Available online at www.sciencedirect.com
ScienceDirect
INTEGRATION, the VLSI journal 41 (2008) 171–182
INTEGRATION the VLSI journal
www.elsevier.com/locate/vlsi

Buffering global interconnects in structured ASIC design ☆

Tianpei Zhang*, Sachin S. Sapatnekar

Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA

Received 1 August 2006; received in revised form 1 April 2007; accepted 4 April 2007

Abstract

Structured ASICs present an attractive alternative to reducing design costs and turnaround times in nanometer designs. As with conventional ASICs, such designs require global wires to be buffered. However, via-programmable designs must prefabricate and preplace buffers in the layout. This paper proposes a novel and accurate statistical estimation technique for distributing prefabricated buffers through a layout. It employs Rent's rule to estimate the buffer distribution required for the layout, so that an appropriate structured ASIC may be selected for the design. Experimental results show that the buffer distribution estimation is accurate and economic, and that a uniform buffer distribution can maintain a high degree of regularity in design and shows a good timing performance, comparable with nonuniform buffer distribution.

Keywords: Structured ASIC; Rent's rule; Buffer insertion; Interconnect; Physical design

1. Introduction

In the nanometer regime, cell-based ASIC designs are increasingly hard-pressed to produce affordable design solutions, due to the challenges associated with skyrocketing mask costs and manufacturability issues for complex designs [1]. As an alternative, field programmable gate arrays (FPGAs) may be used since they provide a regular prefabricated structure that can avoid many of the manufacturing problems associated with ASICs. However, for many designs, the performance gaps, in terms of speed, power and area, between FPGAs and cell-based ASIC designs are too large for an FPGA to be a realistic alternative. In this context, structured ASIC design has emerged as a promising new design style to fill the gap.

Structured ASICs are composed of regular arrays of prefabricated standard building blocks, with fixed mask structures. Design with structured ASICs involves many fewer masks than for cell-based ASICs. One paradigm that is used involves *via-configurability*, where the building blocks and interconnect skeletons are prefabricated and are then connected with appropriate via connections by programming only a small number of masks [2–6]. A standard building block is composed of combinational logic and memory elements (such as flip-flops) and has enough flexibility that it can be programmed to various functions through via configurations. The wires that electrically connect the building blocks are also prefabricated regular fabrics, and the routing of nets can be configured with a via definition. This strategy provides a low NRE (non-recurring engineering) cost, and yet with relatively high-performance solutions. Additionally, since structured ASICs use well-characterized logic blocks and regular, fixed interconnect structures, they are well suited to combat problems associated with manufacturability, yield, noise, or crosstalk.

However, many of the other problems of nanometer design continue to plague structured ASICs. Most importantly, as in cell-based ASICs, the dominant role of interconnect in determining the system performance remains a major hurdle, and a key issue in overcoming this is through the use of buffer insertion along global wires. Buffer insertion cannot only improve timing performance, but can also effectively reduce functional noise by recovering noise margins [7]. The number of

☆ This research is partly supported by NSF award CCF-0205227.
*Corresponding author.
E-mail addresses: zhangt@ece.umn.edu (T. Zhang), sachin@ece.umn.edu (S.S. Sapatnekar).

0167-9260/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.vlsi.2007.04.002

buffers necessary to achieve timing closure and meet noise requirements continues to rise with decreasing feature sizes, and it is projected that at the 32 nm technology node, a very large proportion of cells will be buffers [8]. Although interconnects are generally buffered in the later phases of physical design, buffer resources must be planned earlier in the design process, so that enough resources are available for the later insertion phase. Several buffer resource planning strategies for cell-based ASIC design have been proposed in [9,10].

This problem is more acute for via-programmable structured ASICs, where, in addition to the basic standard building blocks, buffers must be prefabricated and distributed in the layout. Although it is possible to reconfigure the basic building blocks to work as buffers, this is not an economical approach since (a) the number of buffers can be very large, (b) the sizes of the buffers are typically larger than the sizes of regular gates, (c) configuring a large and general standard building blocks as buffers is an inefficient use of resources and (d) these blocks do not have the driving ability of dedicated buffers. Therefore, it is essential to distribute dedicated buffers in structured ASICs, and to plan for them well, prior to the fabrication of the chip.

