



TERAFLUX.EU



Exploiting Dataflow Parallelism in Teradevice Computing



University of Augsburg



a Proposal to Harness the Future Multicores

Roberto Giorgi – University of Siena (coordinator)
Edinburgh – HiPEAC Computing Week
05/05/2010



University of Cyprus



INRIA



University of Manchester

What is about



- 1000 Billion- or 1 TERA- device computing platforms poses new challenges:
 - (at least) programmability, complexity of design, reliability
- TERAFLUX context:
 - High performance computing and applications (not necessarily embedded)
- TERAFLUX scope:
 - Exploiting a less exploited path (DATAFLOW) at each level of abstraction
 - Applications (WP2)
 - Programming model (WP3)
 - Compilation tools (WP4)
 - Reliability (WP5)
 - Architecture (WP6)
 - Common Platform (WP7)

2

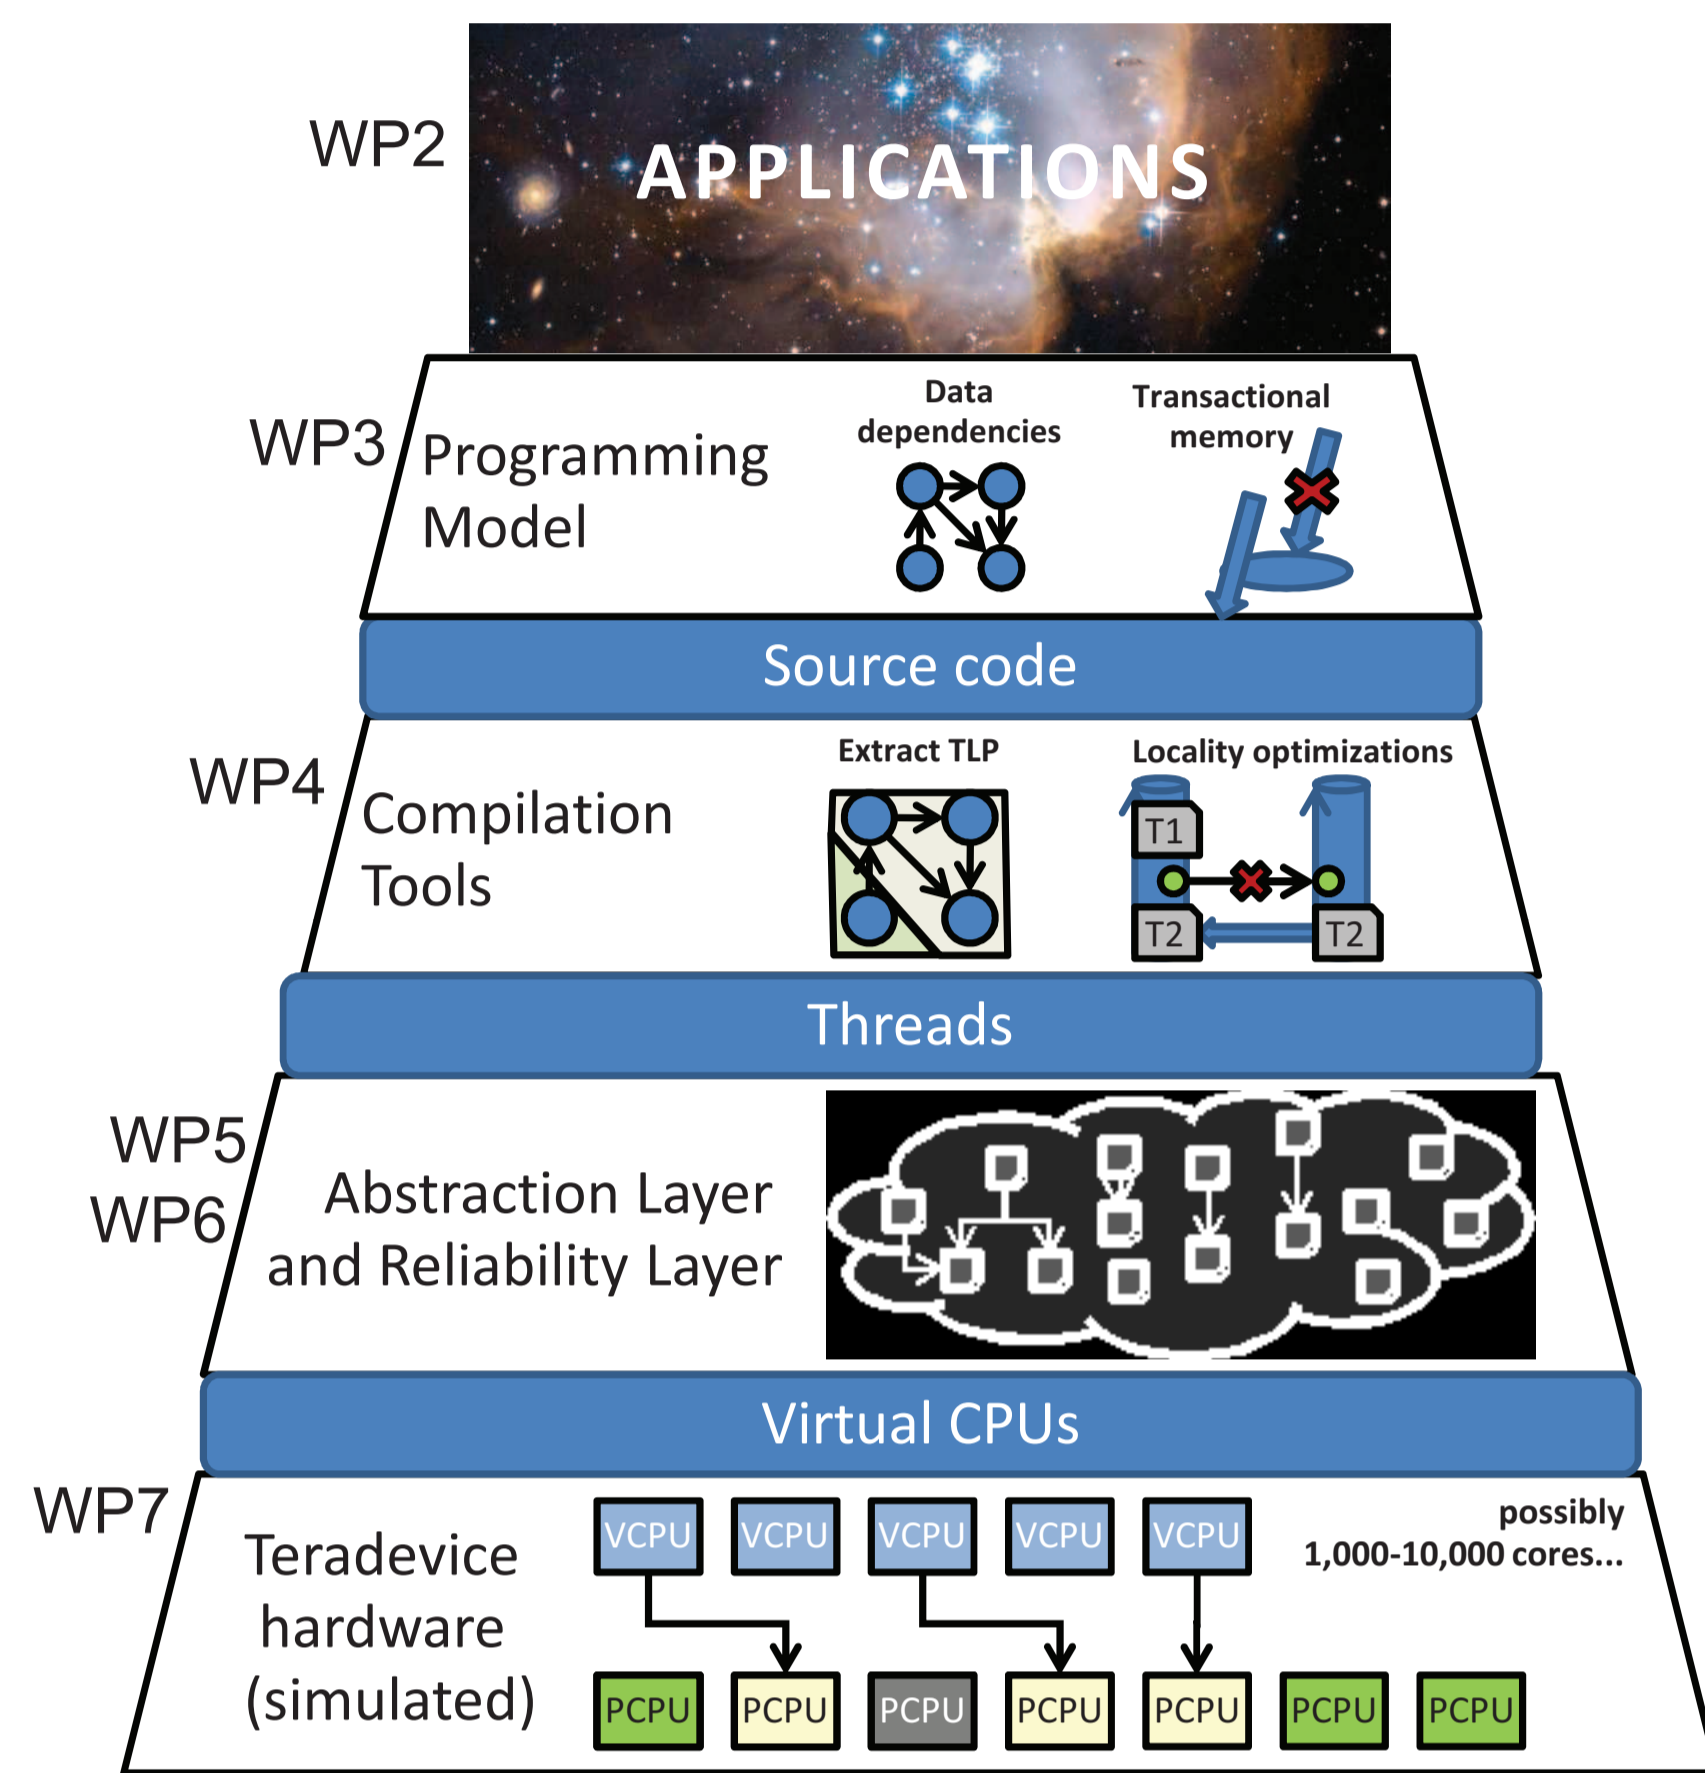


Integration Platform (WP7)

- Goals:
 - provide an OPEN-SOURCE common platform that models a teradevice system
 - Integrate workpackage results (ALL partners are contributing)
- Close to it: 1000 billion-device cores as modeled in the HP COTSon platform
 - Already established and validated (e.g., CORONA paper at ISCA-2008) *
 - Already tested for 1000 complex cores running SPLASH2 benchmarks **
 - PROPRIETARY code has to be removed/rewritten
- Workpackages will plug-in their research results into the platform and test/validate them
 - The new concepts will expand the platform toward a more programmable, less complex, more reliable teradevice system

* Vantrease, D. Schreiber, R. Monchiero, M. McLaren, M. Jouppi, N.P. Fiorentino, M. Davis, A. Binkert, N. Beausoleil, R.G. Ahn, J.H. "Corona: System Implications of Emerging Nanophotonic Technology", ISCA 2008, pp. 153-164, June 2008
 ** Matteo Monchiero, Jung Ho Ahn, Ayose Falcon, Daniel Ortega, and Paolo Faraboschi, "How to simulate 1,000 cores", HPLABS, Tech Report: HPL-2008-190

