PRINCIPLES OF ASYNCHRONOUS CIRCUIT DESIGN – A Systems Perspective

Edited by

JENS SPARSØ
Technical University of Denmark

STEVE FURBER
The University of Manchester, UK

Kluwer Academic Publishers
Boston/Dordrecht/London

Contents

Preface

This book was compiled to address a perceived need for an introductory text on asynchronous design. There are several highly technical books on aspects of the subject, but no obvious starting point for a designer who wishes to become acquainted for the first time with asynchronous technology. We hope this book will serve as that starting point.

The reader is assumed to have some background in digital design. We assume that concepts such as logic gates, flip-flops and Boolean logic are familiar. Some of the latter sections also assume familiarity with the higher levels of digital design such as microprocessor architectures and systems-on-chip, but readers unfamiliar with these topics should still find the majority of the book accessible.

The intended audience for the book comprises the following groups:

- Industrial designers with a background in conventional (clocked) digital design who wish to gain an understanding of asynchronous design in order, for example, to establish whether or not it may be advantageous to use asynchronous techniques in their next design task.
- Students in Electronic and/or Computer Engineering who are taking a course that includes aspects of asynchronous design.

The book is structured in three parts. Part I is a tutorial in asynchronous design. It addresses the most important issue for the beginner, which is how to think about asynchronous systems. The first big hurdle to be cleared is that of mindset – asynchronous design requires a different mental approach from that normally employed in clocked design. Attempts to take an existing clocked system, strip out the clock and simply replace it with asynchronous handshakes are doomed to disappoint. Another hurdle is that of circuit design methodology – the existing body of literature presents an apparent plethora of disparate approaches. The aim of the tutorial is to get behind this and to present a single unified and coherent perspective which emphasizes the common ground. In this way the tutorial should enable the reader to begin to understand the characteristics of asynchronous systems in a way that will enable them to ‘think outside the box' of conventional clocked design and to create radical new design solutions that fully exploit the potential of clockless systems.

Once the asynchronous design mindset has been mastered, the second hurdle is designer productivity. VLSI designers are used to working in a highly productive environment supported by powerful automatic tools. Asynchronous design lags in its tools environment, but things are improving. Part II of the book gives an introduction to Balsa, a high-level synthesis system for asynchronous circuits. It is written by Doug Edwards (who has managed the Balsa development at the University of Manchester since its inception) and Andrew Bardsley (who has written most of the software). Balsa is not the solution to all asynchronous design problems, but it is capable of synthesizing very complex systems (for example, the 32-channel DMA controller used on the DRACO chip described in Chapter 15) and it is a good way to develop an understanding of asynchronous design 'in the large'.

Knowing how to think about asynchronous design and having access to suitable tools leaves one question: what can be built in this way? In Part III we offer a number of examples of complex asynchronous systems as illustrations of the answer to this question. In each of these examples the designers have been asked to provide descriptions that will provide the reader with insights into the design process. The examples include a commercial smart card chip designed at Philips and a Viterbi decoder designed at the University of Manchester. Part III closes with a discussion of the issues that come up in the design of advanced asynchronous microprocessors, focusing on the Amulet processor series, again developed at the University of Manchester.

Although the book is a compilation of contributions from different authors, each of these has been specifically written with the goals of the book in mind – to provide answers to the sorts of questions that a newcomer to asynchronous design is likely to ask. In order to keep the book accessible and to avoid it becoming an intimidating size, much valuable work has had to be omitted. Our objective in introducing you to asynchronous design is that you might become acquainted with it. If your relationship develops further, perhaps even into the full-blown affair that has smitten a few, included among whose number are the contributors to this book, you will, of course, want to know more. The book includes an extensive bibliography that will provide food enough for even the most insatiable of appetites.

JENS SPARSØ AND STEVE FURBER, SEPTEMBER 2001

Acknowledgments

Many people have helped significantly in the creation of this book. In addition to writing their respective chapters, several of the authors have also read and commented on drafts of other parts of the book, and the quality of the work as a whole has been enhanced as a result.

The editors are also grateful to Alan Williams, Russell Hobson and Steve Temple, for their careful reading of drafts of this book and their constructive suggestions for improvement.

