

TERAFLUX

Exploiting dataflow parallelism in Teradevice Computing

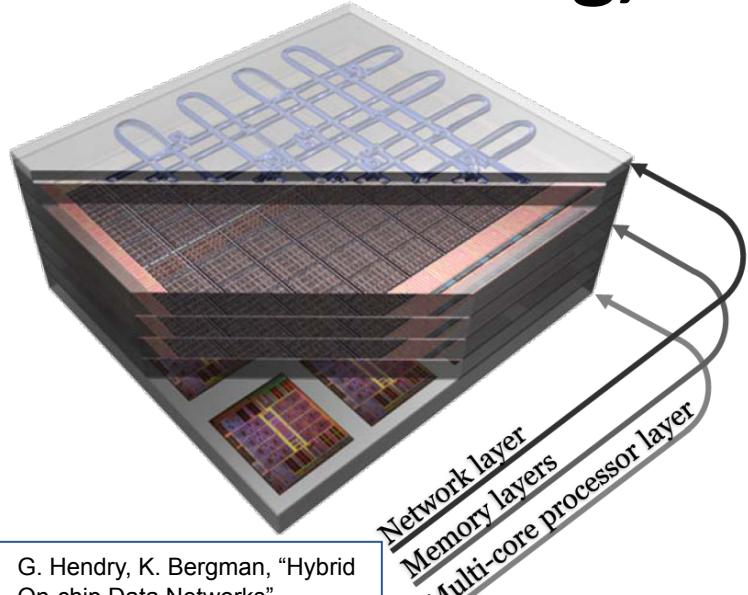
YEAR 2

What is TERAFLUX about

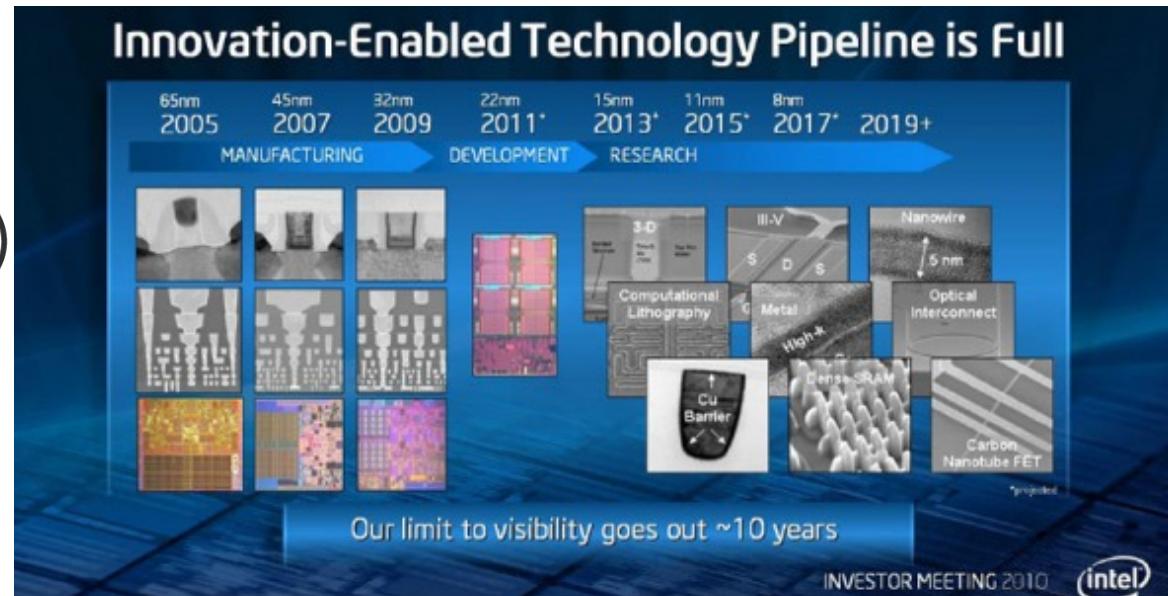
Architecture+Programmability+Reliability
of
Future (single chip)
Many-cores
(targeting 1000+ cores)

