3.1 Publishable summary

The NANOC project aims at developing an innovative design platform for future network-on-chip (NoC) based multi-core systems. Strict component-oriented (and reuse) architectural design is targeted in the project, which enables faster design methods. The component-oriented method is out of reach of current design methods and tools due the introduction of the NoC concept in new and future designs and mostly because of new design constraints. Enhanced dynamism and flexibility in NoC composition is required to tackle the new requirements of future systems like virtualization, power management, thermal management, and application management. In addition, the continuous scaling of technology opens the door to new challenges like reliability and variability. Reliable systems will need to be designed out of unreliable components. NANOC provides design methods and prototype tools to address such challenges and embeds them into the final design platform taking the NoC as an additional component in the system, thus enabling the componentization approach. NANOC is built around a vertical design approach where system, circuit and processes are co-developed with silicon-aware decisions considered at every layer. A common exchange format (CEF) standard is co-developed in the project (and intends to span out of the project) to enable communication among all layers of the design platform.

The project builds around five technical WPs, each one with a particular and contributing goal to the final target of the project. The first two WPs deal with design methods to tackle the new challenges that systems are facing (manufacturing faults, process variation, run-time network management …). WP3 and WP4 deal with tools for the final platform at the front-end (logic level) and back-end (physical level), respectively. WP5 targets the final integration of both design methods and tools into the final design platform. While WP2 and WP5 are scheduled for the second and third year of the project, the first year (under current review) is focused on WP1, WP3, and WP4, addressing three main aspects:

- Development of design methods for supporting static irregularities in NoCs (WP1)
- Development of frontend NoC topology synthesis tools (WP3)
- Development of backend physical level design tools for NoC composability (WP4)

Design methods, developed in WP1, have been focused on tolerating irregularities in the NoC due to manufacturing defects. As technology scales down, the probability of manufacturing defects increases, affecting yield and thus manufacturing costs. There is an urgent need to tolerate such defects and, by that, increase the yield. In the first year, design methods have been designed to tolerate such defects in the NoC. In particular, full coverage of manufacturing defects has been achieved by the different LBDR and FDOR mechanism versions designed in WP1. Both unicast (one to one traffic) and multicast/broadcast (one to many/all traffic) have been supported in the presence of any failure pattern in an initial regular mesh network. All these mechanisms have been conceived taking into account the area and power design constraints present in NoCs. An extension to FDOR to support dynamic irregularities (e.g., those caused by temporary shut down of network components) has been designed. This method, iFDOR, will be one of the starting points for the support of dynamic topology changes (to be developed in the second year in WP2).

The effect of manufacturing process variability has a fundamental impact on NoC performance. It affects the maximum frequency that different parts of the system interconnect can achieve. An investigation regarding this effect has been performed for several existing and future process nodes. A new variability modelling framework has been

elaborated, which is well suited for on-chip interconnection networks and for the system-level assessment of variability effects based on their characterization at a lower abstraction layer (physical links and logic gate delays). This provides valuable understanding of the impact of variability looking forward to advanced technology nodes, and helps to design architectural level techniques that reduce the uncertainty of post-silicon performance. Indeed, two strategies (channel width reduction and space multiplexed channels) have been considered and partially developed to address this issue. The impact of process variation on routing in NoCs has also been covered. Indeed, a variability-aware process mapping algorithm has been designed to improve the overall chip multiprocessor (CMP) performance.

In WP3, tools for the assisted or even automated design of NoCs are to be developed during the first and second years of the project. Significant advances have been achieved and the tools are available or their development is underway. Different challenges are being addressed by these tools. For example, the INOCS NoC synthesis toolchain has been extended to support vertically stacked ("3D") chip designs, instantiating in this case NoCs with vertical communication links. Vertical connectivity is tuned to optimize multiple metrics, such as performance, power consumption, chip area and chip yield. The tools have been demonstrated to support four or more layers, as demanded by projections for 2015 designs. Routing design methods developed in the project (LBDR and FDOR in WP1) have also been evaluated for the use in 3D stacked chips as well.

In order to approach the goal of transparent IP reuse in SoC design, two of the most significant challenges are the partitioning of designs into ever increasing numbers of clock domains and the variety of interface protocols used by today's IP cores. Both trends negatively impact the chip design and verification times. In WP3, the NANOC partners have started to tackle both issues. In terms of multiple clock domain support, the work is well underway with tool support for three new logic blocks: a serial/parallel converter, a mesochronous cell, and a dual-clock packet FIFO. To enable plug-and-play IP core integration, investigations are ongoing to support multiple protocol interfaces (AMBA AHB, AMBA AXI, OCP) and data widths (16 to 64 bits), thus enabling the design of heterogeneous systems. The research is aimed at designing NoC blocks which, at a low implementation cost, can seamlessly interface IP blocks with different port protocols with the NoC and thus with each other, negating the need for dedicated bridges and simplifying verification.

Simulation is a key tool for the platform, as it allows a preliminary evaluation of the different design methods under many different situations (traffic changes, topology changes). In the first year of the project a tool chain has been developed to perform automatic design space exploration (DSE). With the DSE tool, the huge design space is quickly and accurately cropped thus reducing tremendously design time.