Part I of the book has been used as a course text and the quality and consistency of the content improved by feedback from the students on the spring 2001 course “49425 Design of Asynchronous Circuits” at DTU.

Any remaining errors or omissions are the responsibility of the editors.

The writing of this book was initiated as a dissemination effort within the European Low-Power Initiative for Electronic System Design (ESD-LPD), and this book is part of the book series from this initiative. As will become clear, the book goes far beyond the dissemination of results from projects within in the ESD-LPD cluster, and the editors would like to acknowledge the support of the working group on asynchronous circuit design, ACiD-WG, that has provided a fruitful forum for interaction and the exchange of ideas. The ACiD-WG has been funded by the European Commission since 1992 under several Framework Programmes: FP3 Basic Research (EP7225), FP4 Technologies for Components and Subsystems (EP21949), and FP5 Microelectronics (IST-1999-29119).

I

ASYNCHRONOUS CIRCUIT DESIGN – A TUTORIAL

Author: Jens Sparsø
Technical University of Denmark
jsp@imm.dtu.dk

Abstract Asynchronous circuits have characteristics that differ significantly from those of synchronous circuits and, as will be clear from some of the later chapters in this book, it is possible exploit these characteristics to design circuits with very interesting performance parameters in terms of their power, performance, electromagnetic emissions (EMI), etc.

Asynchronous design is not yet a well-established and widely-used design methodology. There are textbooks that provide comprehensive coverage of the underlying theories, but the field has not yet matured to a point where there is an established currriculum and university tradition for teaching courses on asynchronous circuit design to electrical engineering and computer engineering students.

As this author sees the situation there is a gap between understanding the fundamentals and being able to design useful circuits of some complexity. The aim of Part I of this book is to provide a tutorial on asynchronous circuit design that fills this gap.

More specifically the aims are: (i) to introduce readers with background in synchronous digital circuit design to the fundamentals of asynchronous circuit design such that they are able to read and understand the literature, *and* (ii) to provide readers with an understanding of the "nature" of asynchronous circuits such that they are to design non-trivial circuits with interesting performance parameters.

The material is based on experience from the design of several asynchronous chips, and it has evolved over the last decade from tutorials given at a number of European conferences and from a number of special topics courses taught at the Technical University of Denmark and elsewhere. In May 1999 I gave a one-week intensive course at Delft University of Technology and it was when preparing for this course I felt that the material was shaping up, and I set out to write the following text. Most of the material has recently been used and debugged in a course at the Technical University of Denmark in the spring 2001. Supplemented by a few journal articles and a small design project, the text may be used for a one semester course on asynchronous design.

Keywords: asynchronous circuits, tutorial

II

BALSA - AN ASYNCHRONOUS HARDWARE SYNTHESIS SYSTEM

Author: Doug Edwards, Andrew Bardsley
Department of Computer Science
The University of Manchester
{*doug,bardsley*}*@cs.man.ac.uk*

Abstract Balsa is a system for describing and synthesising asynchronous circuits based on syntax-directed compilation into communicating handshake circuits. In these chapters, the basic Balsa design flow is described and several simple circuit examples are used to illustrate the Balsa language in an informal tutorial style. The section concludes with a walk-through of a major design exercise – a 4 channel DMA controller described entirely in Balsa.

Keywords: asynchronous circuits, high-level synthesis

III

LARGE-SCALE ASYNCHRONOUS DESIGNS

Abstract In this final part of the book we describe some large-scale asynchronous VLSI designs to illustrate the capabilities of this technology. The first two of these designs – the contactless smart card chip developed at Philips and the Viterbi decoder developed at the University of Manchester – were designed within EU-funded projects in the low-power design initiative that is the sponsor of this book series. The third chapter describes aspects of the Amulet microprocessor series, again from the University of Manchester, developed in several other EU-funded projects which, although outside the low-power design initiative, never-the-less still had low power as a significant objective.

The chips descibed in this part of the book are some of the largest and most complex asynchronous designs ever developed. Fully detailed descriptions of them are far beyond the scope of this book, but they are included to demonstrate that asynchronous design is fully capable of supporting large-scale designs, and they show what can be done with skilled and experienced design teams. The descriptions presented here have been written to give insight into the thinking processes that a designer of state-of-the-art asynchronous systems might go through in developing such designs.