The buffer insertion problem for structured ASIC design has not been fully addressed in publications so far. The only work we know of that considers this issue is [11], where it is assumed that a uniform distribution of dedicated buffers is placed throughout the layout, and there is a ratio of 2:1 between the number of logic cells and buffers everywhere in the circuit. However, this does not recognize that the demand for buffers depends on the interconnect complexity of circuits, and assuming a single 2:1 ratio for all kind of circuits may result in a large waste of buffer resources.

Clearly, the choice of this ratio should depend on the topology and structure of the circuit that is being mapped to the chip and the number of interconnects. On the other hand, given that this ratio must be predetermined in a structured ASIC, it is clearly not possible or realistic to tune each chip individually to a design, and a "good" set of buffer-to-logic-cell ratios must be chosen.

This paper uses Rent's rule to develop a family of "good" ratios and proposes a distributed buffer insertion methodology for the use of dedicated buffers in structured ASIC design. For each range of (p,k) values, where p is the Rent's exponent and k the Rent's coefficient, our algorithm (described in Section 3.3) finds a statistical estimate of the buffer distribution for circuits falling in that range. Thus for each range, we can prefabricate an off-the-shelf structured ASIC chip with buffers preplaced according to the estimated buffer distribution of that (p,k) range.

In the implementation phase, a designer may choose the appropriate prefabricated chip for implementing custom design according to values of (p,k) for the design. Experimental results show that the buffer resource estimation is accurate and adequate for interconnect buffering purpose of circuits in each range, and with an average uniform buffer distribution based on the estimation, we can maintain a good timing performance as well as a highly regular structured ASIC.

The organization of the paper is as follows. In Section 2, we describe the buffer insertion methodology. Section 3 provides a statistical buffer distribution estimation based on Rent's rule, and a classification method of circuits based on their Rent's exponent and coefficient. In Section 4, we present the experimental results, which is followed by conclusion in Section 5.

2. Buffer insertion scheme for structured ASIC design

Various buffer insertion models have been proposed for cell-based ASIC designs, prominent among which are the buffer block approach [9], in which blocks of buffers are placed *between* the building blocks of the circuit, and the distributed buffer approach [10], in which buffers are interspersed *within* the building blocks, with their exact location being undetermined until later in the design process. The distributed buffer insertion model has several advantages over the buffer block model: it spreads the routing wires around and avoids excessive routing congestion, while satisfying the requirement of buffer insertion.

For structured ASIC design, in the same spirit of interspersing buffers with logic units across the circuit, we adopt a buffer insertion model in which the prefabricated buffers are scattered through the structured ASIC, and the distribution of buffers should be adequate enough for buffering global wires. We also refer to this as a distributed buffer insertion model to capture the distributed nature of this scheme, although unlike [10], the buffers are actually prefabricated in structured ASICs. We define the structured ASIC to be a two-dimensional array composed of via-configurable standard building blocks, denoted as *via-configurable cells* (VCCs), which are connected by via-configurable interconnect wires. To facilitate the distributed buffering model, we divide the circuit into an array of *tiles*, and each tile is a square area containing $m \times m$ VCCs as well as a predetermined number of dedicated buffers for interconnect buffering usage. For a tile positioned at (i,j), we refer to this number as the *buffer capacity*, denoted as $B_{i,j}$. If $B_{i,j}$ is uniform for all (i,j), we refer to this as a *uniform* buffer distribution; otherwise the buffer distribution is said to be *nonuniform*. A routing solution may use some or all of these buffers: the actual number of utilized buffers in tile (i,j) is referred to as the *buffer usage*, denoted as $b_{i,j}$. If $b_{i,j} > B_{i,j}$, the routing solution is invalid, and we refer to this situation as *buffer overflow*. Fig. 1(a) shows a 6×6 tile graph for a structured ASIC design, with buffers distributed within tiles. The corresponding buffer capacity of each tile in the circuit with a nonuniform buffer distribution is shown in Fig. 1(b).

In this paradigm, buffers are prefabricated into the layout but are connected to the global lines that require buffering only in the later phases of physical design. To enable this, we utilize a via-defined buffer insertion (VDBI) scheme, which allows a buffer in a tile to be inserted along