



OBJECTIVE ANALYSIS

Semiconductor Market Research

NoC INTERCONNECT IMPROVES SoC ECONOMICS

Initial Investment is Low Compared to SoC Performance and Cost Benefits

As systems on chip (SoCs) have grown in complexity and cost, one of the more critical factors for an SoC's economic success in the marketplace has become its interconnect.

The interconnect has a significant impact on SoC costs because it influences four key factors of SoC design: die size, power consumption, design time, and performance.

- Die size can increase if conventional interconnect routing wire and gate area requirements explode due to the increasing number of IP blocks in a SoC.
- Power consumption can mushroom if an SoC's interconnect cannot easily be configured for advanced power management schemes like dynamic frequency and voltage scaling (DVFS).
- Project design time can extend if the SoC's interconnect becomes difficult to configure and verify. This can slow downstream designs if a platform-based design methodology is used as a basis for derivative SoC designs.
- Performance can suffer if an interconnect approach cannot adapt to changing requirements over the SoC's design cycle and product life for changing SoC IP blocks, QoS, bandwidth, latency, security and clock frequency.

As the one piece of IP that touches all other functional IP blocks in an SoC, the

interconnect, along with its configuration, verification, and utility, is the critical path to ensuring the proper operation of a modern SoC.

Typically, design teams approach the interconnect using either IP that is provided for free by the suppliers of IP blocks, or they create their own in-house solution. These methods cannot easily adapt to changing SoC feature and process technology requirements, driving the team to invest in interconnect technology and products that will meet their needs into the foreseeable future.

Network on chip (NoC) technology is a relatively new approach to signaling that enables not only more efficient interconnects but also more efficient design and verification processes for modern SoCs. NoC is an approach to signaling that matches the needs of the signal to various communications protocols in a way that reduces the complexity of the chip's interconnect. Slow or low bandwidth signals can be multiplexed onto a single line with other signals, while only the highest speed, highest bandwidth signals communicate directly over space-consuming parallel paths. The technical advantages of a NoC interconnect improve SoC economics.

Wire savings result from the use of a serialized packet transport. The architect can control the number of wires in a master-to-slave link based on connection latency requirements, using more wires for low latency connections such as CPU to DDR traces, a medium number of

wires for high bandwidth connections and a minimum number of wires for connections to I/O peripheral blocks where latency is not a major consideration.

Although NoC interconnects are sold for a fee, they are often lower-cost options than using alternative interconnect methods, even if these other methods have little or no initial cost.

A Real Example

One NoC provider, Arteris, has commercialized network on chip technology and has refined the technology into a mature product, called “FlexNoC,” through its work with 20 leading semiconductor companies including Texas Instruments, Qualcomm, and Samsung.

Arteris has provided Objective Analysis findings based on over 50 SoC designs of varying complexities, allowing us to look over a broad mix of several key variables:

- Die size
- Timing convergence
- Routing congestion
- IP flexibility savings
- Power utilization improvements
- Performance and bandwidth headroom expansion.

NoC Technology Benefits

In actual customer use, Arteris found that chips designed using FlexNoC require roughly half as many global interconnect wires and 30% fewer gates for complex SoC interconnects. (See table.)

The gate count reduction comes from an elegant NoC transport implementation (with relatively small amounts of buffer-

ing in the NoC transport layer) and from the elimination of bus bridges, multiplexers and other extraneous logic.

This 50% reduction in wires and a 30% reduction in gate count will shrink the SoC’s die size by a varying amount depending upon the design’s complexity and whether or not the chip is pad-limited. Typically an application processor SoC would see a die savings of 2-5%. On average this would translate to savings of 10 cents per chip, which in a manufacturing run of ten million chips would result in production savings of \$1 million.

Typical Design Cycle

A NoC approach simplifies signal routing compared to standard interface methods. Instead of routing signals through multiplexers and switches the architect decouples the IP blocks from the interconnect structures, adding protocol conversion blocks at the edge of the network.

