

# Contrasting Topologies for Regular Interconnection Networks under the Constraints of Nanoscale Silicon Technology

Daniele Ludovici, Georgi N. Gaydadjiev  
Computer Engineering Lab., TUDelft  
2628 CD Delft, The Netherlands

Francisco Gilabert,  
Maria E. Gomez  
GAP, University of Valencia

Davide Bertozzi  
ENDIF, University of Ferrara  
44100 Ferrara, Italy

## ABSTRACT

Nowadays, system designers have adopted Networks-on-Chip as communication infrastructure of general-purpose tile-based Multi-Processor System-on-Chip (MPSoC). Such decision implies that a certain topology has to be selected to efficiently interconnect many cores on the chip. To ease such a choice, the networking literature offers a plethora of works about topology analysis and characterization for the off-chip domain. However, theoretical parameters and many intuitive assumptions of such off-chip networks do not necessarily hold when a topology is laid out on a 2D silicon surface. This is due to the distinctive features of silicon technology design pitfalls. This work is a first milestone to bridge this gap, in fact, we propose a comprehensive analysis framework to assess  $k$ -ary  $n$ -mesh and  $C$ -mesh topologies at different level of abstractions, from system to layout level, while capturing implications of system and layout parameters across design hierarchy. When a certain topology proves to be slow due to long links crossing the chip, pipeline stages have been inserted to cope with such slow-down. Furthermore, costs of such speed-up technique have been evaluated to draw a comprehensive performance/area figure.

## Categories and Subject Descriptors

B.7.1 [Hardware]: Integrated Circuits—VLSI

## General Terms

Design, Performance

## Keywords

System-on-Chip integration, Network topologies, Link design techniques, Network-on-Chip

The gap between the constraints driving the design of on-chip vs. off-chip interconnection networks (and hence the gap between the final network architectures selected for use in each domain) is increasingly widening even more as an effect of the relentless pace of technology scaling to the nanoscale regime. New physical effects come into play and may either degrade performance/power in an unpredictable way or even affect feasibility of the design at hand or of specific architecture design techniques.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NoArc '10, December 4, 2010, Atlanta, Georgia, USA  
Copyright 2010 ACM ...\$10.00.

Examples concern the large buffering cost associated with techniques borrowed from the off-chip domain (e.g., for congestion management strategies or for deadlock-free and multicast friendly switching mechanisms) which are not affordable in the on-chip domain. Moreover, the interconnect reverse scaling is making designs for on-chip integration increasingly interconnect-dominated, due to the delay associated with the shrinking cross-section area of on-chip interconnects. This effect becomes increasingly severe at each technology node and tends to widen the gap between post-synthesis and post-place&route performance figures and even to move critical path delays from logic blocks to large global wires.

Selection of the topology connectivity pattern in the early stages of network design is a decision which is extremely sensitive to both the effects illustrated above. In fact, topologies for on-chip networks must match the 2D silicon surface, while off-chip realizations are dictated by board/rack organization. The 2D mapping constraint raises implementation issues such as wire crossings, wires of uneven length or the decrease of switch operating frequency with the number of I/O ports. As an ultimate consequence, topologies borrowed from off-chip networks should be reassessed in the on-chip environment and validated against the design pitfalls in this domain.

This is the motivation that lies at the core of this paper. We are aware that many regular topologies feature better abstract properties (e.g., diameter, bisection bandwidth) than a 2D mesh, however their implementation in an on-chip setting is very challenging. The objective of this paper is to quantify to which extent their inherently better abstract properties is impacted by the degradation effects of the physical synthesis on nanoscale silicon technologies. Proving whether these topologies are still efficient (or even feasible) after the physical degradation mechanisms is a non-intuitive task. This paper takes on this challenge. Previous work in the open literature features frameworks able to evaluate network topologies only from a pure theoretical viewpoint thus neglecting all the physical effects of nanoscale technologies. On the other hand, other works focused on the physical modeling of interconnection networks but only limited to small scale systems mainly due to the unaffordable time and memory requirements for the synthesis of such systems.

Therefore, our contribution consists of:

- an area and network critical path modeling framework able to accurately analyze performance of  $k$ -ary  $n$ -mesh and  $C$ -mesh topologies with layout awareness. Our proposed methodology scales easily to large size systems as only a few sub-systems of the whole network need to be analyzed to draw comprehensive area and performance figures.

