NaNoC Summer School 2012

June - 11

S1.1: Interconnection networks in next- generation many-core SoC platforms.
Speaker: Luca Benini (Università di Bologna)
On-chip networks have been touted for many years as "one-stop solution" to address the key on-chip communication scalability issues. This viewpoint is very simplistic. The hard reality of commercial deployment demonstrates that even though NoCs are the only way to go, we are far from having a complete and future-proof solution, and that lots of effort has to go in designing networks that are providing many services and functions. In this talk I will survey the steps that are being taken in NoC design technology and architecture to provide industry-viable solutions as energy efficiency, virtualization, fault-tolerance, ... Configurability takes the center stage and cannot be "simplified away".

S1.2: Energy-efficiency in Network-on-chips: An Industry Perspective
Speaker: Sriram Vangal (PhD,Principal Research Scientist, Intel labs, Intel Corporation, Hillsboro, Oregon, USA)
Commercial designs currently integrate tens to hundreds of embedded functional and storage blocks in a monolithic SoC, and the number is expected to increase significantly in the near future. The possibility of such high levels of integration in a single chip necessitates the design of high-bandwidth, low-power interconnect architectures. The talk shares key learnings obtained from silicon implementations on two generations of NoCs- the "80-tile Polaris TeraFLOP processor" and the "48-iA core Single-chip cloud-computer" (SCC) with a key focus on improving energy-efficiency in the interconnect. The data confirms that in this age of interconnect-centric VLSI, computing locally is inexpensive but moving information becomes the limiting bottleneck. The discussion presents challenging trends and reviews promising solutions to improving NoC energy-efficiency and provides suggestions for future research.

S1.3: Industry Requirements for Network-on-Chip Interconnect Solutions.
Speaker: Helmult Reinig (Intel Mobile Communications)
This talk will highlight industry Requirements to a Network-on-Chip solution from Intel Mobile Communication. Such requirements do not only deal with the features of the NoC for communication (multi-protocol support, latency, throughput, scalability, routability, area, power consumption). Test and debug features are of similar importance, as well as reliability issues in small geometries. And not to forget the tooling to assemble and optimize NoC solutions with tight time-to-market requirements, which has to fit into existing implementation tool flows from RTL to GDS.

S1.4: The NaNoC project overview.
Speaker: José Flich (Universitat Politècnica de València)
In this presentation, the overall picture of the NaNoC project will be provided. The basic goals and objectives of the project and how the project is structured will be presented. The talk will briefly present the different design methods and tool extensions performed within the project and how they will be combined.

June - 12

S2.1: Fault-Tolerance support in Unreliable NoCs (unicast communication).
Speakers: Samuel Rodrigo (Simula Research Lab.) / Frank Olaf Sem-jacobsen (Simula Research Lab.)
In this session, the basics of routing in interconnection networks will be addressed, focusing on key aspects as deadlock issues and efficient implementations of the algorithms. The topic will be focused on networks-on-chip, which may experience manufacturing defects. This topology change has severe implications on how packets should be routed to avoid deadlocks and excessive resource needs. The session will review the bibliography and will focus on current efficient solutions that address these problems in an implementation-wise manner, enabling the existence of manufacturing defects.

S2.2: Fault-Tolerance support in Unreliable NoCs (collective communication).
Speakers: Samuel Rodrigo (Simula Research Lab.) / Frank Olaf Sem-jacobsen (Simula Research Lab.)
Collective communication enables fast communication between sets of nodes, where typically one sender sends the same message to a set of destination end nodes (or vice versa). Collective communication is highly demanded in chip multiprocessors with cache coherence protocol implementations and barrier synchronization primitives.
In this session, in a first step the need for such communication means will be explored and then efficient solutions will be described and detailed. The session will focus on an unreliable environment where manufacturing defects can be introduced thus changing the network to an unexpected topology at design time.

S2.3: A cooperative testing strategy for NoCs.
Speaker: Alessandtro Strano (Università di Ferrara)
A key requirement to sustain yield in future parallel hardware platforms is to be able to isolate defective components so that system operation is preserved. Detecting the presence of faults in the NoC is a non-trivial and expensive process, given its distributed nature throughout the chip and its lack of directly accessible I/O pins. These features play against traditional scan-based approaches and make a built-in-self-testing strategy appealing. This session reviews traditional testing strategies for digital circuits and proves their inefficiency in a NoC setting. For this reason, a cooperative approach to NoC testing is then presented, exploiting the links as a means of exchanging test pattern and/or test responses between neighboring switches. Key features of cooperative testing are the parallelization of the testing procedure, the fault coverage for test pattern generators and the marginal test execution time for use not only in post-production testing but also at each system bootstrap.

S2.4 Error detection and notification infrastructure for self-reconfiguring NoCs.
Speaker: Alessandtro Strano (Università di Ferrara)
Aggressive CMOS scaling accelerates transistor and interconnect wearout, resulting in shorter and less predictable lifetimes for microprocessors. Studies show that wearout faults have a gradual onset, manifesting initially as timing faults before eventually leading to hard breakdowns. This session presentes design methods for detecting wearout faults as they begin to affect normal operation of NoC switches. Upon detection of such events, this session presents notification techniques (and system infrastructure support) to a centralized manager, which is in charge of reconfiguring the network around the malfunctioning component (link, switch or switch sub-blocks). This task is challenging since the notification mechanism should be itself fault-tolerant and marginally affect NoC area footprint and power. Also, reuse opportunities of the system infrastructure for other notification purposes will be illustrated.

