

Chapter 49

Adaptive Networks for On-Chip Communication

Mário Pereira Vestias

Instituto Politécnico de Lisboa, Portugal

ABSTRACT

The second generation of network-on-chips (NoC) are dynamic or adaptive providing a new set of benefits in terms of area overhead, performance, power consumption, fault tolerance, and quality of service compared to the previous generation where the architecture is decided at design time. To improve resource efficiency and performance, the NoC must consider adaptive processes at several architectural levels, including the routing protocols, the router, the network interface, and the network topology. This chapter focuses on adaptive networks-on-chip, namely adaptive topologies and adaptive routers.

INTRODUCTION

Networks-on-Chip (NoC) are a scalable interconnection network for on-chip communication capable to integrate a high number of processing elements. Scalability, energy efficiency and reliability are among the most important advantages of this new communication paradigm. Hundreds or even thousands of cores can be integrated in a single device using a NoC structure without facing the non-scalability problems associated with bus-based structures or point-to-point connections which are usually irregular and harder to route. Global on-chip communication with long wires thus not scales down with increasing clock frequency. The new communication paradigm decouples the cores from the network, reducing the need for global synchronization, reduces the number of global wires and the energy consumption of cores can be individually controlled. NoC are also reliable since fault-tolerant techniques can be implemented from hardware redundancy to adaptive routing protocols that look for alternative paths for a communication.

The first generation of NoC solutions considers regular topologies, typically 2D meshes under the assumption that the wires' layout is well structured in such topologies. Routers and network interfaces between IP cores and routers are mainly homogeneous so that they can be easily scaled up and facilitate modular design. All advantages of a NoC infrastructure were proven with this first generation of NoC solutions.

However, soon, the designers started to be worried about the two main disadvantages associated with NoCs, namely, area and speed overhead. Routers of a NoC need space for buffers, routing tables,

DOI: 10.4018/978-1-5225-7368-5.ch049

switching circuit and controllers. On the other side, direct bus connection is always faster than pipelined connections through one or more routers since these introduce latency due to packaging, routing, switching and buffering.

In a first attempt to consider area and latency in the design process, designers considered that regular NoC structures may probably be adequate for general-purpose computing where processing and data communication are relatively equally distributed among all processing units and traffic characteristics cannot be predicted at design time. But, many systems developed for a specific class of applications exhibit an intrinsic heterogeneous traffic behavior. Since routers introduce a relative area overhead and increase the average communication latency, considering a homogenous structure for a specific traffic scenario is definitely a waste of resources, a communication performance degradation and an excessive power consumption.

Application specific systems can benefit from heterogeneous communication infrastructures providing high bandwidth in a localized fashion where it is needed to eliminate bottlenecks (Benini & De Micheli, 2002), with sized communication resources to reduce area utilization, and low latency wherever this is a concern.

Homogeneous and heterogeneous solutions of first generation NoCs follow different design methodologies but have one thing in common: their architectures are found at design time and are kept fixed at runtime, i.e., the topology and the architecture of the routers are fixed at design time. Apparently, this is not a problem, but since several applications may be running with the same NoC, the same topology and router will generally not be equally efficient in terms of area, performance and power consumption for all different applications. The efficiency of both homogeneous and heterogeneous solutions can be improved if runtime changes are considered. A system running a set of applications can benefit from the runtime reconfiguration of the topology and of the routers to improve performance, area and power consumption considering a particular data communication pattern. Customization of the number of ports, the size of buffers, the switching techniques, the routing algorithms, the switch matrix configuration, etc. should be considered in a reconfigurable NoC. Both general purpose and application-specific System-on-Chips (SoCs) will benefit from using dynamically reconfigurable NoCs since the performance and power consumption of data communication can be optimized for each application.

The second generation of NoCs are dynamic or adaptive providing a new set of benefits in terms of area overhead, performance, power consumption, fault tolerance and quality of service compared to the previous generation where the architecture is decided at design time. To improve resource efficiency and performance, the NoC must consider adaptive processes at several architectural levels, including the routing protocols, the router, the network interface and the network topology.

This article focuses on adaptive Networks-on-Chip. Adaptive topologies and adaptive routers are described in the following sections.

BACKGROUND

More than a hundred of proposals of NoC architectures can be found in the literature (Salminen, Kulmala & Hämmäläinen, 2008). These NoC proposals differ in the used topology, the routing and the switching schemes, the design metrics and the target application. Routers have been also extensively studied, designed and implemented with different flit widths, buffer sizes, switching and routing mechanisms

10 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/adaptive-networks-for-on-chip-communication/213167

Related Content

Technology Satisfaction in an Academic Context: Moderating Effect of Gender

A. Y. M. Atiquil Islam (2017). *Research Paradigms and Contemporary Perspectives on Human-Technology Interaction* (pp. 187-211).

www.irma-international.org/chapter/technology-satisfaction-in-an-academic-context/176116

Augmented Reality Interfaces for Smart Objects in Ubiquitous Computing Environments

A. W. W. Yew, S. K. Ongand A. Y. C. Nee (2014). *Human-Computer Interfaces and Interactivity: Emergent Research and Applications* (pp. 208-229).

www.irma-international.org/chapter/augmented-reality-interfaces-for-smart-objects-in-ubiquitous-computing-environments/111758

Student Development of E-Workbooks: A Case for Situated-Technology Enhanced Learning (STEL) Using Net Tablets

Selena Chan, Katrina Fisherand Peter Sauer (2016). *Human-Computer Interaction: Concepts, Methodologies, Tools, and Applications* (pp. 447-468).

www.irma-international.org/chapter/student-development-of-e-workbooks/139048

Open Source Applications for Image Visualization and Processing in Neuroimaging Training

Juan A. Juanes, Pablo Ruisoto, Alberto Prats-Galinoand Andrés Framiñán (2016). *Human-Computer Interaction: Concepts, Methodologies, Tools, and Applications* (pp. 1319-1332).

www.irma-international.org/chapter/open-source-applications-for-image-visualization-and-processing-in-neuroimaging-training/139094

The Fashionable Functions Reloaded: An Updated Google Ngram View of Trends in Functional Differentiation (1800-2000)

Steffen Roth, Carlton Clarkand Jan Berkel (2017). *Research Paradigms and Contemporary Perspectives on Human-Technology Interaction* (pp. 236-265).

www.irma-international.org/chapter/the-fashionable-functions-reloaded/176119