

HIPEAC ROADMAP – Interconnects

30/06/2008

Interconnects play a critical role in all computer systems, both on-chip –connecting the multiple components– and off-chip –building networked systems. The era when buses sufficed for the required bandwidth is gone – nowadays most applications need the parallelism offered by networks of switches or routers connected using multiple point-to-point links. While multi-chip networking has been studied for many years, researchers recently also turned their attention to networks-on-chip (NoC). Although the general principles are the same in all networks, constraints vary widely between the on-chip and the off-chip worlds, resulting in radical differences between these two different network types. A lot of research is still needed in NoC's, because this is a young field, while multi-chip networking also evolves and requires its own research efforts as well. Below, we identify and catalogue what we believe to be the main research challenges in interconnects, today:

Challenge 1: Network and Interfaces Microarchitecture

There is a large variety of options when implementing an interconnection network. Nevertheless, truly scalable solutions that support a wide variety of application with diverse traffic characteristics, while satisfying tight area and power budgets, call for a continuing effort in optimizing network's topology, routing mechanisms and switch microarchitecture.

- 1.1. Latency reduction techniques.
- 1.2. Energy-delay optimized switch microarchitectures and buffer reduction techniques.
- 1.3. New switching strategies involving scheduling and flow control for virtualization and QoS.
- 1.4. Power-performance optimized flow control and congestion management techniques.

Lightweight network interfaces are also of paramount importance since they provide the connection path between computational elements and the communication medium.

- 1.5. New mechanisms for efficient support of interprocessor communication and/or cache coherence protocols.
- 1.6. Support for multiple outstanding transactions and out-of-order completion.
- 1.7. Reduction of protocol conversion and packetization overhead.
- 1.8. Design for reusability and ease of integration.

Network components, besides handling communication, should be augmented with additional features that will improve overall system design and performance. How easily a system is debugged and tested are critical parameters of a successful design that catches the short time-to-market frame.

- 1.9. Appropriately designed network and protocols that offer increased system observability.
- 1.10. Hardware/software primitives for design error detection and diagnosis.
- 1.11. Hardware support for network traffic monitoring and analysis. Very useful to application-software optimization and tuning.

Challenge 2: Design tools

Increasing design productivity and improving system quality clearly depends on the available design tools that will automate the design process and let the designer focus on critical decision making. Many different approaches have appeared in the recent years and significant progress has been made in this area. Nevertheless, a lot more effort is needed before these tools become a commodity.

- 2.1. Unifying platforms for complete design-space exploration (from architectural concepts to circuit implementations).
- 2.2. Full-system simulation with details of the interconnection network and network interfaces integrated.
- 2.3. Methodologies and tools for application-driven traffic modelling and analysis.

- 2.4. Low complexity application-specific interconnection networks and the corresponding design automation tools.

Challenge 3: Power Efficiency

Power efficiency needs to be addressed at all levels of design. At the moment, interconnect power efficiency is mostly addressed at the circuit level with some few exceptions that offered solutions at higher levels of abstraction. A focused and combined approach is needed that would bring power issues at the architectural level combining both the hardware and the software perspective of interconnect.

- 3.1. New metrics need to be introduced that would appropriately mix classical network performance metrics such as latency and throughput with energy efficiency.
- 3.2. Interconnect dynamic power management (per link or per switch methods) taking into account the time-evolving workloads.
- 3.3. Distributed power management policies. How power-off or drowsy modes of operations are decided locally and communicated globally via the network.
- 3.4. Techniques that divide the total power budget between computation and communication on a per application basis.
- 3.5. How switches are interfaced to processors that follow independently a dynamic voltage/frequency scaling policy to reduce their power. How this affects flow control and buffering requirements.
- 3.6. Efficient management of thermal issues. Congestion/temperature analogy. Try to avoid hot areas or reduce incoming traffic to a region with high-temperature.

