Arteris™

THE NETWORK-ON-CHIP COMPANY

From “Bus” and “Crossbar” to “Network-On-Chip”

Arteris S.A.

www.arteris.com

Contact information

Corporate Headquarters

Arteris, Inc.
1741 Technology Drive, Suite 250
San Jose, CA 95110
U.S.A.

Tel.: +1 (408) 625-6000
Tel. Sales: +1 (408) 313-7281
Fax: +1 (408) 625-6040

European Engineering Center

Arteris SA
6 Parc Ariane – Immeuble Mercure
Boulevard des Chênes
78284 Guyancourt Cedex
France

Tel.: +33 (0)1 61 37 38 40
Fax: +33 (0)1 61 37 3

From “Bus” and “Crossbar” to “Network-On-Chip”

Network-On-Chip (“NoC”) technology is rapidly displacing traditional bus and crossbar approaches for SoC on-chip interconnect. What defines a NoC ? Many terms are being used in the industry: On-Chip Networks, Interconnect Fabrics, Networks-On-Chips and so on. Let’s use the term “on-chip interconnect” as an umbrella name for all approaches. The Network-on-Chip is one specific architecture and is defined as “an on-chip interconnect with decoupled transaction layer, transport layer and physical layer”. This article explains the underlying concepts of Network-on-Chip technology, and details why NoC technology is superior to bus and crossbar approaches.

Bus and Crossbar Approaches

The first generation of on-chip interconnects consisted of conventional bus and crossbar structures. In a bus, the interconnect is mostly just wires, interconnecting IP Cores, combined with an arbiter that manages the access to the bus.(Figure 1). System-level latency and bandwidth constraints led to a natural evolution towards a multi-tiered bus architecture, typically consisting of a high-performance low-latency processor bus, a high-bandwidth memory bus and a peripheral bus. Example bus architectures are the AMBA family (AHB/APB) from ARM, Inc, and the CoreConnect architecture from IBM.

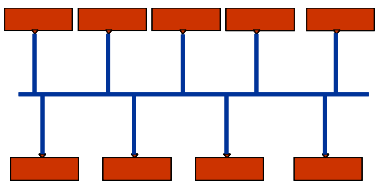

Figure 1: Traditional Bus Interconnect

Architecturally, busses have several limitations. Latency issues arise when high priority accesses are stalled by transactions in progress on the bus. Bandwidth is limited by clock-frequency, which itself is limited by physical design parameters such as the length of the wires. It is very difficult to overcome such physical limitation as pipelining a bus interconnect is very challenging. Faced with increasing IC performance requirements, designers started to implement cross-bar structures (Figure 2). These structures improve latency predictability and significantly increase aggregate bandwidth, at a not-insignificant cost, of course, of a much larger number of wires

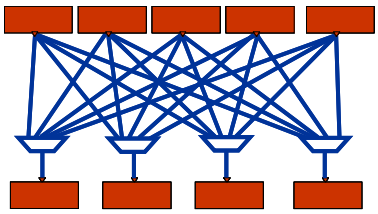

Figure 2:Traditional Crossbar Interconnect

"Decoupled" Busses and Crossbars

The advent of the SoC, incorporating tens to hundreds of IP cores created a significant integration challenge. The above described busses and cross-bars are "coupled" solutions. The interfaces of all IP cores connected to a single bus or cross-bar must all be exactly the same, both logically (signals) and in physical design parameters (clock-frequency, timing margins). This turned out to be significant obstacle to rapid integration and re-use of existing IP into increasingly complex SoCs. The ad-hoc solution was to either change the IP core socket to make it compliant to the latest choices of bus specifications, or, to develop set of bridges and protocol converters. This integration challenge was architecturally addressed by decoupling the

interface of the IP Core from the actual interconnect. This class of on-chip interconnect can be described as a "decoupled". The IP Interface Protocol, the *Transaction Layer*, is decoupled from the interconnect communication protocol or the *Transport Layer*; that is, different protocols are used to (1) communicate (1) from the IP to the interconnect and to (2) transport the data within the interconnect. In effect, all "bus"-related logic was moved away from the IP core into an active interconnect, making the IP Core "bus-agnostic" by absorbing some of the bridging, synchronizing and bus-control functions into the interconnect and thus improving interoperability both at logical and physical levels (Figure 3). The industry saw an emergence of IP Interface protocols, most notably AHB-Lite and AXI from ARM, Inc., and OCP from OCP-IP. These "decoupled" interconnect solutions address the IP re-use issue and logical integration challenges, but no significant advances were made with regards to the actual interconnect architecture. Simply stated, these technologies can be seen as busses and crossbars with configurable interfaces.

