

| | | | | | |
|------------------|--------------|-------------------|-----------------|------------------|-------|
| <i>date/date</i> | 1-10-2009 | <i>ref./réf.</i> | NoC round table | <i>page/page</i> | 1 / 5 |
| <i>from/de</i> | Laurent Hili | <i>visa/visa</i> | | | |
| <i>to/à</i> | | <i>copy/copie</i> | | | |

Subject/objet NoC round table synthesis

**Network on Chip round table
ESA – ESTEC
17th / 18th of September 2009**

Synthesis note

Introduction:

The first NoC round table co-organised by ESA (Laurent Hili) and CNES (Philippe Perdu) was the opportunity to bridge the NoC designers / researchers community with the Space community (micro electronics and data handling engineers). The ever increasing complexity of on board systems put ever more constraints on ASICs and SoCs embedded in those equipments. The advent of Deep Sub Micron technology for Space (ST 65nm) will open new horizons for integration but will also set new challenges. How the Space Community will tackle tomorrow SoC designs exhibiting complexity beyond 20 to 30 Millions gates?

Current SoC solutions are built upon bus topologies, but buses do not scale well when the number of IP blocks increases. Multi layers buses solutions may address the short / medium term complexity increase (tens of IPs) but will certainly not be able to cope with the medium / long term trend where hundreds of IP blocks might be integrated on a single die.

NoC paradigm has been an active field of research over the last decade. This field of research has grown with the necessity of higher design productivity and its corollary technology scale down (65nm, 45nm, 32nm, 22nm ...). How can we maximize IP blocks reuse (plug & play) and how can we overcome some of the cumbersome aspects linked to bus topologies (place & route congestions, timing closure, signal integrity).

Through the various presentations given by Research Institutes and private companies, the ESA-CNES round table had for objective to identify how existing solutions developed in Europe could potentially be adapted and transposed to the Space sector. What might be the benefit for ESA and CNES to invest in such a technology and what might be the delta efforts needed to mature this technology for Space applications.

Why NoC? Advantages over bus based solutions:

During the round table several topics were addressed. There was a common agreement to state that decoupling the processing functionality from the communication functionality as proposed in NoC paradigm offers significant advantages over conventional bus topologies approaches. Bus is a shared medium which doesn't scale well when complexity increases. The performances of bus (throughput) collapses with the loads or number of IPs blocks interconnected. On the contrary NoC is a point to point interconnects structure allowing therefore a lower load and higher throughput while authorising several initiators (masters) to be active at the same time.

Synchronous or asynchronous NoC?

The choice of one versus another is not straight forward and obvious. Synchronous solutions are probably simpler to implement and better suited with conventional design flow methodologies (logic synthesis). A synthesisable NoC IP could also target various ASICs or FPGAs technologies. On the other hand synchronous solutions do not maximise speed offered by advanced CMOS technologies since signals have to be pipelined, timing slacks taken and finally synchronous solutions can hardly cope with process, voltage and temperature variations (PVT). Asynchronous solutions require extra care with the re-synchronisation and buffering but they can on the other hand alleviate some of the limitations mentioned here before. An efficient implementation of an asynchronous NoC IP will usually involve a full custom design (hard IP), which means no portability from one technology to another.

Real time and quality of service (QoS):

Other debates focused on the facts that we may have to dissociate NoCs dedicated for satellites platforms (TM/TC, AOCS) from NoCs dedicated to satellites payloads. The first category involves hard real time constraints whereas the second category can cope with soft or no real time requirements. Discussions also took place with respect to quality of service (QoS), how can we guarantee bandwidth allocation for hard real time applications for instance? Virtual channels combined with circuit switching techniques could obviously bring some elements of answer. Multicast techniques might also be needed to distribute time (time tagging) in hard real time applications. On the other hand payloads related applications may cope with less stringent QoS, for instance image processing may afford packets losses in certain circumstances but packets have to be delivered in order whereas other applications may even cope with not in-order delivery. We see here that QoS requirements may vary from one extreme to another leading to different implementations, circuit switch versus packet switch. For a future NoC

implementation we may have to consider the most constraining case, hard real time applications associated with high QoS (guaranteed throughput). Nevertheless whatever the future solution it should have the ability to adapt to various QoS requests.

Network interfaces (NI):

Regarding IP blocks interconnect or network interfaces (NI) two standards have been endorsed by the NoC community, AMBA 3.0 (AXI) and Open Core Protocol (OCP-IP). Both solutions are point to point protocols, AMBA is an ARM proprietary standard whereas OCP as mentioned in the name is an open standard. ESA has undertaken activities to deliver key IPs such as Leon and SpaceWire with OCP interfaces.