When accounting for layout effects, conclusions drawn from high-level theoretical analysis can be highly misleading. Moreover,  $k$ -ary  $n$ -mesh and  $C$ -mesh topologies suffer of a considerable slow down when laid out on silicon. This is mainly due to their long links which represent the speed bottleneck of the whole network. To tackle this problem, pipeline stages are typically implemented in links of the top dimensions. To the best of our knowledge, all the previously pub-

lished analysis frameworks do not take into account the implications of utilizing such technique from an area/timing viewpoint with layout awareness.

Therefore, the second contribution is:

- the enhancement of our modeling framework with the capability of accurately capture the impact of using link pipelining from both the area and timing point of view accounting for physical effects. Interestingly, when considering also such layout implications of link pipelining utilization, some topologies previously considered low speed turn out to be competitive.

Last contribution can be summarized as follows:

- our previous work only considered systems with  $k$ -ary  $n$ -mesh where cores and network speed was constrained by an integer divider thus limiting the overall performance of the system. In this work, we extended our analysis to systems implemented as globally asynchronous locally synchronous (GALS) system where cores and network speed ratio can be any. Interestingly, the adoption of a GALS approach has considerable consequences on the performance/area figures of various topologies that were not competitive at all in the previously investigated scenarios. In order to achieve this objective, our transaction level simulator has been enhanced with dual-clock FIFO interfaces for cores and network frequency decoupling.

The remainder of this paper is organized as follows. Section 1 reviews previous work regarding topology mapping for NoC systems. Section 3 describes the modeling methodology utilized to characterize the topologies under analysis. Physical layout results of such section are utilized in Section 4 to carry out a system-level exploration with layout awareness. Finally, conclusions are drawn in Section 5.

## 1. RELATED WORK

Although it has been widely used across a number of Network-on-Chip tile-based embedded and high performance microprocessors designs [9, 10], the 2D-mesh NoC topology features well-known drawbacks in the communication latency scalability and concentration of the traffic in the center of the network[6]. This has motivated works in the open literature that come up with optimized NoC topologies while keeping regularity properties as much as possible. A novel interconnect topology called *spidergon* was proposed in [21], where each core is connected to the clockwise, counterclockwise and diagonal node. A traditional wormhole-routed mesh augmented by a hierarchical ring interconnect for routing global traffic is illustrated in [11]. NOVA is a hybrid interconnect topology targeted at FPGA and is compared in [12] with star, torus and hypercube topologies. Gilabert et al. propose in [13] to use high-dimensional topologies, using different metal layers to reduce long link delay and trading-off dimensions with the number of cores per router. However, this is not backed by any physical synthesis run. The work in [6] proposes a concentrated mesh architecture with replicated subnetworks and express channels.

Topology exploration is an active research area due to the large scale of on-chip networks and to the feasibility challenges posed by nanoscale technologies [14, 16, 17]. Unfortunately, as technology scales to the nanometer regime, topology analysis and exploration needs to be performed with novel methodologies and tools that account for the effects of nanoscale physics, largely impacting final performance and even feasibility of many NoC topologies. A general guideline driving network-on-chip (NoC) design under severe technology constraints consists of silicon-aware decision-making at each hierarchical level [18]. This is likely to result in less design re-spins and in faster timing closure. In this direction, new tools are emerging that guide designers towards a subset of most suitable candidates for on-chip network designs while considering the complex trade-offs between applications, architectures and technologies [19, 20].

Our previous work in [2, 3, 5] presented silicon-aware topology analysis and comparison for networks with 16 nodes. In all these works, the exploration of the design space is performed through a transaction-level simulation environment that is able to back-annotate key parameters (frequency, latency, area) from the results of physical synthesis. When extending the analysis to larger 64-tile networks, the unaffordable time and memory requirements for the synthesis of such systems makes a comprehensive exploration based on post-layout figures unfeasible. This is the reason why our previous work was limited only to 16-tile systems. In order to extend the exploration to larger 64-tile networks, in this work we devise (i) a novel modeling methodology based on selective synthesis runs that is able to capture the key post-layout parameters of a large scale topology such as, maximum frequency and switch cell area. (ii) Moreover, our framework is able to capture the impact of link buffering and link pipelining from the timing and area cost viewpoint. (iii) Furthermore, by utilizing such physical parameters in our transaction-level simulator, we are able to perform a layout-aware system-level analysis. This way, overall area and performance figures can be drawn. Differently from our previous work in [2], the simulator has been enhanced with the implementation of dual-clock FIFO interfaces thus enabling the modeling of systems where cores and network are completely decoupled from the frequency viewpoint. Interestingly, achieved results may look counterintuitive at a first glance when compared with commonly known theoretical properties of the investigated topologies.