A CEF (Communication Exchange Format) file specification which extensively encompasses the requirements and outcomes of a NoC design flow has been defined and agreed by all the NANOC partners. The 1.0 specification includes the ability to describe numerous attributes, requirements and outputs of the chip interconnect design flow, and it is designed to become the universal file format for exchanges among NANOC project partners and beyond. The main capabilities of the format allow for the description of:

- Frequency and voltage domains in which the chip is partitioned

- IP cores belonging to the design, and to be interconnected
- Communication requirements, including communicating pairs of cores and required performance
- Architecture of the chip interconnect
- Communication routing
- Physical properties of the design, e.g. basic floorplan, wire length
- Additionally, the file format is designed to be extensible with custom tags, so that each user of the format can customize it with additional specifications.

In WP4, physical level issues of the NANOC design platform are addressed. Tools and methods are developed to address such issues. This WP spans the first two years of the project, thus being right now in the middle with no major deviations.

Improving the hold-time margins of circuits, in order to improve robustness to process variations, is a highly valuable feature in advanced IC design flows for large-scale designs. Currently, hold-time fixing is done as a post-implementation ECO and as such it is highly sensitive to a number of physical-level design issues. An optimization technology is being developed which provides, for the first time ever, the flexibility of applying early stage netlist modification, in order to improve hold-time robustness.

Also, the ability to analyze dynamic power grid performance at an early design stage, and with a fast turn-around-time, is essential for maintaining physical design convergence in complex NoC-based SoC designs. This is not possible with existing analysis tools, as these are focused on use at the sign-off level. Thus, system-level power grid modeling for SoCs is being developed, which will allow designers to perform extremely fast, early-stage feasibility studies of SoC floorplans, assessing dynamic voltage drop in a matter of seconds rather than minutes or even hours, as is possible with existing tools.

A substantial part of the work aims at constraining the non-determinism of post-place&route link statistics and at coming up with link-level area/performance trend, as the synthesis constraints and topology requirements are varied, has also been started. In the second year of the project, such trend will be characterized for two major link inference techniques: repeater insertion and link pipelining.

Furthermore, the role of non-routable obstructions will be assessed on the derived results, thus accounting for the more or less limited routing channels available in actual chip layout for network link routing depending on the physical space budget. Moreover, the area/performance/power implications of utilizing such techniques will be evaluated not just for the link in isolation, but for the network topology as a whole.

As a summary of achievements for the first year of the project, we can list the following ones:

- Design and implementation of a static design method (LBDR) that achieves 100% coverage of manufacturing defects in system designs based on regular NoC topologies (deliverable 1.1).
- Support, with 100% coverage of manufacturing defects, for collective communication (bLBDR design method).
- An application-specific routing methodology capable of providing quality of service guarantees with minimal routing, which is deadlock free in terms of both

packet and protocol deadlocks. Design of a dynamic design method (iFDOR) to adapt the NoC to sudden irregularities (deliverables 1.1 and 1.5).

- Tools to design NoCs for vertically stacked ("3D") chips having four or more layers (deliverable 3.1).
- Support in 3D chips for LBDR design method (deliverable 3.1).
- Initial multiple DVFS domain support in tools.
- Design of a tool-chain for Design Space Exploration to highly reduce exploration time (deliverable 3.3).
- Development of the first version of the Communication Exchange Format (CEF) standard to allow all the NANOC tools to exchange data in a common standard format (deliverable 3.4).
- Development of a system-level power grid modeling tool for SoCs enabling dynamic power grid performance at an early design stage and with fast turn-around-time.
- Development of an accurate variability model that takes into account the main sources of process variation (deliverable 4.3).
- Prediction of what the network link contribution will be to total area figures depending on the constraints of the physical synthesis (deliverable 3.3).

All the work performed in the first year of the project is in line with the initial schedule and no major deviations and roadblocks have been found. Therefore, the partners expect the best achievements obtained during this year will be fully integrated in the final design platform in the third year of the project.

European universities and companies are still competitive in research and development of complex systems-on-a-chip. Especially the use of ESL tools is accepted more openly in European companies compared to the US. The design tools of the NaNoC project may help to keep this edge throughout the next years. This improved productivity is vital to be able to achieve first-time-right SoC designs with more and more processors and other components communicating over more complex interconnect solutions facing more variations caused by shrinking process geometries. Such tools help to do such designs cost-effectively in European countries despite of their higher cost of labour. This has a double effect on economy and society: first, the directly associated jobs can be kept in Europe, and second, a lot of related jobs will remain in Europe as well in industries that have a growing part of electronics contributing to their success (e.g. automotive, medical equipment, industrial control).

With respect to the dissemination effort of the consortium during the first year, different publications have been published in highly-rated and peer-reviewed conferences (e.g. DATE conference, NOCS conference). A book has been edited with five chapters involving the NaNoC research, thus potentially reaching a larger audience. Several press publications and press interviews have been performed by UPV and UNIFE. Poster presentations of the project have been held also in important venues like the HiPEAC network of excellence. The dissemination effort can be considered as high in the present year.

The project consortium maintains a website (http://www.nanoc-project.eu) with the most recent achievements and the list of public deliverables. All the NANOC-related info can be found at the website.