
School of Electrical, Electronic & Computer Engineering



Automated Generation of Processor Architectures in Embedded Systems Design

Maxim Rykunov, Andrey Mokhov, Alex Yakovlev, Albert Koelmans

Technical Report Series
NCL-EECE-MSD-TR-2010-158

July 2010

Contact:

maxim.rykunov@ncl.ac.uk

andrey.mokhov@ncl.ac.uk

alex.yakovlev@ncl.ac.uk

albert.koelmans@ncl.ac.uk

Supported by EPSRC grants EP/G037809/1

NCL-EECE-MSD-TR-2010-158

Copyright © 2010 University of Newcastle upon Tyne

School of Electrical, Electronic & Computer Engineering,
Merz Court,

University of Newcastle upon Tyne,
Newcastle upon Tyne, NE1 7RU, UK

<http://async.org.uk/>

Maxim Rykunov, Andrey Mokhov, Alex Yakovlev, Albert Koelmans

Microelectronics System Design Group, Newcastle University, UK

Abstract

Automated design of processor architectures has traditionally been focused on the clocked pipeline organisation consisting of a fairly standard datapath and control logic. Control logic has been normally generated from the architectural description of a processor using the conventional techniques based on Finite State Machines (FSMs). As the area of processor design automation is becoming increasingly inclusive of system paradigms that are heterogeneous in terms of timing, such as multicycle and asynchronous circuits, there is a need for appropriate models and associated synthesis algorithms.

This paper approaches the problem of designing control logic for a processor using Conditional Partial Order Graphs (CPOGs). The new method allows composing a large set of microarchitectural algorithms (instructions) into a compact relational form, which opens way for various transformation and optimisation procedures leading to an efficient implementation of control logic. The paper presents a CPOG-based design methodology and demonstrates how it can be applied to it on synthesis of control logic for Intel 8051 microcontroller.

1 Introduction

Nowadays microprocessor designers face new challenges associated with the rapidly changing requirements for power consumption, high performance, area efficiency, etc., as well as demand for shorter time to market and greater productivity. To date, there exist approaches and associated EDA (Electronic Design Automation) support for automated generation of Application Specific Instruction set Processor (ASIP) [1][2]. They show a general processor design environment based on micro-architectural description of instruction algorithms. For example, the PEAS-III system [3] allows the generation of a complete tool-suite and includes a compiler, assembler, linker and simulator; the MetaCore system [4] is a benchmark driven ASIP development system, based on a formal representation language (see also the Xtensa system [5], the EPICS system [6], etc.). While these methodologies have reached significant levels in automating the whole design process, they are limited to design of synchronous systems. However, modern processors are increasingly diversified in timing modes, which calls for new ways of handling concurrent and asynchronous interactions.

In principle, one could apply a recently developed desynchronisation method, where the synchronous design is converted automatically into an asynchronous one [7][8][9]. Desynchronisation has the advantage of minimal changes in the conventional design flow, and may be an appropriate way towards design reuse, i.e. obtaining new asynchronous implementations from existing synchronous designs. In the event of a rapidly evolving processor specification, a more “direct” way of incorporating new instructions and their microarchitectural algorithms will be needed.

There are direct methods in asynchronous design and they have been applied to CPUs. However they are known to lead to inefficient implementations, especially in terms of area and performance [10][11][12]. This is caused by the fact that multiple algorithmic procedures for a large set of instructions require explicit enumeration of all alternatives. Syntax-directed methods instantiate every branch as a hardware control thread, thereby leading to large overheads.