Next section will describe such abstract properties which will be later on put in discussion by the physical implementation part of this work.

## 2. HIGH-LEVEL TOPOLOGY EXPLORATION

In this section a high-level comparison of topology performance is provided. However, this analysis will only give the high-level perspective and is agnostic of physical implementation effects. Nonetheless, it may be used in the early stages of system design to select the subset of the most promising topology candidates.

We restrict our focus to large 64-tile systems. The number of cores attached to each switch has been limited to four as a higher number of connected cores would introduce serious performance and feasibility issues. In fact, the topology would have a very low bisection bandwidth. Moreover, the placement of cores around the switches would not be a trivial task since the length of the injection/ejection links would increase. This would significantly limit overall NoC performance [3].

Table 1 summarizes the values of the properties of all 64 cores configurations considered for each topology. The analysis includes two different configurations of the CMesh network. From a pure topology viewpoint, a CMesh can be seen as a classical 2-D mesh with express links, regardless of the number of cores attached to each switch. As the investigated systems sizes are quite large, several topology configurations are possible and need to be taken into account. The best solution for high traffic loads is represented by the 2-ary 6-mesh. Moreover, this topology has one of the lowest hop counts (6), thus making it well suited for latency sensitive systems and applications. However, it requires the highest amount of resources: 64 switches of degree 7 and 384 unidirectional links. On the other hand, from a low-latency viewpoint, the best solution is either the 2-ary 4-mesh or the 4-cmesh, which again are completely equivalent from a high-level view-point.

Overall, the best topology would be the 2-ary 6-mesh, as it provides four time more bisection bandwidth than the low-latency solutions, while requiring only two hops more (6 hops in the 2-ary 6-mesh versus 4 hops in both low-latency solutions). The only drawback of such topology lies in the high number of required resources. Finally, when system specifications do not require such a high bisection bandwidth, the 2-ary 5-mesh solution becomes a good trade-off that provides twice the bisection bandwidth of the low-latency solutions (while increasing the number of hops by one). Clearly, by

Topology	Switches	Cores/ switch	Max. degree	Unidir. links	Bisection bandwidth	Hop count	Connect.
8-ary 2-mesh	64	1	5	224	16	14	2
4-ary 3-mesh	64	1	7	288	32	9	3
4-ary 2-mesh	16	4	8	48	8	6	2
2-ary 6-mesh	64	1	7	384	64	6	6
2-ary 5-mesh	32	2	7	160	32	5	5
2-ary 4-mesh	16	4	8	64	16	4	4
8-cmesh	64	1	5	256	32	8	4
4-cmesh	16	4	8	64	16	4	4

Table 1: High level parameters of topologies with 64-tile.

blindly relying on this table and upon the underlying theoretical analysis, a designer would easily discard the 2D mesh (8-ary 2-mesh) as candidate topology.

The remainder of this paper will prove that theoretical properties of such topologies are put in discussions when layout considerations are taken into account and may even lead to counterintuitive final results.

Next section will present the characterization methodology that is at the core of our modeling framework. Such methodology will be used to extrapolate key physical parameters to be back-annotated in the transaction-level simulator, thus enabling a layout-aware system-level exploration.

### 3. PHYSICAL MODELING FRAMEWORK

The *xpipesLite* [8] switch was used as the basic building block to construct the 64-tile topologies under test. However, exploring the design space of topologies with such a large number of cores with full physical synthesis proved impractical due to synthesis time and memory capacity requirements. Therefore, next section will present a way to cut down on the number of physical synthesis tests while still characterizing the full topology with high accuracy.

All the analyzed topologies of this work have been laid out by means of a backend synthesis flow leveraging industrial tools. The topology specification is fed to the *xpipescompiler* tool[22], resulting in the generation of self-contained SystemC code for RTL-equivalent simulation and for synthesis. Synopsys Physical Compiler is used for placement-aware logic synthesis. The technology library is a low-power low-V<sub>th</sub> 65nm STMicroelectronics library available through the CMP project [7]. Placement and routing have been performed with Cadence SoC Encounter.

