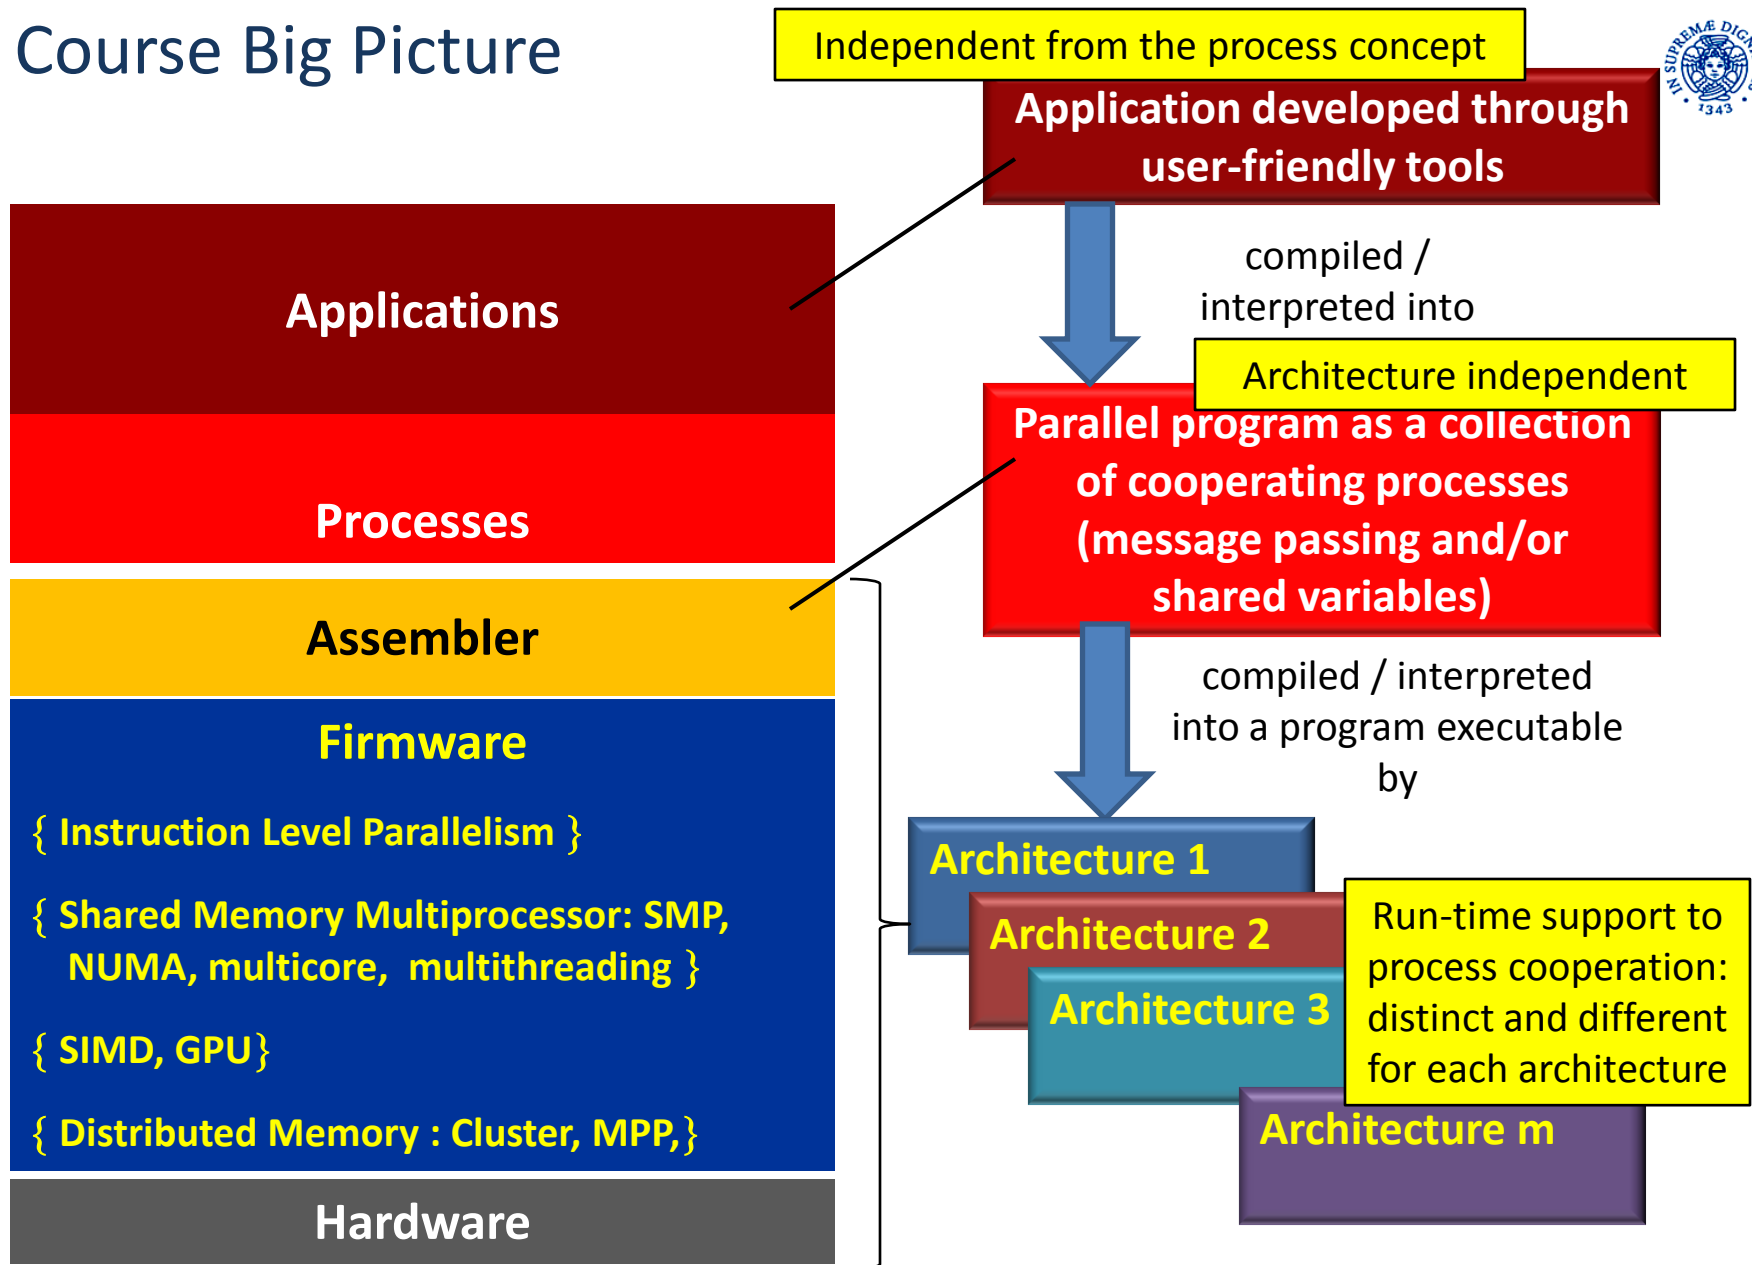Master Program (Laurea Magistrale) in Computer Science and Networking