Challenge 4: Reliability, Fault Tolerance, Yield Improvement

Very-deep sub-micron technologies impose several other challenges that affect interconnection network design and overall performance. Offering acceptable levels of reliability and fault-tolerance needs additional design effort. Much work has been done in this area and many useful techniques exist. However, the increasing technology limitations call for a continuous improvement.

- 4.1. Robust interconnection network design and protocols in the presence of transient faults.
- 4.2. Efficient error detection and correction mechanisms.
- 4.3. Practical reconfiguration techniques (either at run-time or at idle periods) of routing algorithms.

In the same context, additional techniques are required to alleviate the effects of increased process variations that degrade the yield of the designs.

- 4.4. Adaptive inter-node synchronization.
- 4.5. Flow control adaptation to tackle variable switch performance throughout the chip.
- 4.6. In-network process variation monitoring.

Challenge 5: Virtualization, QoS, Security

System virtualization, where a single system is viewed as a set of virtual entities that operate independently, is a critical demand for current and future computer systems. Each of the independent “virtual” networks must be able to meet pre-specified Quality-of-Service (QoS) levels. Also, building scalable systems of very-large complexity calls for efficient implementation of protection and isolation mechanisms.

- 5.1. Energy-delay exploration of virtualization: Which approach is more effective? Physical partitioning of the network or allowing virtual logical partitions of the network to share the common physical components (switches and links).
- 5.2. New flow control and congestion management approaches that provide quality-of-service to a large number of virtualized systems with minimal hardware overhead.
- 5.3. Network interface design that handles protection.
- 5.4. How can topologies and routing algorithms help in system partitioning and isolation.

Secure computation and communication is a critical aspect of modern high-performance and embedded systems. The goal is to achieve the desired levels of security without compromising performance, size, cost, or

energy consumption. This seems feasible when security is considered at the beginning of system design and not later as an afterthought. Interconnect design should provision all necessary features so that security-preserving actions can be efficiently implemented.

5.5. Network-level threat detection.

5.6. Secure end-to-end data transfers with the synergy of new hardware primitives and software approaches.

5.7. Efficient system notification broadcast mechanisms.

Challenge 6: Off and On-Chip Interconnection Networks Integration

As more and more processors are present on the same chip they increase the need for off-chip bandwidth in order to interact with large capacity off-chip memories. This is mostly true for multimedia or graphics workloads where on-chip memory hierarchy does not help much in reducing the demand for off-chip accesses. The on-chip interconnection network will be the vehicle for delivering the large amount of data required by the on-chip processors.

6.1. Efficient integration of memory controllers (or network interfaces) with the on-chip network to carry out several simultaneous requests to the off-chip networks.

6.2. New congestion management and flow control techniques that will focus on efficient inter-network data transfers.

6.3. Simplified network interfaces that reduce the overhead of protocol translation and packetization between off and on-chip networks.

6.4. Maybe, definition of a common standard interface.

Challenge 7: Benchmarking

In order to clearly determine the benefits and the shortcomings of future interconnects architectures we need the definition and the development of a new benchmarking methodology that will cover all aspects of network design.

7.1. Intensive parallelization effort of the most computation demanding applications used in the industry (e.g., signal processing, multimedia).

7.2. Selection of a suitable programming paradigm for on-chip networks. Today is almost impossible to find common operating conditions and a suitable application suite that can be used to compare NoC architectures.

7.3. Realistic performance/power models of network architecture and components.

ROADMAP – Questions & Answers

The following questions are related to the domain of your cluster:

Based on the categorization shown in the first part of the preliminary Interconnects Roadmap, we split the research challenges according to the timeframe we believe they will emerge as critical research questions as follows:

What are the main problems/challenges you see in the next 5 years?

