PROJECT RESULT

EDA for SOC Design and DFM

MEDEA+

2A718: Tera-scale multi-core processor architecture (TSAR)

Towards powerful future parallel computing

Parallel computing techniques based on multi-core processing are gaining ground, but have run into the sand at the threshold between multi-core and many-core architectures, where hundreds of processors can be incorporated into a single chip. The markets are crying out for such computing power, especially for applications like video streaming. But industry lacks the knowledge and tools to design, program and run such complex architectures. The TSAR project, tasked with bridging this knowledge gap, has come up with scalable many-core architectures that can include up to more than 10^3 processors.

European industry is striving, in a competitive global market, to improve data-processing power, streaming applications, wireless and mobile access, remote medical diagnosis, artificial intelligence and more. Manufacturers have given up extracting ever greater performance from a single processor core; industry is moving to multi-core, and eventually many-core, architectures.

Parallel computing based on multi-core processing, or assigning multiple processors to speed up processing of very large data files, is not new. Technology advances have already pushed the capability frontiers in multi-core processing to between 10 and 100 processors within a chip. Now the latest leading-edge research is focusing on many-core processing with between 100 and 1,000 processors on a single chip.

Bridging knowledge gap

Yet this many-core era is still only emerging. There remains a lack of knowledge about how to build, program or manage systems of 64 to 1024 processors, and the computer-architecture community lacks the infrastructure tools to carry out this research. The MEDEA+ 2A718 TSAR project set out to bridge this knowledge gap.

The consortium – including a chipmaker, a major computer manufacturer, systems developers and both research centres and academics from across Europe – was tasked with exploring and developing massively parallel processing architectures based on much more sophisticated many-core processors.

Key targets were to:

- Investigate massively parallel processing to support both shared-memory systems with cache coherence and streaming with message-passing models;
- Prove a scalable multi-processing concept in hardware/software, processor architecture and protocols, design, validation and implementation; and
- Produce product and system prototypes for demonstration as real industrial applications.

The project targeted the principal technological problems involved in many-core processing, especially for architectures with tens or hundreds of processor cores. As classical shared buses do not have the ability to scale, the project built on a distributed, packet-switched network-on-chip (NoC) approach.

Co-operative multi-processing

TSAR focused on co-operative multi-processing. No master processor is required; all processors are at the same level, process the same data and use the same memory space. The challenge was to develop the protocol to make this possible.

The work required in-network management of cache coherence and memory consistency. Design of protocols and NoC architectures, development of appropriate software – including operating systems (OSs), parallel-system programming and compilation – and virtual and field programmable gate array (FPGA) based prototyping were all key themes.

By the end, the project had developed scalable many-core architectures which could include up to 1,000 processors. The architectures could support a wide range of memory models – including message passing and shared memory – and allowed trade-offs between performance, parallel-programming efficiency and power consumption. The first TSAR virtual prototype is a cycle-accurate generic SystemC model ranging from 4 processors in 1 cluster to 512 processors in 128 clusters.

Advances included:

- An NoC-based architecture supporting a write-through cache-coherence strategy and able to use a range of commodity OSs. An advanced in-network cache-coherence protocol was proven on the NoC-based architecture;
- A new OS concept, supporting the non-uniform memory access characteristics of such many-cores architectures;
- Efficient and fully-automated placement of applications on the 2D mesh by the advanced locality management operating system (ALMOS) OS;
- Compilation and execution of key applications, including SPLASH benchmarks on the multi-processor virtual prototype; and
- Demonstration of a hardware FPGA prototype containing one complete TSAR cluster – four processor cores.

Philips used the TSAR concept to develop a medical-image filtering software application. This sequential C program is a multi-threaded application where the number of threads can be adjusted to fit the number of processors. The scalability of the distributed hybrid cache-coherence protocol implemented in the TSAR many-cores architecture was demonstrated, providing a quasi-linear speedup, up to 512 processor cores.

Many-core architectures

Applications which can use this level of processing power already exist. Advanced nuclear-reactor thermal-hydraulic modelling for example, requires 32 million processor hours, climate-change models can absorb five million processor hours and high-resolution modelling of the Earth system can easily use 20 million processor hours.

Co-operative processing based on many-core processors is the future; the success of TSAR marks the presence of Europe in this crucial and highly competitive arena. Several partners – such as ACE, Bull, Compaan, NXP, Philips and UPMC/Lip6 – have already incorporated parts of the TSAR many-core technology into the design of architectures and protocols for new hardware and software applications.

TSAR has also enabled partners such as Philips and Thales to parallelise existing key applications to take advantage of faster many-core processor architectures. However further investment will be needed to change the way key applications are managed and processed to take full advantage of this new approach.

The results of TSAR are also feeding into a new CATRENE project, SHARP. This aims to develop a heterogeneous high-performance computing platform to accelerate specific applications. One of these will be the TSAR architecture.

EDA for SOC Design and DFM

2A718: Tera-scale multi-core processor architecture (TSAR)

PARTNERS:

ACE
Bull
CEA-LETI
Compaan Design
FZI
NXP Semiconductors
Philips Healthcare
Thales Communications
TU Delft
Uni Leiden/LIACS
UPMC/LIP6

PROJECT LEADER:

Huy-Nam Nguyen
Bull

KEY PROJECT DATES:

Start: June 2008
End: May 2011

COUNTRIES INVOLVED:

France
Germany
The Netherlands

CATRENE Office
9 Avenue René Coty
F-75014 Paris
France
Tel.: +33 1 40 64 45 60
Fax: +33 1 43 21 44 71
Email: catrene@catrene.org
http://www.catrene.org

EUREKA Σ!

MEDEA+ Σ!2365 is the industry-driven pan-European programme for advanced co-operative R&D in microelectronics to ensure Europe's technological and industrial competitiveness in this sector on a worldwide basis.

MEDEA+ focuses on enabling technologies for the Information Society and aims to make Europe a leader in system innovation on silicon.