High Performance Computing Systems and Enabling Platforms

Marco Vanneschi

Course Conclusion

Course Big Picture

Independent from the process concept

Application developed through user-friendly tools

compiled / interpreted into

Architecture independent

Parallel program as a collection of cooperating processes (message passing and/or shared variables)

compiled / interpreted into a program executable by

Applications

Processes

Assembler

Firmware

{ **Instruction Level Parallelism** }

{ **Shared Memory Multiprocessor: SMP, NUMA, multicore, multithreading** }

{ **SIMD, GPU**}

{ **Distributed Memory : Cluster, MPP,**}

Hardware

Architecture 1

Architecture 2

Architecture 3

Architecture m

Run-time support to process cooperation: distinct and different for each architecture

Lectures given *vs* Course Syllabus

- Main topics
  - ILP basics, cost model and techniques for programming tools
  - Shared memory architectures, interconnection structure, caching, cost model
  - Interprocess communication run-time support
- Many hours dedicated to prerequisites (over 33%)
- Some topics could have been developed more deeply
  - Implementation of parallel programs on parallel architectures, Compilation and configuration tools
  - ILP: superscalar, implementation of multithreading
  - On chip interconnection networks
  - General-purpose *vs* specialized architectures, networks processors, GPU, Heterogeneous architectures
- Second edition (a.a. 2010-11)
  - SPA at first semester of first year (SPM at second semester)
  - Additional hours for prerequisites ("pre-courses" in PA, ALE and SPA)
  - Integration of parallelism fundamentals (first coat of paint) and architectures
- Feedbacks are welcome