In the case of the Arteris FlexNoC product, small Network Interface Units (NIUs) at the edge of the interconnect automatically translate common IP protocols into Arteris’ proprietary NoC transport format. NIUs have been designed to support all ARM protocols including ARM AMBA4 ACE, AXI, AXI Lite, AHB and APB, all versions of OCP, Tensilica PIF, BVCI and a number of proprietary formats.

Once the SoC architecture and topology are defined, the software automatically generates a gate level representation of the actual NoC, then automatically runs verification coverage tests and generates test benches. The automatic generation of the gate-level representation shaves days or even weeks from the SoC design cycle, and the verification and test fea-

Actual User Wire and Gate Savings			
	User A	User B	User C
Wire Reduction	35%	>50%	>50%
Gate Reduction	30%	17%	29%
Freq Obtained	> 400 MHz	469MHz	438 MHz

ture accelerates SoC verification by weeks or even months.

Cutting Project Costs

The RTL and layout development budgets for a typical SoC design project will cost around \$20 to \$30M. A NoC approach can help cut this cost by shortening development cycles, minimizing or eliminating routing congestion, and simplifying timing closure. This reduces the number of front-end and back-end engineering iterations required to meet SoC requirements.

Since NoCs use point-to-point connections rather than the more typical mixture of multiple fan-out meshes, it is simpler to add or remove IP blocks. While a more standard hybrid bus will contain a centralized crossbar and numerous wires that create difficult-to-route congestion points, an NoC's distributed communication architecture minimizes routing congestion points to simplify the chip's layout effort. This can cut weeks or even months from an SoC design project schedule as is illustrated in Figure 1.

Certain advanced tool sets offer the SoC architect design specification and exploration capabilities powerful enough to allow a large chip's complexity to be mastered in less time than can be done using other high-level interconnect design tools. Arteris tells us that their FlexArtist tool set helps the architect to optimize a design by modeling data traffic bottlenecks, congestion points, and performance issues. This

toolset generates multiple levels of SystemC models for the interconnect IP which can be used to get a more complete system level model of the SoC.

The NoC configuration tool allows fine-grained pipeline insertion to accelerate the process of timing closure. This can also subtract weeks from a project schedule and minimize the risk of an SoC missing its timing specification. Easier layout and fewer iterations not only speed project schedules and shrink costs, but they also can reduce project risk.

More Design Flexibility

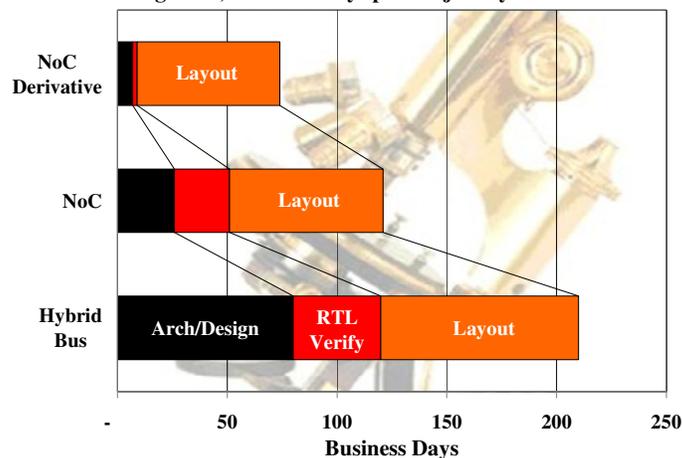
For many designers the use of a NoC results in more flexibility to optimize their SoCs. Prior to using a NoC approach these designers only had time to change their interconnect about 5-6 times per project. Today these same designers are able to refine their interconnect

9-10 times per project This gives them greater flexibility to accommodate changes coming from both engineering and customers while also providing more time for system-level power and performance optimization.

Improved Time to Market

Because the SoC's interconnect IP is available earlier, emulation can also start earlier, accelerating system verification and software development. Certain observability features such as Arteris' statistics collector allow the interconnect to be used to observe the inner workings of

Figure 1.) Business Days per Project By Phase



the SoC. This helps accelerate debug and performance optimization on early silicon.

