WHITE PAPER – April, 2015
Aims Technology Inc.
3080 Olcott St. Ste#D235, Santa Clara, CA 95054

Future of SoC Designs: Network-on-chip

By Purna Mohanty, Co-founder and CEO

Introduction

History speaks the future. Since late 90s through now, *System-on-chip (SoC)* designs have come a long way from its predecessor name *Application Specific Integration Circuits (ASIC).* As a matter of fact, ASIC derived its name from IC (Integrated Circuits) – when multiple components went into a single component. Then when an IC got tailored for a particular application as opposed to general purpose, it became known as an ASIC. Multiple ASICs on a platform were used to build systems. Thanks to the deep sub-micron semiconductor technologies, we were able to put in an entire gamut of ASICs used to build a system on one ASIC itself, and coined a new name – System-on-chip (SoC). So now what is next? Is it *Network-on-chip (NoC)?*

Trend

Usually in early ASIC days, the processor itself was a chip, and was sitting alongside one or more ASICs communicating via peripheral bus like PCI/PCI-X. As semiconductor technology enabled higher level of integration, processors also got into ASICs, small and large SRAMs got into ASICs, and that is when world of embedded software development began. This is where the life of hardware engineers became challenging, and always they live in the world of innovation.

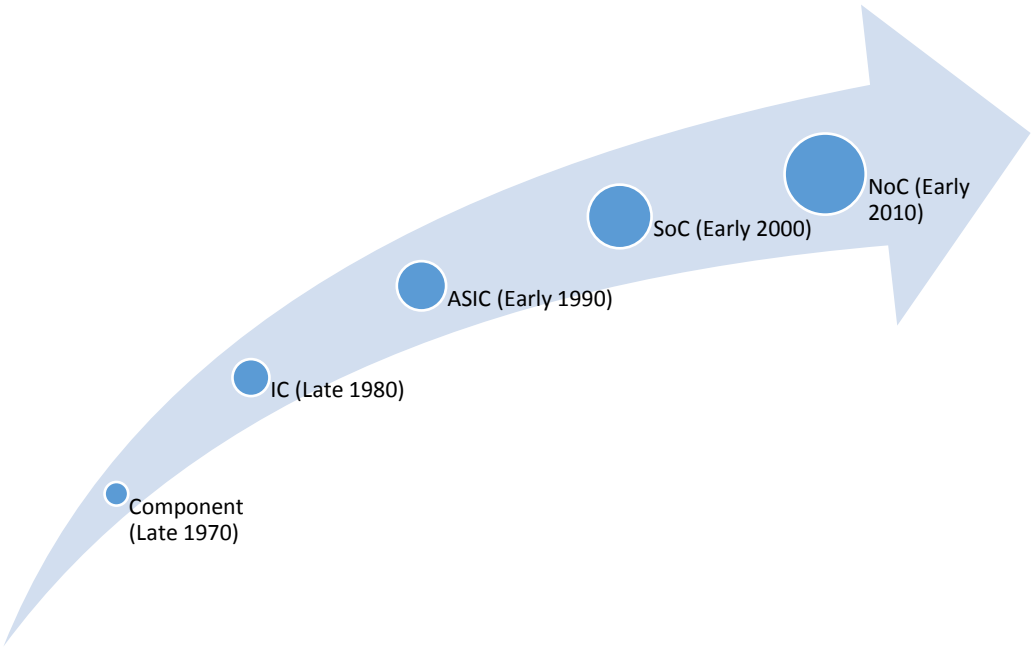

As the SoC became more complex with the increasing number of low speed, medium speed, and high speed peripherals, this gave way to the design of on-chip interconnect IP required to facilitate effective communication amongst all IO blocks and internal blocks of a SoC.

The on-chip interconnect is alternatively termed as Network-on-chip (NoC). As you can see, when the number of communicating agents is limited, and the traffic pattern is homogeneous, it is easy to design the NoC IP, and facilitate an effective communication among the stakeholders. However, the complexity goes up when one or more of the following things happen –

- Number of agents goes up from single digit to double and even triple digits
- The max bandwidth (bus width and the clock) of each agent is different from the other.
- The traffic pattern is unique for some components. Some I/Os require isochronous traffic.
- The bandwidth requirement may change based on the time of the day.
- When multiple agents strive to access a central resource like DDR or SATA

This is where the innovation in scheduling comes in so that NoC is fair to all, meets their need, none is starved. On the top of all these complexities, the use of such devices in mobile or battery environment has given rise to the requirement of power savings and management.

Reduced power consumption means higher battery life, and thus higher useful life of the gadget. So different techniques of power savings comes in – switch off the clock for the unused logic and/or switch off the power to the unused logic. The switch-on and switch-off of such controls is accomplished through software. Now is the time to go the next level of details.