S2.5: Communication Exchange Format: from Design Intent to Implementation.
Speaker: Federico Angiolini (iNOCS)
The design of a NoC is an exercise that requires multiple inputs, from different design teams working at different stages of the design flow, and produces numerous outputs. Furthermore, incremental NoC redesign and optimization is a reality in industrial flows, where specifications may be updated on the fly, software requirements change in parallel with hardware design, verification results may mandate design changes, and feedback from the back-end team is expected close to the tapeout deadlines.
The NaNoC project devised CEF, or Communication Exchange Format, to ensure agile interoperability among NoC EDA tools, such as e.g. synthesizers, routers, simulators, verifiers and floorplanners. CEF is an XML-based file format aimed at describing the relevant inputs and outputs of a NoC design flow, at the level of abstraction most suitable for interconnect designers. CEF can model IP cores in the design, communication requirements among them, chip operating modes, floorplan, and of course NoC topology, architectural knobs, and routing. NoC tools can add, modify, reset or annotate information as the design process unfolds, promoting iterative and incremental flows in addition to linear ones, and easing automation.

June - 13

S3.1: Addressing Virtualization and Power Efficient Support in Future NoCs
Speaker: José Flich (Universitat Politècnica de València)
The network-on-chip component, as it becomes more mature, will migrate to a dynamic approach, where it reconfigures to the changing requirements of higher system levels. This is the case of mapping different applications at the same time onto the same chip. Different applications will have different requirements and traffic isolation will be needed. Also, as the system size increases, the number of unused components for large periods of time will increase, demanding power-efficient management of such resources. The network-on-chip will need means to power down these resources. In this session, dynamic mechanisms to enable such behavior will be reviewed. Focus will be on a complete reconfiguration mechanism that will enable the smooth and transparent transition.

S3.2: Addressing Congestion and Process Variation in Future NoCs
Speakers: Samuel Rodrigo (Simula Research Lab.) / Frank Olaf Sem-jacobsen (Simula Research Lab.) / Federico Silla (Universitat Politècnica de València)
As we keep further stressing the chip with more components, it is clear that new challenges will arise. On one side, at the architectural level, the network-on-chip will become stressed with higher levels of communication needs, thus driving the network to a possible congestion situation. This drives the entire performance perspective to low values. In this session, we will address the problem of congestion in high-performance networks and will differentiate the main sources of congestion as well as the different typical approaches followed in the literature. On the other side, at the transistor level, process variation will play a contaminating role since transistors with different operating frequencies will drive different performance values. In this respect, we will review which are the main sources of process variation and how they affect the components in the chip. We will also address how to minimize the effects of process variation on links and routers, as well as how to lower its impact on overall CMP performance.

S3.3: The NaNoC switch.
Speaker: Bertozzi Davide (Università di Ferrara)
Current switch architectures are conceived to meet the requirements of current technology nodes. In general, they exhibit poor robustness to network irregularities that may arise as an effect of unwanted effects (manufacturing faults, process variations) or intentional network configuration strategies (application, power and thermal management policies). Moreover, the fully synchronous timing paradigm is often assumed. NoC design platforms in the next three to ten years will require network switching fabrics with fundamental advances with respect to routing mechanisms, reconfiguration circuits, testing strategies, synchronization paradigms and virtual channel implementation. This session presents a new switch architecture that embodies the innovative design methods of the NaNoC project, aiming at meeting the requirements of a nanoscale NoC. Emphasis will be given to the interaction between different design requirements and to the area, power and latency overhead that design techniques incur.

S3.4: The Teklatech NoC/IR-Drop-aware Floorplanning Tool
Speaker: Mikkel Stensgaard (Teklatech)
At the architectural level of large-scale ASIC design, logical and physical design decisions are closely intertwined. Since system-level modules will ultimately be laid out on a 2D surface, the placement of these modules in the physical floorplan and the communication infrastructure must be taken into account concurrently. Meanwhile physical-level challenges such as power integrity (IR-drop and power grid noise) and pad placement places constraints on the floorplan. Current ASIC backend tool flows fall short in including more abstract, higher level design requirements; fully wired systems are expected as input. This causes a need for manual iterations in the design flow, in order to arrive at a valid floorplan that matches a valid communication infrastructure, and fulfills given physical design constraints. Teklatech has developed a floorplanning tool which places system-level blocks before the system-level communication infrastructure connectivity is defined. It takes as input abstract communication requirements, and places blocks, between which there is a high degree of communication, closely together. A prototype power delivery network can also be specified, and power integrity is taken into account concurrently, by a fast power rail analysis engine which is embedded directly into the floorplanner. The floorplanning tool will be demonstrated in the context of an automated NoC synthesis flow.

S3.5: The iNoCs NoC Synthesis Toolchain
Speaker: Federico Angiolini (iNOCS)
NoC designers are faced with a large number of constraints (protocol interfaces, frequency domain crossings, bandwidth, latency, power, wirelength) and an even larger set of design options (NoC architecture, architectural parameters, topology, routing, placement). Although traditional design methods, born in the shared bus era, have relied on the skills of a system architect to account for all the variables, the problem scale is growing to the point where design automation can help the architect converge more quickly towards a working solution that also improves power, performance and area metrics. iNoCs has developed a NoC synthesis toolchain that can take as inputs a set of system constraints, for example the communication requirements and floorplan, and can explore the NoC design space, especially in terms of topology, routing, architectural parameters and NoC component placement. As an outcome, a set of automatically-generated topologies is presented to the user, who can then choose the best options based on rankings by power, latency, wirelength, area, etc. A demo will be shown to illustrate the iNoCs toolchain operation.