4

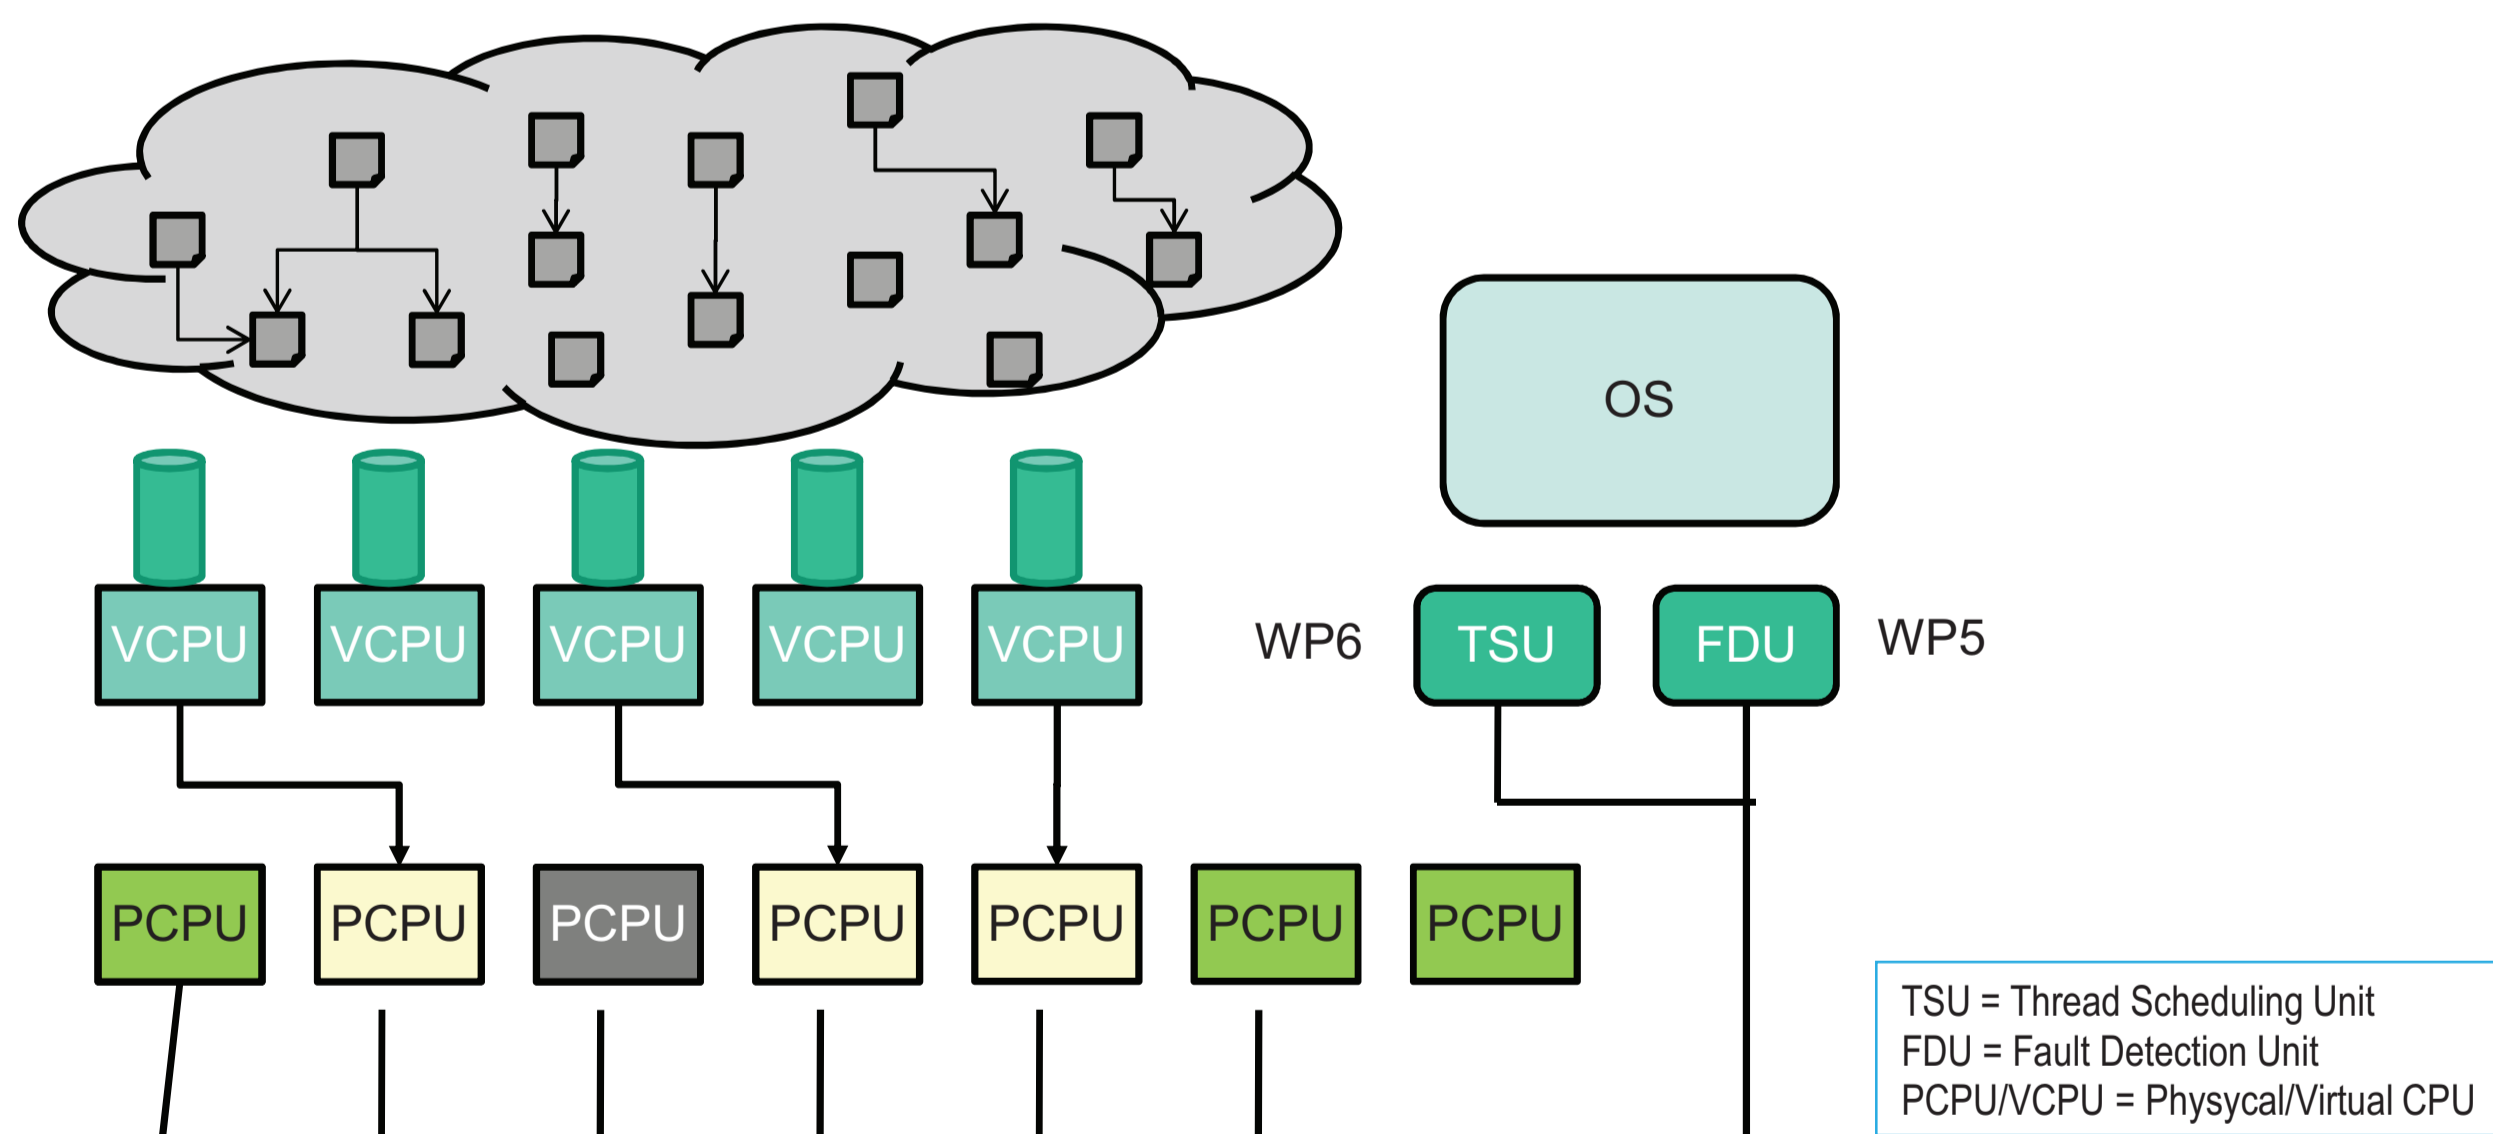


3



Reliability (WP5) and Architecture (WP6)

- Goals:
 - Creating a substrate of fine grain threads that will "flood" the architecture (WP6) with less pressure, e.g., to the memory subsystem
 - lowering the number of faults by 90% compared to the same overall multi-/many-core processor without reliability techniques (WP5)



6




Programmability (WP3) and Compilation Tools (WP4)

- Goals:
 - Achieving a more scalable implementation of given applications
 - More efficient code
- Integration of applications written in several programming styles
 - Productivity programmer (90% of cases)
 - Efficiency programmers
 - Ad-hoc dataflow programming
- Integration of dataflow and the Transactional Memory
 - Providing a more efficient implementation
- Advanced dataflow optimizations in the code
 - Aggressively tackling the CDG of an application



Applications (WP2)

- Goals:
 - Choose and characterize representative applications
 - Port applications to the new programming models
- E.G. NAMD models molecular dynamics and interactions
 - Used in biomedical research and pharmaceutical industry (drug design)
 - 3D space is modeled as grid of cells
- TERAFLUX can avoid highly inefficient Message Passing (tested on BSC supercomputer) or Shared Memory (tested on a 128-node Altix)
- Scalability is an important issue on previous machines: dataflow scheduling (e.g. DMA assisted data communication and hardware thread scheduling) will manage more parallelism, more performance, more complex biostructures
- TERAFLUX will explicit inter-cell dependencies to exploit thread-level dataflow parallelism, schedule computation and communication, thus offering higher performance



FET joint activities

- Joint Workshop
 - A yearly (December) TERA-COMP focused workshop
 - Increase the potential of interaction (e.g. HiPEAC, ARTIST, ...) communities (High Performance, Compilation, Embedded, ...)
- Joint strategy and roadmapping
 - Integration with IAB expertise
- Joint strategy and roadmapping
 - Seeking for more international cooperation, e.g., with NSF
 - Seeking for interaction with HiPEAC companies and members

8

THANKS FOR YOUR ATTENTION!