Network-on-chip (NoC)

As we saw previously, complexity in connecting internal agents of a SoC gave rise to the development of Network-on-chip IPs. As the future of SoC largely depends on NoC IPs, let us try to capture the future of such critical IPs, and thus delineate the future of SoCs.

As you see in the illustration below, the functionality – bandwidth, latency etc. requirement of each agent is achieved first by the interconnect IP. Once this is achieved, the next bottleneck came from the physical side of the design, where the datapath and control path wires need to be routed amongst different agents. With the ever increasing need for higher bandwidth, the width of datapath and clock frequency keeps going up. Now-a-days, the datapath width has gone up from simple 32b wide to 512b wide, and the clock frequency has gone up from 125/250MHz to 3.0GHz and beyond. Just imagine the multiplying complexity of the NoC design

as a result of such requirements in terms of physical routing of wires across agents and meeting the stringent timing requirement.

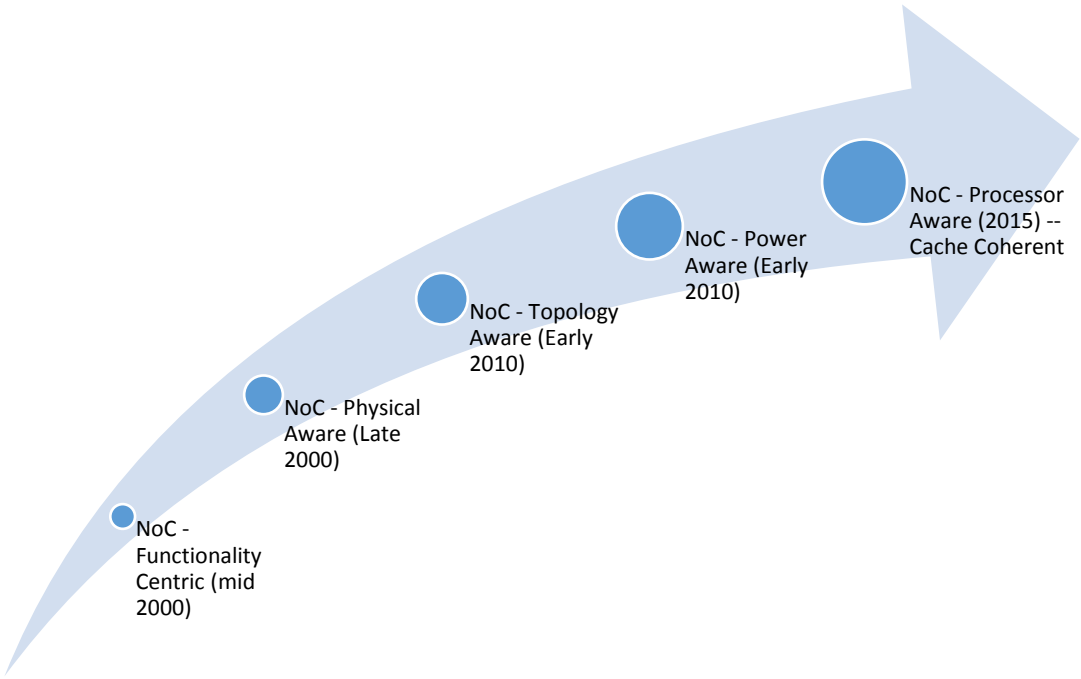

In order to meet functional requirement along with the mitigation of timing and routing congestion, topology of NoC became a key driving factor. The entire NoC can be broken down into island NoCs of ring type, mesh type and cross bar type, and each island can communicate with the other through a thin but high speed path.

With the advent of mobile and battery-operated applications, it is required that we design the architecture and micro-architecture which is power aware, and we do logic design which is power aware. We become creative in our understanding of the system level traffic pattern and bring in dynamic control to the power consumption.

Finally, as the compute power keeps going up, and bandwidth of each IO keeps going up, there is a need for NoC to deal with cache coherency across processors on chip so that it limits off-chip memory access and share data within the chip. The NoC is now required to handle both cacheable and non-cacheable access within a single network. This paved the way to develop the next generation of on-chip interconnect – *cache-coherent NoC (CcNoC)* that supports cacheable, cache-aware, and non-cacheable communications.

Conclusion:

At Aims Technology, *we design the CcNoC with the vision of its long haul usefulness with customer's applications in volume production. We keep return-on-investment in mind at customer end during integration, simulation, validation, and we have implemented robust DFM/DFT schemes for yield and the time on ATE.* The IP is designed with future in mind – data collection, security, storage, computing, and digestion of data in emerging compute intensive areas in IoT and IoE evolution.