Keywords: asynchronous circuits, large-scale designs

Chapter 13

DESCALE: [*]

a Design Experiment for a Smart Card Application consuming Low Energy

Joep Kessels & Ad Peeters
Philips Research, NL-5656AA Eindhoven, The Netherlands
{Joep.Kessels | Ad.Peeters}@philips.com

Torsten Kramer
Kramer-Consulting, D-21079 Hamburg, Germany
Kramer@kramer-consulting.de

Volker Timm
Philips Semiconductors, D-22529 Hamburg, Germany
Volker.Timm@philips.com

Abstract We have designed an asynchronous chip for contactless smart cards. Asynchronous circuits have two power properties that make them very suitable for contactless devices: low average power and small current peaks. The fact that asynchronous circuits operate over a wide range of the supply voltage, while automatically adapting their speed, has been used to obtain a circuit that is very resilient to voltage drops while giving maximum performance for the power being received. The asynchronous circuit has been built, tested and evaluated.

Keywords: low-power asynchronous circuits, smart cards, contactless devices, DES cryptography

[*] Part of the work described in this paper was funded by the European Commission under Esprit TCS/ESD-LPD contract 25519 (Design Experiment DESCALE).

Chapter 14

AN ASYNCHRONOUS VITERBI DECODER [*]

Linda E. M. Brackenbury

Department of Computer Science, The University of Manchester

lbrackenbury@cs.man.ac.uk

Abstract Viterbi decoders are used for decoding data encoded using convolutional forward error correcting codes. Such codes are used in a large proportion of digital transmission and digital recording systems because, even when the transmitted signal is subjected to significant noise, the decoder is still able efficiently to determine the most likely transmitted data.

This chapter descibes a novel Viterbi decoder aimed at being power efficient through adopting an asynchronous approach. The new design is based upon serial unary arithmetic for the computation and storage of the metrics required; this arithmetic replaces the add-compare-select parallel arithmetic performed by conventional synchronous systems. Like all Viterbi decoders, a history of computational results is built up over many data bits to determine the data most likely to have been transmitted at an earlier time. The identification of a starting point to this tracing operation allows the storage requirement to be greatly reduced compared with that in conventional decoders where the starting point is random. Furthermore, asynchronous operation in the system described enables multiple, independent, concurrent tracing operations to be performed which are decoupled from the placing of new data in the history memory.

Keywords: low-power asynchronous circuits, Viterbi, convolution decoder

14.1. Introduction

The PREST (Power REduction for Systems Technology) [1] project was a collaborative project where each partner designed a low power alternative to

[*] The Viterbi design work was supported by the EPSRC/MoD PowerPack project GR/L27930 and the EU PREST project EP25242, and this support is gratefully acknowledged.

Chapter 15

PROCESSORS [*]

Jim D. Garside
Department of Computer Science, The University of Manchester
jgarside@cs.man.ac.uk

Abstract Computer design becomes ever more complex. Small asynchronous systems may be intriguing and even elegant but unless asynchronous logic can not only be competitive with 'conventional' logic but can show some significant advantages it cannot be taken seriously in the commercial world.

There can be no better way to demonstrate the feasibility of something than by doing it. To this end several research groups around have the world have been putting together real, large, asynchronous systems. These have taken several forms, but many groups have chosen to start with microprocessors; a processor is a good demonstrator because it is well defined, self-contained and forces a designer to solve problems which are already well understood. If an asynchronous implementation of a microprocessor can compare favourably with a synchronous device performing an identical function then the case is proven.

This chapter describes a number of processors that have been fabricated and discusses in some detail some of the solutions employed. The primary source of the material is the Amulet series of ARM implementations – because these are the most familiar to the author – but other devices are included as appropriate. The later parts of the chapter widen the descriptions to include memory systems, cacheing and on-chip interconnect, illustrating how a complete asynchronous System on Chip (SoC) can be produced.

Keywords: low-power asynchronous circuits, processor architecture

[*] The majority of the work described in this chapter has been made possible by grants from the European Union Open Microprocessor systems Initiative (OMI). The primary sources of funding have been OMI-MAP (Amulet1), OMI/DE-ARM (Amulet2e) and OMI/ATOM (Amulet3). Without this funding none of these devices would have been made and this support is gratefully acknowledged.