Future Scenarios

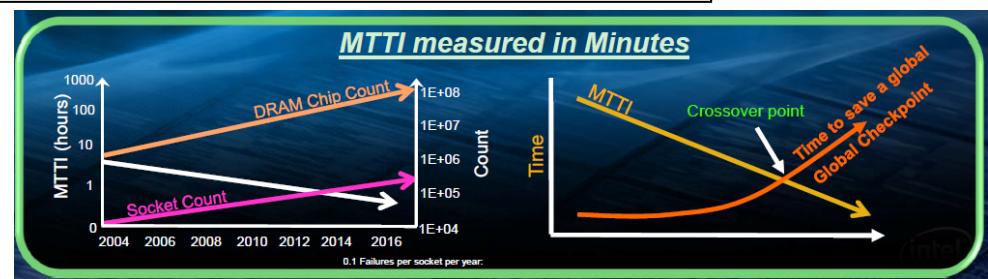
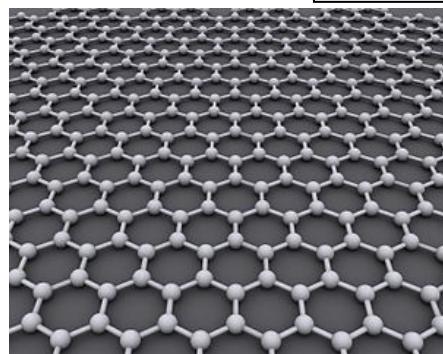
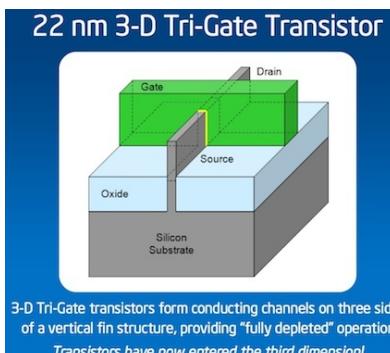
== 3D stacking, 8nm, 3D transistors, Graphene



G. Hendry, K. Bergman, "Hybrid On-chip Data Networks", HotChips-22, Stanford, CA – Aug. 2010



Fab D1X (OR), 42 (AZ) starting the 14nm node in 2013



Pawloski, May 2011, Exascale Seminar, Ghent

Session Agenda

- *TERAFLUX: exploiting dataflow parallelism in teradevice computing - Year 2*, Roberto Giorgi (UNISI)
- *Teraflux Architecture*, Skevos Evripidou (UCY)
- *Reliability aspects in Teraflux*, Theo Ungerer (U. Augsburg)
- *Teraflux, from the programming model to the execution model* Antoniu Pop (INRIA)
- *Panel of questions about Teraflux Project*

DATAFLOW

A Scheme of Computation in which
an activity is initiated by presence of
the data it needs to perform its
function
(Jack Dennis)

Recent Projects/Efforts towards DATAFLOW

- Maxeler (UK) selling “dataflow computer” to J.P. Morgan → about 350x speedup vs. standard x86 cores
- DARPA funding 25M\$ for UPHC program, encompassing:
 - Gao’s dataflow execution model (codelet based) – SWARM by ETI
 - Intel’s Runnamede project

