INTEGRATION, the VLSI journal 52 (2016) 102–121

Contents lists available at ScienceDirect

INTEGRATION, the VLSI journal

journal homepage: www.elsevier.com/locate/vlsi

# Design automation for application-specific on-chip interconnects: A survey

Alessandro Cilardo [1], Edoardo Fusella

*Department of Electrical Engineering and Information Technologies, University of Naples Federico II via Claudio 21, 80125 Napoli, Italy*

ARTICLE INFO

*Article history:*
Received 2 March 2015
Received in revised form
26 July 2015
Accepted 27 July 2015
Available online 28 August 2015

*Keywords:*
On-chip interconnects
Systems-on-chip
Embedded systems
Design automation
Wires

ABSTRACT

On-chip interconnects provide a vital facility for highly parallel MultiProcessor Systems-on-Chip, particularly in data-intensive applications, where the choice of the underlying communication architecture, tailored on the particular application requirements, is critical to the global performance. This survey focuses on the design automation of a broad class of communication architectures, here referred to as structured on-chip interconnects, the predominant choice in most real-world systems. Such interconnects benefit from well-established standards, CAD compatibility, predictable performance, and are highly scalable for many types of applications. However, in spite of their importance for current MPSoCs and their recent technology advancements, the design methodologies for structured on-chip interconnects have never been exhaustively surveyed so far, unlike application-oblivious interconnect solutions like Networks-on-Chip. The essential aim of this paper is to fill this gap by presenting an extensive review of state-of-the-art design automation techniques for application-specific on-chip interconnects. The paper goes through the main options available for building different on-chip interconnect topologies, discussing the details of hierarchical buses, crossbars, and cascaded crossbars as well as the approaches that can be adopted to formalize the description of such topologies and the related parameters of interest. Then, the paper surveys the most relevant techniques proposed in the literature to analyze a given interconnect solution, i.e. quantify parameters such as latency, bandwidth, area cost, power consumption, operating frequency, followed by an in-depth review of the main approaches for interconnect synthesis, including several advanced aspects such as co-synthesis of memory and communication architectures, joint scheduling and interconnect synthesis, floorplanning, dynamic configuration, multi-path communication. After presenting the above approaches, the paper discusses the potential impact that the body of research in the area of on-chip interconnects may have on current trends and emerging interconnect technologies.

© 2015 Elsevier B.V. All rights reserved.

**Contents**



*E-mail addresses:* acilardo@unina.it (A. Cilardo), edoardo.fusella@unina.it (E. Fusella).
[1] Tel.: +390817683852; fax: +390817683816.

http://dx.doi.org/10.1016/j.vlsi.2015.07.017
0167-9260/© 2015 Elsevier B.V. All rights reserved.

## 1. Introduction

Electronic system design is being revolutionized by the widespread adoption of the MultiProcessor System-on-Chip (MPSoC) paradigm [1]. In fact, to support aggressive performance challenges, the semiconductor industry has moved from the single processor frequency scaling entering the era of parallelization. In a few years, we will be able to build heterogeneous many-core systems with thousands of small general and special purpose cores as well as large memory elements (MEs) such as shared and local caches [2]. In this scenario, on-chip communication architectures provide a vital facility, enabling the computation to be distributed among the different processing elements (PEs) and data to be scattered across the MEs in order to maximize performance.

To exploit the MPSoC potential benefits to the fullest, suitable design methodologies are required for addressing two different facets of system design [3]. First, it is essential to properly map the application's computation requirements to a set of PEs like CPUs, DSPs, application specific cores, etc. Second, it is equally necessary to map the system's communication requirements onto an optimized communication architecture possibly tailored on the specific application. In particular, the choice of the underlying communication architecture in data-intensive applications, driven by the particular application requirements, is a critical design step since the amount of communication among functional blocks critically affects the global performance [4]. Furthermore, to meet the tight time-to-market constraints and efficiently handle the design complexity, we also need suitable computer-aided design (CAD) tools supporting the automation of these tasks.

This survey focuses on the application-driven automated design of a broad class of communication architectures, here referred to as *structured on-chip interconnects*. While this class includes usual shared buses and full crossbars, suffering from either limited performance or poor scalability, it also embraces complex scalable high-performance interconnects with customizable topologies such as hierarchical architectures, rings, cascaded crossbars as well as ad-hoc heterogeneous networks made up of a combination of shared buses and crossbars. In this section, we will briefly review the background of structured on-chip interconnect technologies. We will then point out the advantages and disadvantages of this architectural paradigm compared to a prominent, usually application-oblivious approach to communication in MPSoCs, i.e. Network-on-Chip (NoC) [5–7].

### 1.1. On-chip communication

*Shared buses* are the traditional and simplest on-chip communication architecture, consisting of a set of shared parallel wires connected to all components in the system. At any given time, only one PE can drive the bus. This limits the achievable concurrency of the system, which makes shared buses non-scalable and unsuitable for highly parallel MPSoC applications. Advanced shared buses appeared during the last years have introduced several specific features, such as separation of address/control and data phases, pipelined operations, burst-based, split and out-of-order transactions, etc. [8]. A few examples include AMBA™ (AHB/APB) [9] by ARM® and CoreConnect™ (PLB/OPB) [10] by IBM®. In addition, several research works have focused on improving some features of the shared bus architecture itself, such as the arbitration schemes [11–15]. However, single shared buses remain a major performance bottleneck for the majority of data-intensive applications due to their inherent lack of parallelism.

To overcome the physical limits of shared-buses, hierarchical architectures consisting of several buses interconnected through bridge components were introduced [8]. In addition to improving the potential bandwidth, this approach introduces a new dimension in the design space, the definition of the interconnect topology, which becomes the main parameter affecting the overall performance of the communication architecture. In fact, based on the positions of the bridges, we can essentially build any topology. As an example, a ring can be built as a group of consecutive bus segments, which can operate in parallel, connected through bridges, which may be mono- or bi-directional. To take this new design dimension into account, several newly introduced design methodologies include automated topology synthesis as a central step for improving performance.

*Crossbars*, also called bus matrices, are another major design alternative. They consist of a multi-layered communication architecture with multiple buses operating in parallel connecting multiple inputs to multiple outputs in a matrix-like scheme. A crossbar can be full or partial depending on the required connectivity between inputs and outputs. If a crossbar with $M$ inputs

Download English Version:

https://daneshyari.com/en/article/542690

Download Persian Version:

https://daneshyari.com/article/542690

Daneshyari.com