#### 3.1 Characterization Methodology

In order to accurately characterize the switch and link buffering cell area of the topology under analysis, we propose to utilize the methodology depicted in Figure 1. In fact, as already reported in our previous work [2, 3, 5], the performance bottleneck of a topology lies in its longest switch-to-switch communication channel. Aware of this, from a high-level topology specification we build a sub-system composed of two communicating switches at the maximum possible distance in the topology. This way, the critical link delay can be extracted. Such delay (which is the critical path delay of the network) is then used as the *target delay* to re-synthesize, place and route all the possible switch-to-switch sub-systems for each different inter-switch link length. The reason for this is that our goal is to accurately capture the switch cell area at a certain distance and at a certain target speed. It is well known from logic synthesis theory that as the target speed is decreased, large area can be saved. In this direction, it would make no sense to synthesize switches for maximum performance when a long link limits overall network speed (unless decoupling techniques like link pipelining are used, as we will see later on). Please note that each switch of the built sub-system has been pre-characterized standalone with the input/output delay that is able to tolerate from its neighbor communicating block. These parameters were set in such a way that the communication link delay is optimized as much as possible thus shortening the critical path of the switch-to-switch modeling architecture.

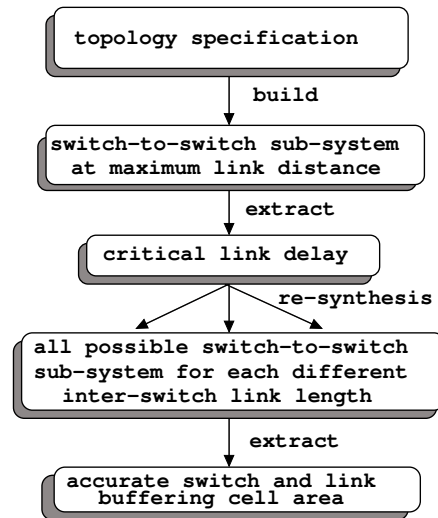


Figure 1: Characterization methodology flow.

With this methodology, only a few selected synthesis runs for each topology need to be performed to characterize its delay and area as a whole. The approximation lies in the availability of enough routing channels for regular routing of NoC links and in the balance preservation of relative wire delays in links that undergo bending in the actual layout.

Moreover, with this method we are also able to capture the link buffering cost, in fact, by leveraging the report of the utilized physical synthesis tool, we are able to trace the inferred buffers of the switch-to-switch channel. In order to be as accurate as possible when characterizing a topology, two communicating ports of both switches in our subsystem were left unconnected. They are the ports connecting to the processing cores, which are typically placed close to their switch and therefore feature minimum capacitive load. Should we fail to model this (even by simply leaving an output port unconnected), the input and output buffer of the switch would be incorrectly sized by the synthesis tools by using larger driving strengths than actually needed for the switch-to-core links.

A further step of our work is the estimation of the number of required pipeline stages for each link to speed up a topology. For this purpose, such retiming stages are instantiated along the communication link thus breaking the switch-to-switch critical path. By incrementing the number of pipeline stages, we were able to achieve timing closure bringing back the critical path to the second link dimension. In fact, as mentioned later, in order to limit area overhead, our pipeline stage insertion criteria consisted of adding such stages only from the third link dimension onwards.

The next section starts by commenting physical synthesis results achieved for 64-tile topologies without link pipelining. Consequently, the analysis is shifted to pipelined systems. Section 4 will utilize the obtained physical results to carry out a system-level exploration with layout-awareness.

### 3.2 64-tile topologies

As reported in Table 2, the range of possible switch radix per topology spans from a reasonable 6 to a large 12 that is even more difficult to place and route as a stand-alone block without DRC (design rules check) violations [4]. Post-synthesis frequency results reflect the increasing trend with the switch radix, in agreement with the analysis of [4]. After placement and routing, the effect of the long links comes noticeably into play. Most of the topologies suffer from long switch-to-switch channels that need to be routed along the chip. For the sake of the analysis, only the longest link per topology is reported in the 5th column. By comparing such column with the 4th one, it is possible to recognize a clear correlation between the increasing link length and the decreasing operating speed of the topology under analysis. In fact, the critical role of the interconnect is a major factor limiting the performance of a topology. It should also be observed that also some logic gates end up in series to the critical links close to the far-ends. They are associated with flow control management and further contribute to the critical path delay. The trend above is even more apparent when we consider larger topologies. In fact, only topologies with short links (e.g., 8-ary 2-mesh and 4-ary 2-mesh) can work at a reasonable frequency for realistic application scenarios.