Testability, fault tolerance and dynamic reconfigurability:

Other topics of high importance for the Space community, testability and fault tolerance, were also debated during the round table. Concerning the first point, testability, one of the major problems to which SoC designers are confronted today is observability. How could we in complex SoCs monitor the behaviour of a particular IP block and not only basic errors due to manufacturing such as “stuck at fault” (open nets, shorts)? Stuck at fault are usually well handled by known techniques like “scan path”. Nevertheless despite obvious advantages this test has also limitations. One of the limitations is linked to the fact that the test has to be run on a specific tester and not in-situ (on board), second drawback, only basic errors (stuck at) are addressed, scan path can’t easily detect malfunctioning IP blocks due to soft errors (SEE). The inherent network infrastructure made available by NoCs could also be used to enhance testing. Usually when an in flight ASIC is malfunctioning it is simply switched to its redundant part or even redundant equipment if there are no redundancy schemes at ASIC level. NoC could open new possibilities to perform in-situ and on-line (in flight) tests. Thanks to the network infrastructure complex SoC could gain in observability, faulty nodes could be finer analysed while ASIC being still operated in flight. If finer in flight diagnostics could be established then cleverer mitigation strategies could be envisaged; reconfiguration of the faulty node in case of “soft error” or replacement with a spare node in case of “hard error” (permanent error, stuck at fault). Instead of replacing the complete ASIC function or complete equipment then only the faulty node within the NoC could be reconfigured or replaced at the expense of marginal performances decay (graceful degradation). Regarding fault tolerance applied to NoCs, only few research institutes have looked in depth to those aspects. A very interesting presentation was given by Martin Radetzki from Stuttgart University showing how simple mechanisms could be efficiently implemented on NoCs to allow automatic identification of faulty nodes and automatic traffic rerouting.

Complementary activities coordinated by ESA and CNES would be needed to deeper investigate NoCs testing aspects, off line (on ground) and on line (in flight) but also investigate associated dynamic reconfiguration techniques to cope with soft and hard errors (fault tolerance).

Convergence between off and on chip networks:

To conclude this round table some debates took place to assess if the convergence between off and on chip networks would make sense. The answer was, yes, convergence with SpaceWire for instance is perfectly feasible. Investigations in that direction have already been undertaken by Steve Parkes (University of Dundee), Kees Goossens (NXP) and Björn Osterloh (Technology University Braunschweig). The idea of adapting an existing SpaceWire router solution is probably not good since off and on chip constraints might be divergent; nevertheless adapters at interfaces level could be developed to ease this interfacing while minimising the overhead (logic gates). SpaceWire Remote Memory Access Protocol (RMAP) could also be a good candidate to access off and on chip nodes indifferently.

Tools:

Finally tools are also a major issue since no commercial tools exist yet allowing NoCs dimensioning and traffic modelling. Tools so far have been essentially developed by Research Institutes, but question is, how reliable and mature are those tools? In order to allow an efficient deployment of future NoC solutions, SystemC transactional models (TLM 2.0) of all key IP blocks will be needed, interfaces of those IPs will have to be standardised (AXI and/or OCP). Making available VHDL and TLM SystemC models will ease architecture trade off and portioning between HW/SW (co-design). For traffic modelling even higher abstraction models might be needed to speed up simulations and trade offs (time and space exploration). Those models might be SystemC data flow models, less accurate (no bit and cycle accuracy) but much faster at execution time. Finally we will have to ensure a full coherency between those abstraction models, C or Matlab application → SystemC (Data Flow) → SystemC (TLM) → VHDL (RTL).

Conclusions:

This first edition of the NoC round table was a success with the participation of ~ 40 attendees from various horizons most of them having no related links with Space business. Participants came from all around Europe and even from Russia and Israel. The goal of bridging the NoC community with Space community had been a success and new contacts have been established opening the road for future cooperation. Several goals were set at the start of the round table by Philippe Armbruster (Head of the Data Systems Division at ESA).

1. Is NoC an interesting concept for Space applications?
2. Among existing solutions are there emerging candidates?
3. What would be the delta effort needed to adapt those solutions for Space?

To the first point the answer is clearly, yes, NoC could bring significant advantages over conventional solutions in terms of testability, fault tolerance, reliability and dynamic reconfigurability. Regarding point number two, the answer is, no, and deeper

investigations are needed. We will have to answer and assess whether we want to develop a NoC synthesisable and portable IP across various technology platform or if we want a hard IP, non portable but optimised for DSM technology (ST 65nm). Point number three, the delta effort needed to adapt a solution to Space needs, first a solution or a set of solutions have to be identified (point number 2) then the non exhaustive list below describes the key topics which might be addressed in priority.

- Testing (off / on line)
- Reliability and NoC ability to cope with PVT in DSM technologies
- Fault tolerance
- Quality of service
- Convergence off / on chip (SpaceWire)
- Design tools and traffic modelling tools

As a conclusion, the recommendation for ESA and CNES is to place a pre-study to first identify a solution or a set of solutions, then we will have to answer to the question soft or hard IP, then we will have to address the points mentioned in the non exhaustive list here above.