Network and Interfaces microarchitecture: 1.1, 1.2, 1.6, 1.11

Design tools: 2.1, 2.3, 2.4

Power efficiency: 3.1, 3.2

Reliability, Fault tolerance, yield improvement: 4.3

Virtualization, QoS, Security: -

Off and on-chip interconnection networks integration: -

Benchmarking: 7.1, 7.2, 7.3 – (it is a constant challenge although evolving as time passes.)

What are the main problems/challenges you see in the next 5 to 10 years?

Network and Interfaces microarchitecture: 1.3, 1.8, 1.9, 1.10
Design tools: 2.2, 2.4
Power efficiency: 3.3, 3.5
Reliability, Fault tolerance, yield improvement: 4.2, 4.4, 4.6
Virtualization, QoS, Security: 5.2, 5.3, 5.4
Off and on-chip interconnection networks integration: 6.1, 6.2, 6.3, 6.4
Benchmarking: 7.1,7.2, 7.3 - (it is a constant challenge although evolving as time passes.)

What are the main problems/challenges you see after the next 10 years?

At this timeframe much of the progress done in interconnects clearly depends on the technology drivers of that period. For example 3D stacking technologies may change the way we view system design (how many parts of the system will actually be off-chip). Also the architectural enhancements that will emerge throughout these years are expected to give efficient solutions to many of the current hard-to-solve problems. Nevertheless, the issues that we believe will remain critical for the performance of interconnection networks will be practical congestion management techniques and power efficiency. Also, although significant progress is expected in the development of automated network-design tools in the next decade, their impact on long-term evolution is considered very important and they will remain a research challenge even beyond the next 10 years. Additional challenges are the following:

Network and Interfaces microarchitecture: 1.4
Design tools: new tools will be needed for sure
Power efficiency: 3.4, 3.6
Reliability, Fault tolerance, yield improvement: 4.1, 4.5
Virtualization, QoS, Security: 5.1, 5.5, 5.6, 5.7
Off and on-chip interconnection networks integration: 6.4
Benchmarking: 7.1,7.2, 7.3 - (it is a constant challenge although evolving as time passes.)

Name at least one subject in your field that you're not going to work on (because you don't believe it will happen, for example).

--- Extra opinions needed ----

For all of the challenges you enumerate, please give an assessment of:

Who is working on that? (Country, university, industry, ...)

European Universities and Research Institutes

ALARI Lugano, Switzerland
Cambridge University, UK
CEA-LETI, France
Ecole Polytechnique Federal Laussane (EPFL), Switzerland
Foundation for Research and Technology-Hellas (FORTH), Greece
IMEC, Belgium
KTH, Sweden
Politecnico di Torino, Italy
Politecnico di Milano, Italy
Poznan University, Poland
Simula Labs, Norway
Tampere Institute of Technology, Finland
Technion, Israel
University of Bologna, Italy
Uninersity of Catania, Italy

University of Cantabria, Spain
University of Castilla – La Mancha, Spain
University of Ferrara, Italy
University of Heidelberg, Germany
University of Jonkoping, Sweden
University of Murcia, Spain
University Politecnica de Valencia, Spain

Universities outside Europe

Carnegie-Mellon University, USA
Hong-Kong University of Technology
Georgia Tech, USA
KAIST, Korea
M.I.T, USA
Princeton University, USA
Stanford University, USA
University of California at Berkeley, USA
University of Southern California, USA
University of Texas, USA
University of Washington at St. Louis

Companies

AMD, USA
ARM, United Kingdom
Arteris, France
Cray, USA
Fulcrum, USA
IBM, USA and Switzerland
Intel Corp., USA
NXP, Netherlands
Quadrics
ST Microelectronics, Italy and France
Sun Microsystems, USA
Thales Research
Tilera, USA
Xyratex

Are there already any promising solutions? Which ones? Are they similar?

What is the activity of Europe in that? Of the Hipeac partners?

European universities and companies, with the majority of them being HIPEAC members, are highly active in the research area of Interconnects. Many of the research challenges analyzed are already confronted by the HIPEAC community and promising solutions have emerged. A lot of work needs to be done in order to effectively merge new architectural solutions with viable physical implementations.