TOPOLOGY	Radix	post-synthesis frequency	post-P&R frequency	longest link
8-ary 2-mesh	6	1.08GHz	890MHz	1.5mm
8-cmesh	6	1.08GHz	250MHz	6.75mm
4-ary 3-mesh	8	950Mhz	220MHz	6.9mm
2-ary 6-mesh	8	950Mhz	220MHz	6.9mm
2-ary 5-mesh	9	810MHz	230MHz	6.96mm
4-ary 2-mesh	12	720MHz	530MHz	3.0mm
2-ary 4-mesh	12	720MHz	260MHz	6.4mm
4-cmesh	12	720MHz	260MHz	6.4mm

Table 2: Post-place&route results of the 64-tile topologies under test.

From the area viewpoint (see Figure 2), it is interesting to note that this result is influenced by the combination of many parameters such as: number of switches in the topology, their radix and consequently their final working frequency. In fact, as explained above, in order to be accurate, all representative switches in every topology have been re-synthesized at the final working speed of the whole network.




Figure 2: Normalized area for 64-tile topologies.

As an example, let us consider a very slow topology like 2-ary 6-mesh that features a larger area footprint with respect to the 8-ary 2-mesh. Such a network is operating at a frequency much slower than the 8-ary 2-mesh, but since it has an equal number of switches (64) with a higher radix (8 vs. 4, 5 or 6), the overall area figures plays in favor of the 8-ary 2-mesh with a 10% saving.

Another interesting result concerns the 4-ary 2-mesh. This topology has a relatively short link (3mm), thus it does

not suffer from a large speed degradation after place-and-route. As reported in Table 2, this topology is the only one (along with 8-ary 2-mesh) to have a final working speed above 500MHz. Interestingly, the area footprint of such topology has a 20% saving with respect to the 8-ary 2-mesh as it only has 16 switches. Although their radix is 10, 11 and 12, their final working speed along with the number of their instances results to be more area effective than the 8-ary 2-mesh counterpart.

The overall conclusion is that most of the topologies are not competitive with the 8-ary 2-mesh because of their long links that influence the final working speed. A natural way to tackle this problem is to implement link pipelining on such long links but the policy of insertion has to be carefully engineered. In fact, the studied 64-tile topologies feature a high number of long links that could rapidly bring the area cost to an unaffordable budget for a system-on-chip.

### 3.3 Pipeline stage insertion for 64-tile systems

In order to cope with the high speed degradation of most topologies analyzed in the previous section, pipeline stages need to be inserted especially in the top dimensions. By adding pipeline stages, it is possible to partially (if not completely) recover the initial operating frequency of the basic switch block. The criteria that has been adopted for the insertion of pipeline stages is to use them only from the third link dimension onwards. Therefore, topologies such as 8-ary 2-mesh and 4-ary 2-mesh have not been modified. Table 3 collects the results of this experiment. As clearly reported in the 3rd and 4th column, the insertion of pipeline stages is a very effective way to reduce post-place and route frequency degradation. Column 5 reports the number of pipeline stages inferred in each link dimension whereas the 6th column points out the number of links of each topology. The area weight comes from the combination of these two factors and it is reported in the 7th column. Total cell area of the topologies along with the contribution of such retiming stages insertion is reported in Figure 3.

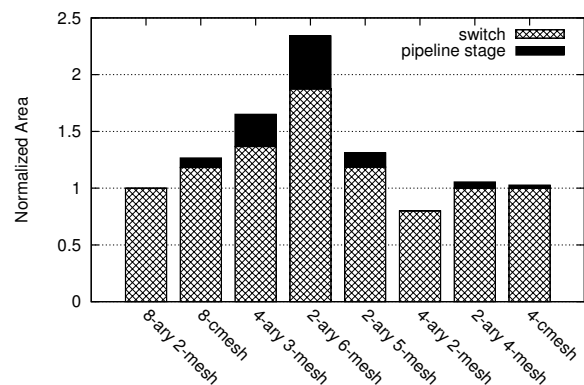


Figure 3: Normalized area for 64-tile topologies with pipeline stages.

Please note that the number of pipeline stages per link depends on the maximum achievable frequency (dictated by the maximum switch radix) along with the link length which is an intrinsic characteristic of each topology. As reported in Figure 3, the 2-ary 6-mesh is the most area greedy topology because it has the highest number of switches (64) and they were placed and routed at the high frequency of 855MHz. Moreover, this topology features 192 links with up to 5 pipeline stages on the longest interconnection channel. The key take away is that, for each topology, there is a different price to pay to restore the possible working frequency allowed by the elementary switch block. For this reason, Section 4 will introduce the throughput/area metric (or area efficiency) that provides a fair assessment of the cost of the achievable bandwidth in each topology (see Figure 6(b)).