Advanced topics and deepening

- Several topics are available for Master Thesis in the Parallel Architecture Group

  - Project on Autonomic Computing, Pervasive / Mobile Grids: programming model, prototypes, emergency management applications

  - Project on high-performance applications for signal processing on multicore

  - Cost models and programming models for shared memory multicore and GPUs

  - On chip interconnection networks - optical networks

  - Network processors: applications and programming tools

  - …

- Second year Study Plan

  - Some recommended courses for students interested in these subjects

  - (Study Plans will be individually discussed on next July and September)

  - Possible connection with stages

Course approach

- Structured approach to computer architecture studying

- Methodology to understand existing and new architectures and trends
  - as opposite to mainly technological approach
  - "formalization" of computer architecture
  - motivations and rationals, fundamentals, models
  - "practical" experiences and deepening delegated to complementary courses (Study Plan)

- "Hardware-software" interrelations are fundamental to understand and to apply advanced architectures

- Also: acquiring an attitude to clear explanation is recommended to students (see Exercizes)

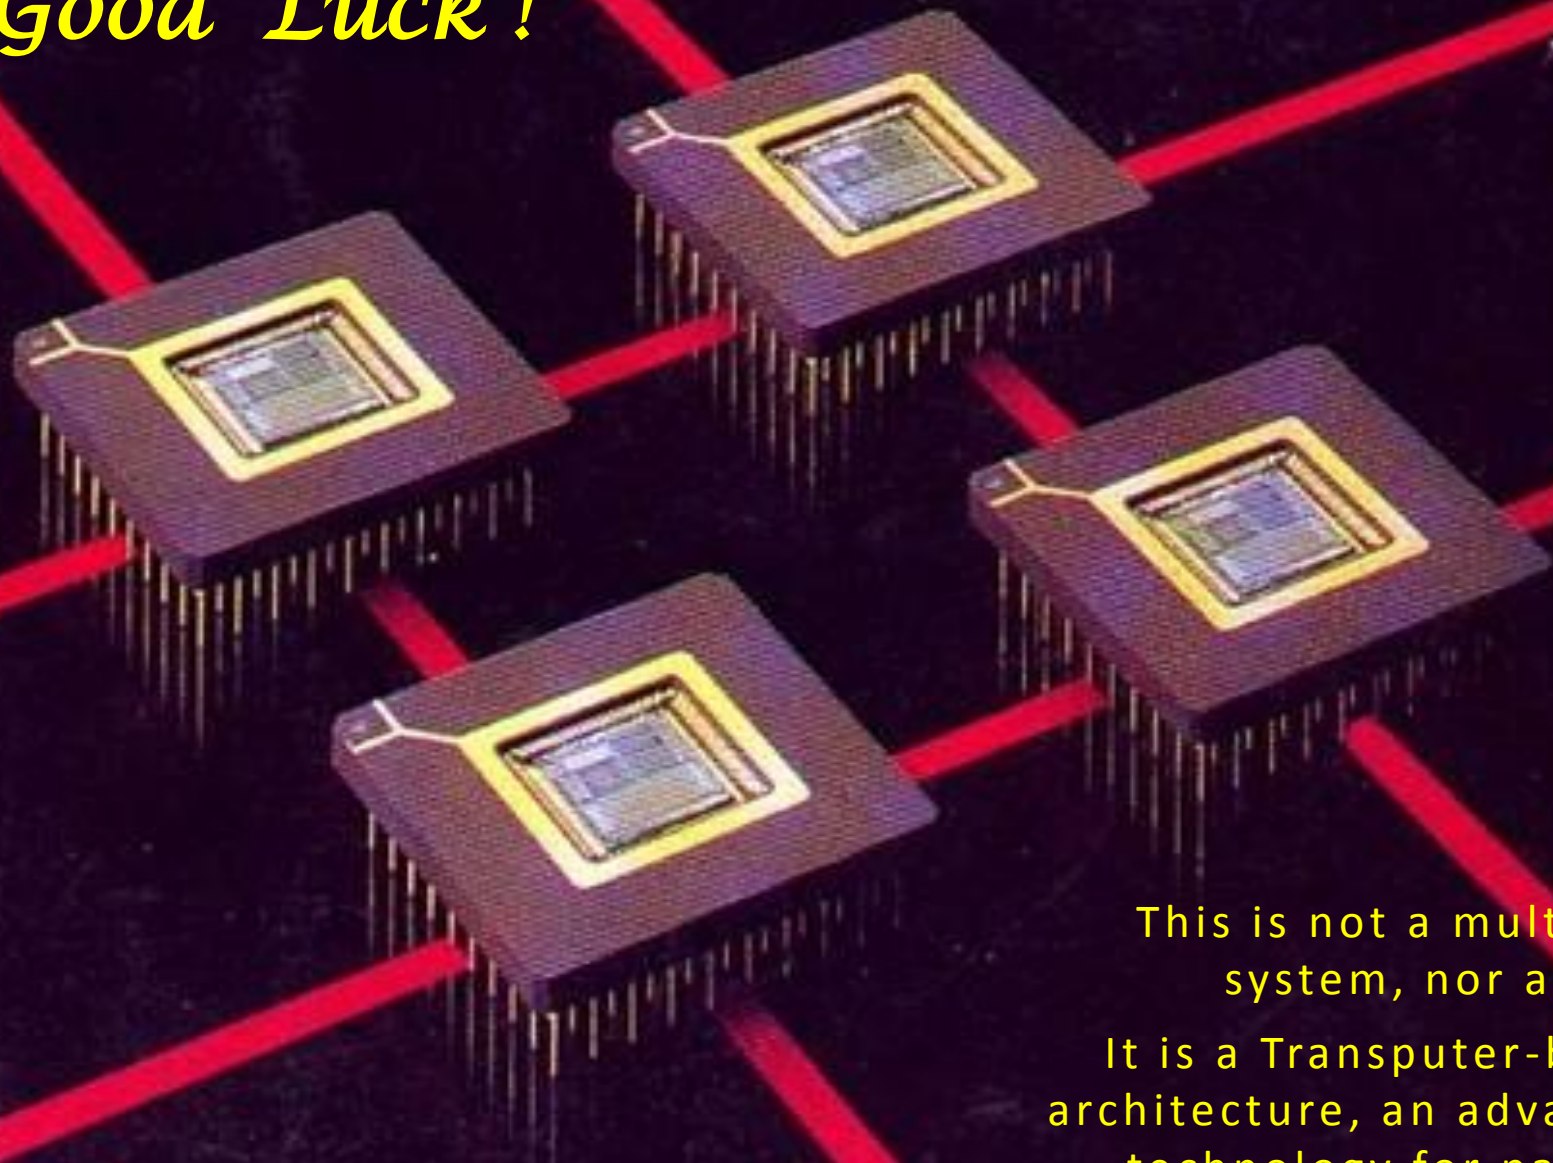- Feedbacks are welcome

Course attendance and exams

- Remarkable participation to lectures

- Few students have discussed Exercizes

- What is the current situation for students?
  – For SPA exam preparation and for all the first-year exams
  – Are advices required ?

- SPA written exam:
  – Questions on studied concepts and techniques, their interrelation
  – Reasoning capacity, synthesis capacity
  – Possibly, small exercizes
  – Fundamental: clear (and readable) explanations of the answer
  – Registration on the Official Course Page

- SPA oral exam:
  – Same approach of written exam, more oriented to deepenings and discussions

Exam planning and "appelli"

Appelli of Academic Year 2009 - 2010 (Fundamental courses)							
	January 2010	February 2010	June 2010	June 2010	July 2010	Sept. 2010	January 2011	February 2011
PA	1	2	3	4	5	6		
ALE	1	2	3	4	5	6		
SPM	1	2	3	4	5	6		
SPA			1	2	3	4	5	6
RMD			1	2	3	4	5	6
TCO			1	2	3	4	5	6

Good Luck !

This is not a multicore system, nor a GPU.

It is a Transputer-based architecture, an advanced technology for parallel processing of the 80s.