In some cases these capabilities result in a time to market improvement of 20-50%, potentially shrinking an 18-month design cycle by 3½ - 9 months. To understand this in fiscal terms, assume that a project uses a design team of 60 hardware designers each with a fully loaded annual cost of \$250K. A 3-month design cycle reduction would yield savings of 180 man months or 15 man years translating into savings of \$3.5M per project.

Better Performing SoCs

The NoC approach doesn't only simplify layout – this approach has side benefits like power savings and improved performance. Lower power translates into longer battery life and a better user experience, leading to stronger sales.

To add to these power savings Arteris has added power conservation support to their FlexNoC product through multi-level clock gating, support for frequency, voltage and power domains, power management and other features. In one design, a 500,000-gate Arteris NoC in a 65nm LP process consumed 0.7mW of idle power compared to a hybrid bus which con-

sumed approximately 7mW under the same conditions. This is a 10x improvement in power consumption and standby battery time.

Cumulative Savings Over Multiple Projects	
	Measure
Annual Unit Production	100,000,000
Die Area Savings per Die	2%
Total Die Area Cost Savings	\$10,000,000
Projects per Year	6
Total Project Savings (8% total)	\$14,400,000
Sales Increase (5%)	5%
SoC Unit Shipment Increase	5,000,000
SoC Revenue Increase	\$25,000,000
Total Value Increase	\$49,400,000

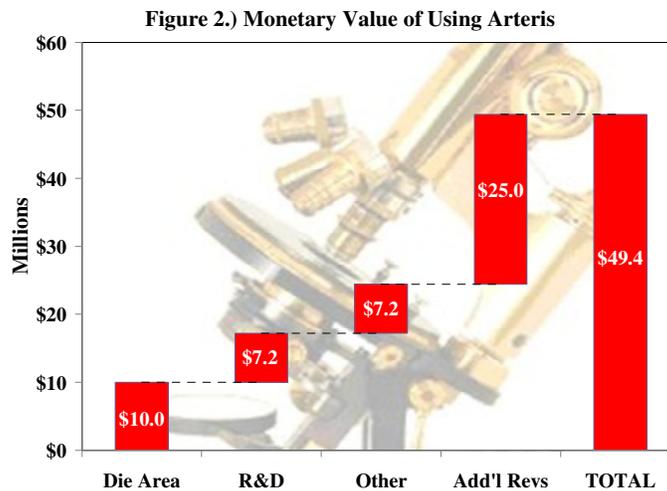
Higher performing SoCs can lead to improved revenues. Current NoC technology supports interconnect frequencies of up to 1GHz in a 40nm process, facilitating its use in high-performance designs. By comparison, conventional hybrid bus and cross-

bar performance tops out at around 300 to 400 MHz.

Some NoC approaches include a multi-port memory scheduler, like Arteris' FlexMem, which supports architectures that require optimized memory bandwidth and latency. This system connects to Synopsys and Cadence-Denali memory controllers and supports link widths ranging from 8-256 bits, delivering over one terabyte per second of bandwidth.

Higher Margins & Sales

Let's look at an example of how a NoC approach could impact the success of a typical SoC.



We will assume a volume chip maker shipping 100M SoCs per year based on six design projects per year. This scenario, shown in the table and Figure 2, is an average based on

Arteris' customers' production shipping experiences over five years.

In this example an SoC maker adopting the NoC interconnect IP platform would save \$49.4M per year based on a shipping volume of 100 million chips and 6 projects. The benefit of saving money by choosing a free software solution pales in comparison to the economic returns of the savings in this table.

What does this mean to a company's stockholders? If we assume that a SoC manufacturer has 100 million shares outstanding – adoption of an NoC interconnect platform could add about 4 cents per share to the annual earnings of this company.

It is clear that NoC interconnect technology could prove to be the most significant SoC design innovation of recent times when considered from the perspective of return on investment.

Jim Handy, April 2011