To conclude the physical implementation part, it is inter-

topology	radix	post-synthesis frequency	post-P&R frequency	# of pipe-stage per dimension	num. links	tot. pipe-stage area (um <sup>2</sup> )	tot switch area (um <sup>2</sup> )	impact of pipe-stage insertion on tot. switch area
8-ary 2-mesh	6	1.08GHz	893MHz	0	112	0	2327712.8	0%
8-cmesh	6	1.08GHz	893MHz	express link $\Rightarrow$ 4	128	193425.9	2752108.8	7.03%
4-ary 3-mesh	8	950MHz	855MHz	dim.3 $\Rightarrow$ 4	144	660216.3	3182953.2	20.74%
2-ary 6-mesh	8	950MHz	855MHz	dim.3,4 $\Rightarrow$ 1, dim.5,6 $\Rightarrow$ 5	192	1087918.1	4362092.8	24.94%
2-ary 5-mesh	9	810MHz	562MHz	dim.3 $\Rightarrow$ 1, dim.4,5 $\Rightarrow$ 3	80	293081.6	2758480.4	10.62%
4-ary 2-mesh	12	720MHz	532MHz	0	24	0	1860718.3	0%
2-ary 4-mesh	12	720MHz	532MHz	dim.3,4 $\Rightarrow$ 3	32	125574.7	2328426.4	5.39%
4-cmesh	12	720MHz	532MHz	express link $\Rightarrow$ 3	32	62787.4	2328426.4	2.69%

Table 3: Post-place&route results of 64-tile topologies with pipeline stage insertion.




Figure 4: 64-tile topologies area overhead for pipeline stage insertion.

esting to observe the result depicted in Figure 4. For each topology, this graph reports area results before and after inserting pipeline stages. Interestingly, the substantial cell area increment in all cases (except topologies where pipeline stages were not inserted) comes from a twofold contribution: pipeline stages insertion (as discussed so far) and the restored higher frequency allowed by such insertion. In fact, the largest contribution in terms of area comes from the switch cell area devoted to achieve the new working frequency of the switch block. This relevant effect is typically overlooked by vast majority of topology exploration frameworks in the open literature.

## 4. SYSTEM-LEVEL EXPLORATION

This section will discuss the gap between high-level and realistic performance predictions, by comparing the former with layout-aware ones.

### 4.1 Experimental setup

In order to obtain accurate performance estimations, this work used the simulator presented in [3], which is cycle accurate with the assumed RTL architecture. The clock domain crossing mechanism implemented in the original simulator was ratio based. In particular, tile frequency was forced to be an integer divider of NoC frequency. However, as discussed in the previous section, when link pipelining is not considered, some topologies present severe critical path degradations. These low frequency topologies cannot remain competitive with ratio-based clock domain crossing mechanism, as it will have a direct impact over the speed of the processing cores (see [2]). In order to allow a fair performance comparison between topologies with extremely different operating frequencies, the simulator was augmented with the implementation of a dual-clock FIFO interface thus enhancing the whole modeling framework. Last but not least, a dual-clock FIFO allows frequency decoupling between a core and the network node it is connected to. In all the cases, tiles are assumed to work at a frequency of 750 MHz.

### 4.2 Experimental results

Figure 5(a) depicts accepted traffic vs. average message latency for a uniform distribution of message destination for

different topologies when considering high-level estimations. Obtained results reflect the conclusions drawn in Section 2 where the 2-ary 6-mesh proved to be the best solution when neglecting layout implications. Figure 5(b) shows the same analysis where each topology works at the operating frequency reported in the previous section (Table 2, without pipeline stages). By comparing Figure 5(a) against Figure 5(b), when link pipelining is not considered, there is a misleading gap between the performance predictions of the high-level analysis and the layout-aware one. In fact, while the theoretical results reported in Figure 5(a) claim that several topologies outperform the 8-ary 2-mesh, this latter topology is proved to be the best solution in the layout-aware results of Figure 5(b). In fact, there is a direct correlation between the operating frequency and the achieved system-level performance: the lower the operating frequency, the higher the average latency and the lower the maximum achievable throughput, regardless of the results obtained in the high level analysis. In practice, poor matching with silicon technology completely offsets the better theoretical properties of the topologies. However, when the impact of wiring complexity over the critical path is alleviated by using link pipelining techniques, different conclusions can be drawn.

