab Time	Lab Time	Lab Time	

Welcome to Parallel Programming Summer 2013