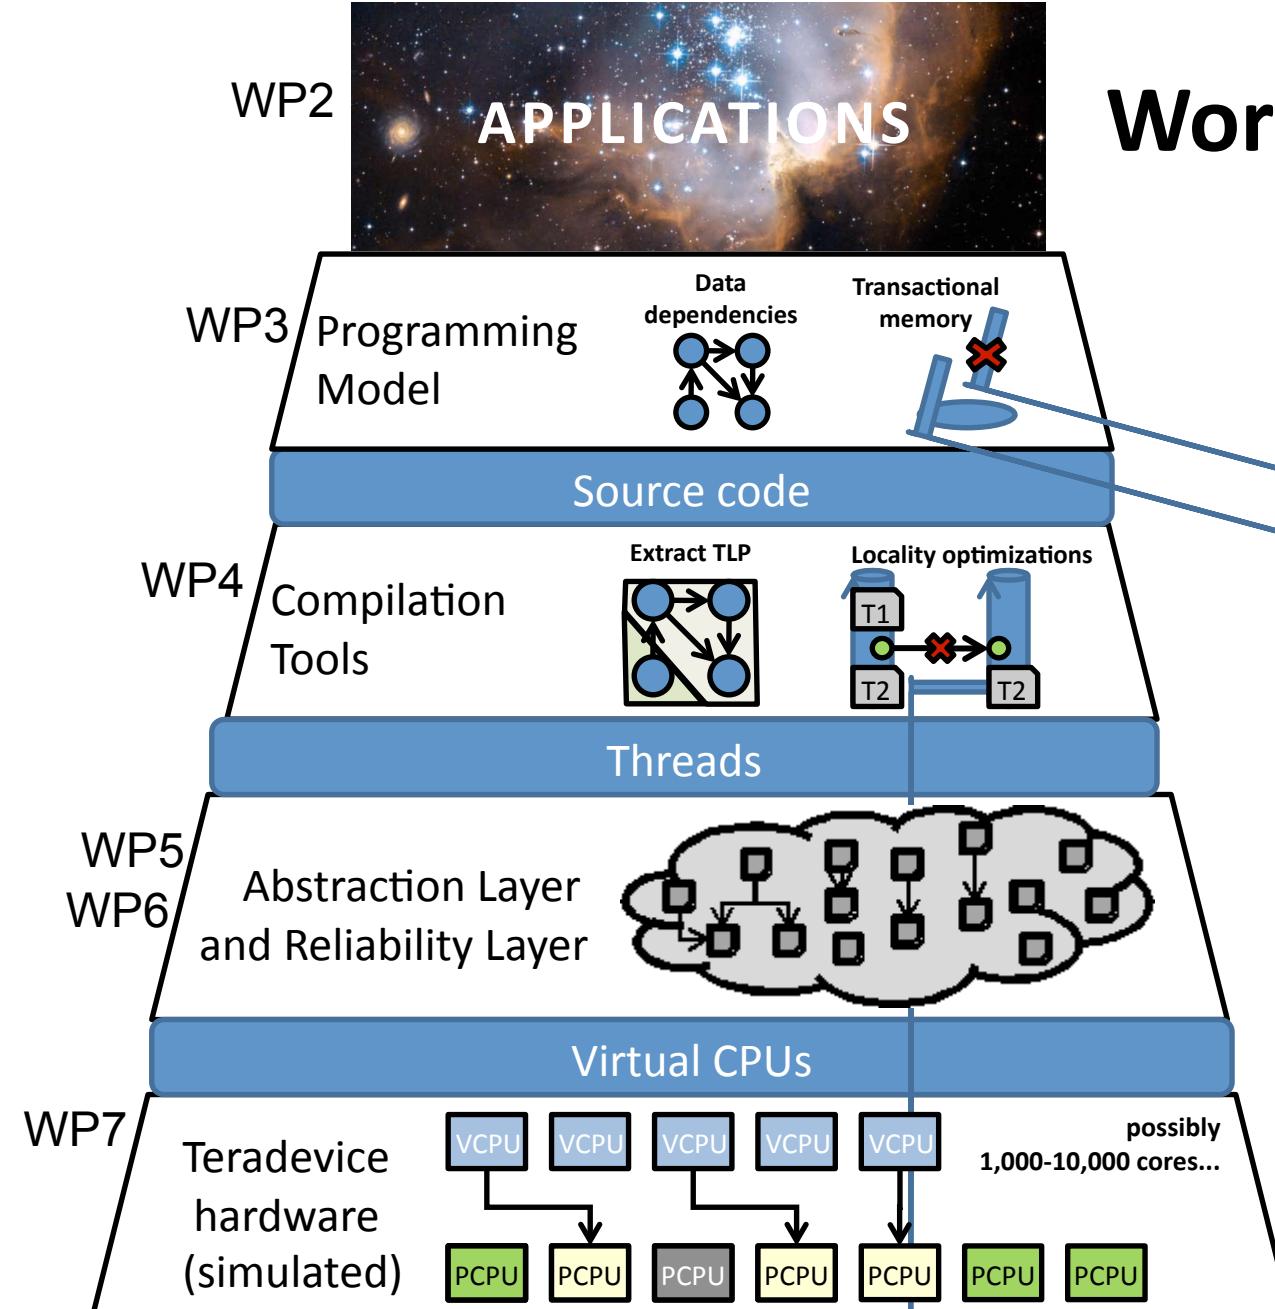


J.P. Morgan Deploys Maxeler Dataflow Supercomputer for Fixed Income Trading
December 15, 2011