What could be the impact if the challenges are not solved?

Interconnection networks lie at the core of any parallel computer architecture. Traditionally, they appeared in high-performance general-purpose machines, while recently they are becoming a necessity in embedded

systems. Therefore, in a general perspective, if the interconnection network does not offer the performance and the additional services required by modern applications, then we will end up with not scalable systems that are developed in an asymmetric manner. Such systems will locally have a lot of computation capabilities but they will not be able to communicate efficiently with each other (or to the associated memory hierarchy) leading to poor overall system performance. In this way any progress made in multi-core systems and in new scalable programming models and operating systems will not be able to show their full promise since the interconnection network will be the performance bottleneck.

What could be the impact if the challenges are solved?

In case that the interconnection networks provide their full potential to the attached processing nodes (and/or memories) then we believe that this will change the view that we have for overall system design. Scalability will not be an issue any more from a communication perspective but it will be up to the software to deliver the expected levels of performance exploiting in the best possible way the underlying hardware infrastructure. Additionally, power issues will be tackled in an effective manner thus alleviating the need for costly thermal management policies and enabling the design of highly sophisticated mobile devices.

Address societal problems like environment, energy, aging

Interconnects are responsible for a large amount of the energy dissipated by modern computer systems. In case of on-chip networks long wires and switches consume more than 30% of chip's power. Also, high-speed channel I/O devices consume a lot of power in order to enable multi-Gbit/s chip-to-chip or board-to-board communication bandwidth. Therefore, power efficiency of interconnection networks is a crucial parameter to their future success. Reducing the number of mWatts dissipated per data transfer (either on-chip or off-chip) will have an immediate impact on system's power consumption thus directly alleviating the energy-related negative environmental effects of computer systems.

Suggest what will create jobs, what might cost European jobs

Markets needs are focused to high-end products that will offer huge and diverse computation capabilities, seamless connectivity, while being in many cases portable handheld devices. The design of such sophisticated systems requires skilled engineers with a broad background and experience. On the other hand interconnection networks play a key role in systems' scalability and ease of integration, thus a being a key tool for the success of future products. Therefore, we believe that engineers will a solid background to interconnection network architectures and interfaces will be more and more needed by high-tech companies. This assumption holds for companies that design final products, IP core providers, as well as the design automation industry that needs to provide the appropriate tools for system level design and integration. Embedded systems industry is already seeking for engineers that will design efficient platforms, which, based on new generic communication paradigms, will execute future demanding applications. Additionally, more investment is needed in the research and development of power efficient interconnection network architectures (from system down to circuit level), in order to sustain Europe's long-standing leading position to power efficient portable devices.

Generic questions:

What will be the most important domain in 10 years from now

Hardware complexity is growing with a fast pace every year both at the chip and at the system level. Designers must integrate an enormous amount of components implementing in many cases heterogeneous functions. This situation already leads to re-using commodity components (processors – memories – specific function blocks) in order to satisfy short time-to-market and keep design complexity manageable. Future systems will be responsible for carrying out even more complex tasks requiring huge amounts of computation power. Applications do not help in this direction since it is hard to express their inherent parallelism to the

hardware. At the same time, technology and market trends (especially that of mobile devices), push for systems that dissipate a small amount of energy (or at least keeping power at a reasonable level satisfying packaging and thermal constraints).

These facts lead us to predict that the most important domain in 10 years from now will be “System integration and customization” along with the design tools needed to support system simulation and verification as well as design-space exploration with full hardware/software interaction. In fact, in this context, Interconnects research will be more prevalent since efficient methods will be required to allow system plug-and-play connectivity along with power/performance efficient communication mechanisms. In parallel, new design tools and methods will come forward in order to deal with system reliability and fault tolerance at early design stages.