Figure 5(c) reports the same analysis results when each topology works at the operating frequency (see Table 3) enabled by the usage of link pipelining. In this case, there are three network topologies that clearly outperform the 8-ary 2-mesh: 2-ary 6-mesh, 2-ary 5-mesh and 4-ary 3-mesh.

Similar curves have been drawn for several traffic patterns for each topology. Those results are summarized in Figure 6(a). This figure shows the normalized maximum throughput of each topology with respect to the 8-ary 2-mesh solution. In this plot, a bar higher than 1 implies an improvement of the maximum throughput over the 8-ary 2-mesh solution. Interestingly, those results follow the same trend as discussed for the uniform traffic pattern. All non-pipelined solutions are clearly worse than the 8-ary 2-mesh, while pipelined solutions follow the same trend reported in the high level analysis: most of the solutions outperforms the 8-ary 2-mesh, with the 2-ary 6-mesh being the best solutions for all the traffic patterns. Although in this case the obtained performance is closer to the high-level estimations, link pipelining techniques may have a great impact over the implementation cost, thus requiring a new metric to assess the real effectiveness of link pipelining techniques.

In particular, we have considered the *area efficiency metric*, defined as *throughput/area*, which correlates the throughput improvement with the area cost that has been paid to achieve that. Results are shown in Figure 6(b), which depicts the area efficiency of each topology normalized with respect to that of the 8-ary 2-mesh. Results are reported with and without pipelining for several traffic patterns. In most of the cases, the area efficiency of both pipelined and non-pipelined solutions is clearly lower than the 8-ary 2-mesh solution. The key take away is that the performance improvements achieved by complex topologies with pipelined links are not cost-effective.

The only exception is when the traffic pattern favors topologies with a low hop count, as in the case of the perfect shuffle traffic. This characteristic, along with the fact that some topologies feature a low area cost, leads to a higher area efficiency with respect to the 8-ary 2-mesh.

## 5. CONCLUSIONS

In this work we presented a comprehensive analysis frame-




Figure 5: Performance of 64-tile systems with uniform traffic.




Figure 6: Normalized performance and area efficiency of 64-tile systems.

work to assess  $k$ -ary  $n$ -mesh and  $C$ -mesh topologies at different levels of abstraction, from system to layout level. Our framework leverages an accurate physical characterization methodology that allows to characterize various topologies from the area and timing viewpoint while pruning the implementation time as well as memory requirements. All the topologies have been evaluated at physical level and their key parameters have been back-annotated for use in a transaction-level simulator that performs layout-aware system-level exploration. The latter has been enhanced with dual-clock FIFO interfaces to allow fully decoupled working frequencies between cores and the NoC.

This paper demonstrated that is possible to evaluate large scale topologies with physical level accuracy while cutting down the analysis time by utilizing only a few selected synthesis and place&route runs. Furthermore, we proved that conclusions drawn by a pure high-level analysis of topology performance can be highly misleading if not enriched by the information provided by the physical synthesis. As an example, let us consider  $k$ -ary  $n$ -mesh, these are very difficult topologies to be realized without link pipelining and the implementation cost of using such technique is typically overlooked. To tackle this problem, our modeling framework has been devised in such a way that is able to capture the utilization of pipeline insertion from an area and timing viewpoint. Leveraging this capability, analyzed topologies, typically considered too slow, turn out to be cost-effective and may represent a valid alternative for a given implementation budget. Last but not least, our work has been extended to consider globally asynchronous locally synchronous (GALS) systems where the cores and network speed is fully decoupled. Utilizing such systems brings back momentum to topologies that were strongly limited by the constraint of using an integer clock divider between cores and network.

## 6. REFERENCES

- [1] L. Benini and G. De Micheli, "Networks on Chip: a New SoC Paradigm". IEEE Computer, 35(1):70-78, January 2002.
- [2] F. Gilibert, D. Ludovici, S. Medardoni, D. Bertozzi, L. Benini, G. N. Gaydadjiev, "Designing Regular Network-on-Chip Topologies under

Technology, Architecture and Software Constraints". Proc. of IEEE MuCoCos, Fukuoka, Japan, 2009.