Figure 3:Decoupled Bus and Crossbar Approaches

The Network On Chip Architecture

Today's SOC designers face a host of new challenges in the design of the on-chip interconnect beyond the natural evolution of an increasing number of IP Cores and higher performance requirements. These include power management, multiple clock domains,

security requirements, error handling and others. Moreover, physical implementation through floorplanning and place-and-route is made increasingly difficult by competing and sometimes conflicting area, wiring, frequency, clock domain and voltage domain constraints. The key to addressing these issues in the interconnect is to decouple the *Transport Layer* from the *Physical Layer*. This is achieved by architecting the on-chip interconnect based on a packet-based Transport Protocol. In such a solution, the requests (and responses) at the IP interfaces are formatted into a simple packet, consisting of a header and a payload, and transported through the interconnect through unidirectional links. (Figure 4).

The decoupling between the *Transport Layer* and *Physical layer* is achieved in two ways: *logically*, because the packets are locally transported over physical links using framing that is most appropriate for the physical constraints such as a simple handshake for local interfaces or credits for low-pin count channeled interfaces; and, *physically*, because the physical attributes of each link (width, frequency) can be selected according to the specific latency and bandwidth constraints of the link

Using NoC technology, the SOC interconnect is constructed by interconnecting any number of NoC interconnect IP elements, such as switches, synchronizers, width converters, power isolator cells and others through a set of links into a user-defined topology.

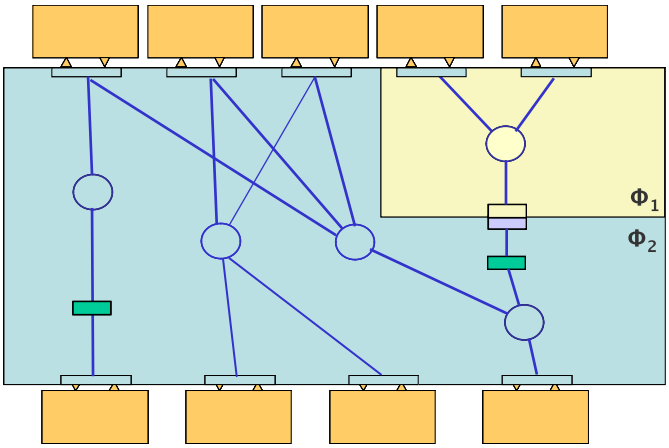

Figure 4: NoC Architecture Enables Flexible Topology

This "true" Network-On-Chip architecture has many advantages over bus and crossbar approaches. It is eminently scalable and flexible. In contrast to the above described legacy solutions where the system architect is forced to map a system architecture onto a set of busses/crossbars, using the NoC approach, the interconnect can be constructed into an user-defined topology matching the designers vision of the optimal architecture.

The NoC architecture allows for an orthogonal approach to an SoC's design challenges. Power domain isolator cells, security IP elements, width converters and synchronizers can be easily inserted on any link, and on a link by link basis. This allows the IC designer to follow an approach where requirements are addressed independently. As an example, a NoC can be constructed by first creating an interconnect topology based on connectivity, QoS and clocking requirements, then focusing on reducing area by selecting the optimal link widths for each link based on latency and bandwidth needs, then, partitioning the design into power domains, and finally, during the physical design phase of the IC, maybe inserting an additional repeater by a simple click of a button.

Because of the elegant decoupling of the *Transport layer* and *Physical layer*, very high clock frequencies can be achieved. Logic gate complexity is minimal since during the transport of the packet, the content of the payload does not need to be evaluated and or transformed. As an example, a 64 to 32 bit width-converter simply sequences every word of 64 bits into two 32-bit words, without ever having to examine any content of the packet (a word being defined as a transfer of data over a link during a single clock cycle).

Since the physical properties of each link are configured solely based on the specific requirements of the link and independently from other links, the total number of wires in the

interconnect is significantly reduced (up to 40%), resulting in smaller designs and considerably faster timing convergence during the back-end.

While, to some people, the idea of a Network-on-Chip for SoC interconnect might still seem futuristic, NoC technology is rapidly being adopted in SoC designs of all levels of complexity. The reason is simple: true NoC implementations of SoC interconnects are consistently proving to be significantly smaller, faster and more power-efficient than traditional or decoupled busses and crossbars.