Do you see other challenges in other domain that could impact (positively or negatively) your domain?

We believe that the most critical domain that will impact the performance of the interconnection network and it will heavily influence their design is the one that will decide the next generation programming models and run-time system’s communication APIs. The way processors communicate has a direct impact on how applications are written and which traffic patterns appear in the network. This software communication model and the corresponding applications will set the barrier for the required network latency, throughput and congestion management strategies. Also, how the software will be able to take full advantage of the virtualization and QoS features offered by the network will clearly determine where more innovation is needed. Finally, although not directly related to HIPEAC research clusters, technological parameters (such as 3D stacking and/or devices) are expected to influence significantly the design of interconnection networks especially for the case of Network-on-Chips.

What do you see as probable key evolutions in your domain in the next 5 years?

The following points describe what we believe will be the most probable evolution in the area of Interconnection networks in the next 5 years:

- The first generation of interconnection-network-specific design tools will emerge and they will be efficiently integrated to other design automation tools.
- Full system simulation environments, which have recently included interconnection network elements, will be more mature and widely used thus improving in this way the validity of the results and clearly distinguishing the parts of the network that needs to be optimized. A first class of network-specific benchmarks will emerge.
- At the micro-architectural level (switch design, topology, routing algorithms design and implementation) a lot of new designs are expected that will bring system throughput at an acceptable level will tackling latency issues.
- New network interfaces are expected to appear that will offer the needed hardware primitives for speeding up interprocessor communication.
- Further new research results in the area of fault-tolerant interconnection networks and reconfiguration strategies.

In the next 5 to 10 years?

For this longer timeframe the issues that will be solved at least at an acceptable level are the following:

- More mature design tools that besides simulation and verification capabilities will automate a large part of network design.
- Power/thermal-conscious interconnection network components, enhanced with network monitoring and debug features.

- In-network support for system virtualization.
- More efficient integration of on-chip and off-chip networks. Maybe definition of new standards that will cover system connectivity in a unified manner.
- New communication APIs and applications that will allow clear network benchmarking.

Beyond 15 years?

New opinions are needed.....

(This is different from the first questions: current and probable evolution could ignore a problem/challenge !)

Give your view of the future in the domain of Hipeac

Other problems? What technical issues can solve them? Technology drivers?

In our view, a key technology driver for the evolution of future computer systems will be power efficiency. All innovation in hardware platforms and software applications will be limited by the amount of energy spent. Therefore, whatever architectural techniques or future implementation devices will help in delivering the complex computation and communication tasks of modern computer systems within a fixed power budget will be the determinant factor of future evolution.

High tech opportunities in the next 10 years?

We believe that in the next 10 years the majority of the research and development effort will be focused in the practical description and exploitation of parallelism in applications. Describing parallelism is a burden that future parallel programming models need to solve in an effective manner allowing the average programmer to take full advantage of the underlying hardware. In turn, scalable parallel hardware architectures that will offer the expected high computation power, while being relatively easy to design will clearly determine the success of computing systems industry in the next decade.

Directions where we lost jobs if we don't invest?

As already clearly identified by the HIPEAC community the new computing systems environment is characterized by the convergence of the commodity, embedded, and supercomputer markets. At the same time, evolution and innovation are constrained by design complexity crisis and by power consumption limitations, where performance gains will no longer be possible by relying just on technology evolution. The technology wave is directed at the commoditisation of the technology base that was previously associated with the high-end (super-) computing platforms. Key building blocks have been identified which collectively reduce the burden of the processing engines and create a truly distributed processing environment. The applications of this technology platform extend into data center, networking, Grid computing, video server, and wireless applications. Therefore, to ensure European competitiveness in such a demanding environment we need to deliver the next generation platforms that remove bottlenecks and constraints from the existing IT infrastructure technology. A combined effort and associated investment is required that will push Europe in leading the design of these new hardware platforms that will be efficiently exploited by novel higher level software applications.