- [3] F. Gilibert, S. Medardoni, D. Bertozzi, L. Benini, M. E. Gómez, P. López, J. Lópezo, "Exploring High-Dimensional Topologies for NoC Design Through an Integrated Analysis and Synthesis Framework". Proc. of International Symposium on Network-on-Chip, pp.107-116, 2008.
- [4] A. Pullini, F. Angiolini, S. Murali, D. Atenza, G. De Micheli, L. Benini, "Bringing NoCs to 65 nm". IEEE Micro Special Issue on Interconnects for Multi-Core Chips, 27(5):75-78, 2007.
- [5] D. Ludovici, F. Gilibert, S. Medardoni, C. Gómez Requena, M. E. Gómez, P. López, G. N. Gaydadjiev, D. Bertozzi, "Assessing Fat-tree Topologies for Regular Network-on-Chip Design under Nanoscale Technology Constraints". Proc. of DATE, pp.562-565, 2009.
- [6] J. Balfour and W. J. Dally, "Design Trade-offs for Tiled CMP On-chip Networks". Proc. of the 20th ICS, pp.187-198, New York, NY, USA, 2006.
- [7] Circuits Multi-Projects, Multi-Project Circuits: <http://cmp.imag.fr>
- [8] S. Stergiou, F. Angiolini, S. Carta, L. Rao, D. Bertozzi, G. De Micheli, "XPipes Lite: a Synthesis Oriented Design Library for Networks on Chips". Proc. of the Design Automation and Test in Europe (DATE), pp.1188-1193, 2005.
- [9] TILE64 Processor Family, online at [http://www.tilera.com/pdf/ProBrief\\_T11e64\\_Web.pdf](http://www.tilera.com/pdf/ProBrief_T11e64_Web.pdf)
- [10] S. Vangal et al., "An 80-Tile 1.28TFLOPS Network-on-Chip in 65nm CMOS". Proc. of ISSCC 2007, pp.98-589, 2007.
- [11] S. Bourduas, Z. Zilic, "Latency Reduction of Global Trajc in Wormhole-Routed Meshes Using Hierarchical Rings for Global Routing". Proc. of IEEE International Conference on Application-Specific Systems, Architectures and Processors, pp.302-307, 2007.
- [12] F. Martinez Vallina, N. Jachimiec, J. Sanjie, "NOVA Interconnect for Dynamically Reconfigurable NoC systems". Proc. of IEEE International Conference on Electro/Information Technology, 2007, pp.546-550, 2007.
- [13] F. Gilibert, M. E. Gomez, P. J. Lopez, "Performance Analysis of Multidimensional Topologies for NoC". ACACES 2007, poster session with proceedings at the Summer School.
- [14] M. Mirza-Aghatabar, S. Koohi, S. Hessabi, M. Pedram, "An Empirical Investigation of Mesh and Torus NoC Topologies Under Different Routing Algorithms and Trajc Models". Proc. of Euromicro Conference on Digital System Design Architectures, Methods and Tools, pp.19-26, 2007.
- [15] L. Bononi, N. Concer, M. Grammatikakis, M. Coppola, R. Locatelli, "NoC Topologies Exploration based on Mapping and Simulation Models". Proc. of Euromicro Conference on Digital System Design Architectures, pp.543-546, 2007.
- [16] H. Wang, L. S. Peh, S. Malik, "A Technology-Aware and Energy Oriented Topology Exploration for On-Chip Networks". Proc. of Design Automation and Test in Europe (DATE), pp.1238-1243, 2005.
- [17] M. Kreutz, C. Marcon, L. Carro, N. Calazans, A. Susin, "Energy and Latency Evaluation of NoC Topologies". Proc. of IEEE International Symposium on Circuits and Systems, pp.5866-5869 Vol. 6, 2005.
- [18] I. Hatimaz, S. Badel, N. Pazos, Y. Leblebici, S. Murali, D. Atenza, G. De Micheli, "Early Wire Characterization for Predictable Network-on-Chip Global Interconnects". Proc. of SLIP Conference, pp.57-64, 2007.
- [19] S. Murali, G. De Micheli, "SUNMAP: A Tool for Automatic Topology Selection and Generation for NoCs". Proc. of the Design Automation Conference (DAC), pp.914-914, 2004.
- [20] Soteriou, V., Eiseley, N., Wang, H., Li, B., Peh, L.S., "Polaris: a System-Level Roadmapping Toolchain for On-Chip Interconnection Networks". IEEE Trans. on VLSI 15(8), pp.855-868, 2007.
- [21] M. Coppola, R. Locatelli, G. Maruccia, L. Pieralisi, A. Scandurra, "Spidergon: a novel on-chip communication network". Proc. of International Symposium on System-on-Chip, pp.16-18, 2004.
- [22] A. Jalabert et al., "xpipesCompiler: a Tool for Instantiating Application Specific Networks on Chip". Proceedings of DATE, pp.884-889, 2004.