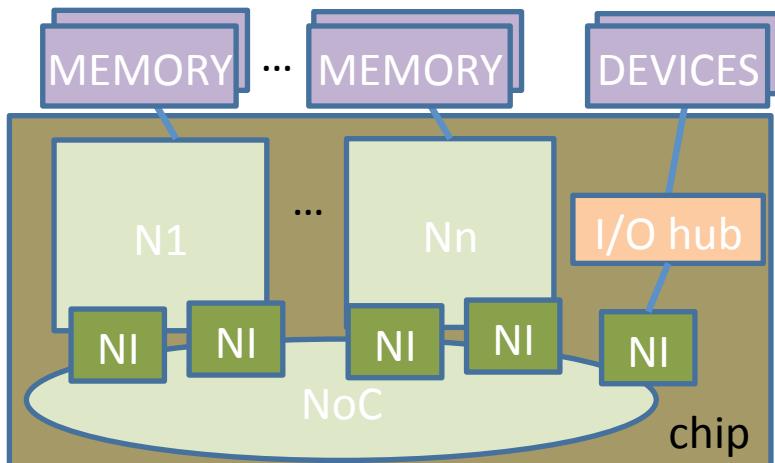
UPHC=Ubiquitus
High-Performance Computing

The Intel-led UHPC team intends to develop new circuit topologies, new chip and system architectures, and new programming techniques to reduce the amount of energy required per computation by between 100x and 1000x compared to today's computing systems. Such dramatic reduction in energy consumption will allow these future systems to take full advantage of the increasing transistor budgets afforded by the steady advances in Moore's Law.

Working Hypothesis



- 1000 Billion- or 1 TERA-device computing platforms pose 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



TERAFLUX Architectural template

LEGENDA:

n = # of nodes

m = # of cores per node

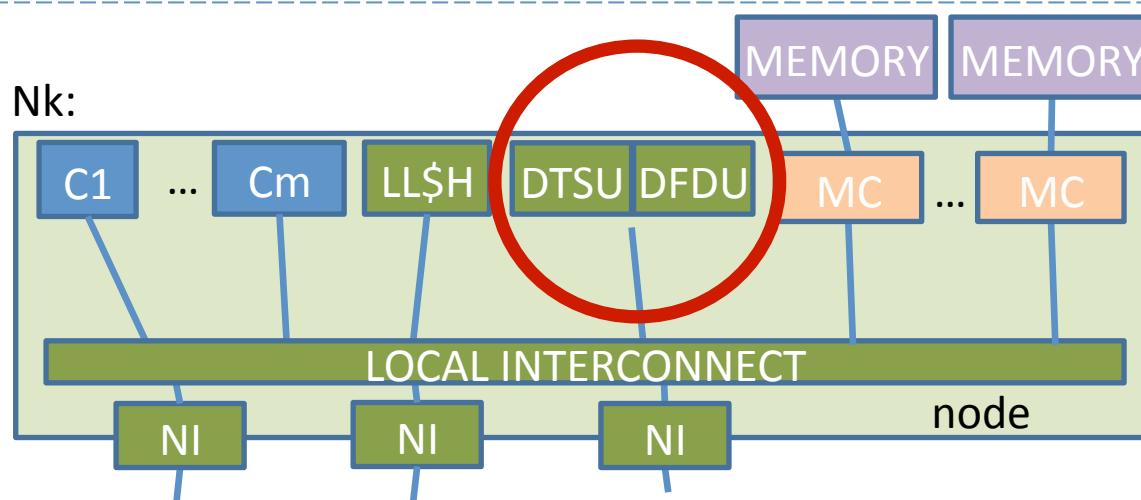
u = # of DRAM controllers insisting on the Unified Physical Address Space

z = # of I/O Hubs

N_k = k-th Node ($k=1..n$)

NI = Network Interface

NoC = Network on Chip



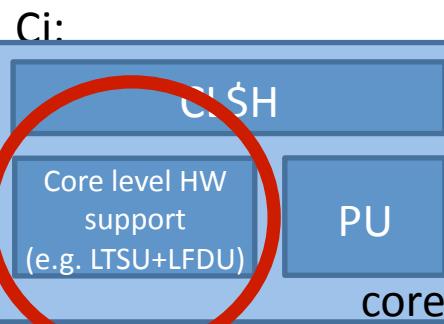
C_j = j-th core ($j=1..m$)

MC = Memory Controller

DTSU = Distributed Thread-Scheduler Unit

DFDU = Distributed Fault-Detection Unit

LL\$H = Last Level Cache Hierarchy

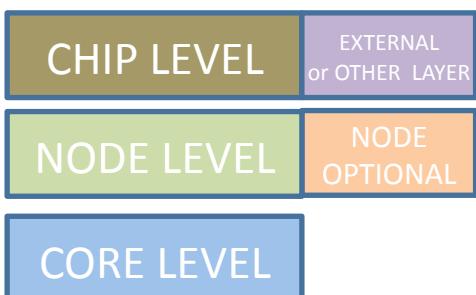


CL\$H = Core Level Cache Hierarchy

PU = Processing Unit

LTSU = Local Thread-Scheduler Unit

LFDU = Local Fault-Detection Unit

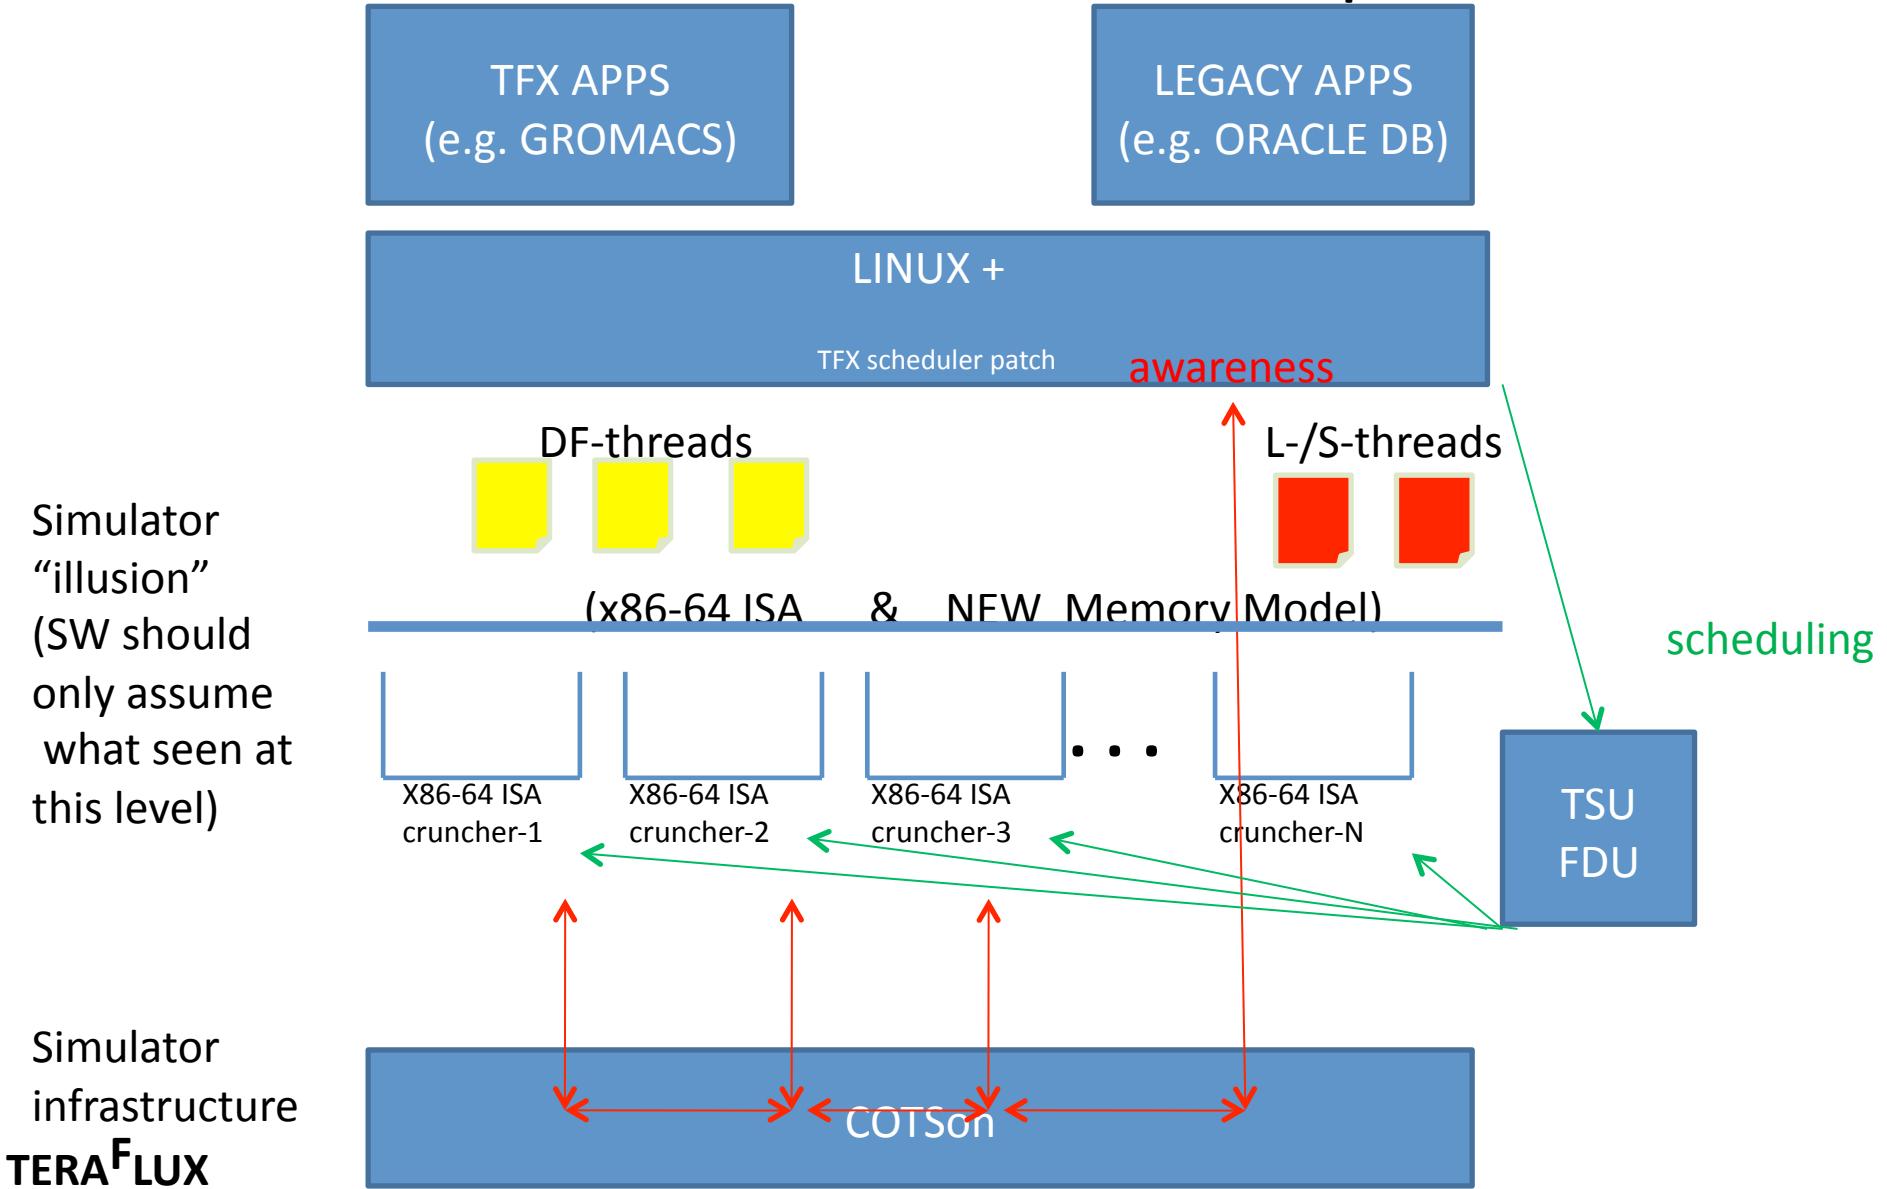


TERAFLUX

Our pillars

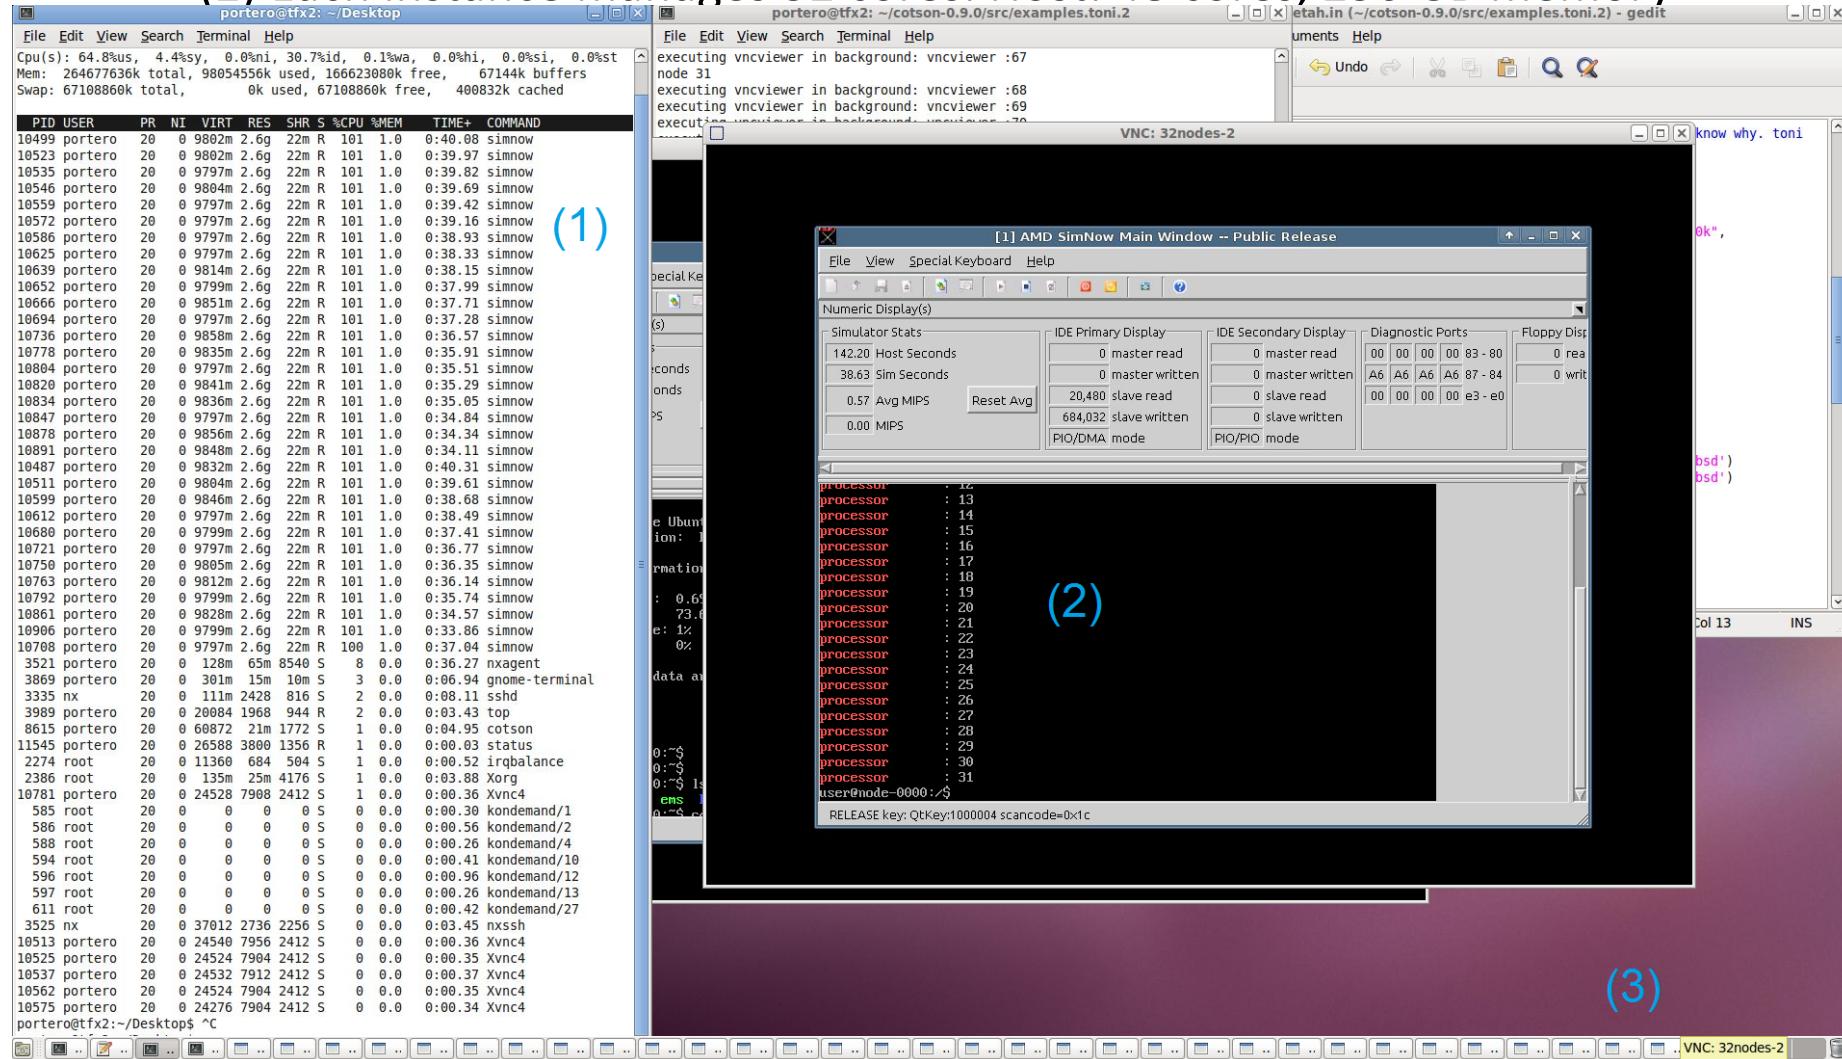
- FIXED and MOST-USED ISA (**x86**)
- MANYCORE FULL SYSTEM SIMULATOR (**COTSon**)
- REAL WORLD APPLICATIONS (e.g. GROMACS)
- SYNCHRONIZATION: **TRANSACTIONAL MEMORY**
- **GCC** based TOOL-CHAIN
- OFF-THE-SHELF COMPONENTS FOR CORES, OS,
NOC, MEMORY HIERARCHY
- **FDU** AND **TSU** (Fault Detection Unit and Thread
Scheduling Unit)

Evaluating a MANY-CORE chip of the future (2020), i.e., 1000+ cores on a chip



Simulation booting up 1024 cores. (1) COTSon execution of 32 SimNow instances.

(2) Each instance manages 32 cores. Host: 48 cores, 256 GB memory



Note: the simulation is PARALLEL at GUEST NODE-LEVEL and it's also possible to distribute the simulation on several HOST NODES

Major Technical Innovations in TERAFLUX

- Fragmenting the Applications in **Finer grained DF-threads**:
 - DF-threads allow an easy way to decouple memory accesses, therefore hiding memory latencies, balancing the load, managing fault, temperature information without fine grain intervention of the software.
- **Possibility to repeat the execution** of a DF-thread in case this thread happened to be on a core later discovered as faulty
- Taking advantage of a “**direct**” **dataflow communication of the data** (through what we call DF-frames).
- **Synchronizing threads** while taking advantage of native dataflow mechanism (e.g. several threads can be synchronized at a barrier)
 - DF-threads allow (atomic) Transactional semantics (DF meets TM)
- A **Thread Scheduling Unit** would allow fast thread switching and scheduling, besides the OS scheduler; scalable and distributed
- A **Fault Detection Unit** works in conjunction with TSU

TERAFLUX SIMULATOR (COTSon)

<http://cotson.sf.net>

HP-Labs COTSon is **OPEN-SOURCE**



FUTURE AND
EMERGING
TECHNOLOGIES
PROJECT N. 249013



SEVENTH FRAMEWORK
PROGRAMME THEME
FET proactive 1 (ICT-2009.8.1)
Concurrent Tera-Device Computing



TERAFLUX

Exploiting dataflow parallelism in Teradevice Computing

PROJECT NUMBER: 249013

<http://terafux.eu>