The key part in a microprocessor control specification is the description of instructions. Each instruction corresponds to a schedule of primitive actions such as data transfer, arithmetic operation, memory

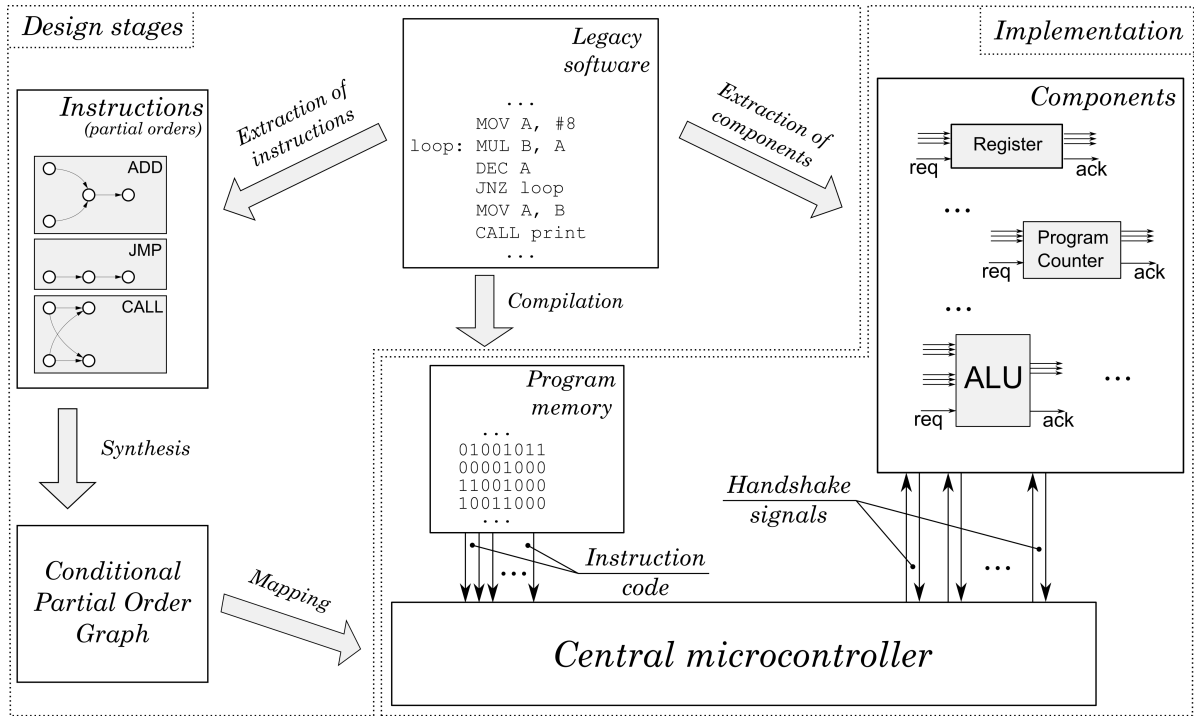


Figure 1: Conceptual view of the design process

access, etc., which are performed by operational units (datapath components). The design of instruction sets for a particular combination of operational units and software requirements is a difficult task [2]. Moreover, it is worth mentioning that there is a significant difference between designing CPU control and other types of hardware. Namely, microprocessor control contains many algorithms, which are associated with a variety of instructions rather than a single algorithm, which is commonly present in random logic design.

The main contribution of this work is a new design flow for automated synthesis of CPU control as a composition of scheduled algorithms of CPU instructions (see Figure 1). It is particularly important to consider this flow in the context of diversifying requirements, such as high performance, low power consumption, high security, etc. This flow is based on a new model: Conditional Partial Order Graphs (CPOG). CPOGs facilitate systematic transformation of a superposition of microarchitectural control algorithms into a netlist for the central control logic of a CPU.

The rest of this paper is organised as follows. In Section 2 we discuss the motivation for a model to describe the microarchitectural level of CPU design and other goals in achieving an effective design flow for synthesis of control circuits for microprocessors. Section 3 provides a quick overview of a new model for system specification and synthesis – CPOG. Section 4 outlines the whole proposed design flow. Section 5 illustrates the application of the flow to the synthesis of control for Intel’s 8051, we summarise the contribution of this paper and discuss the future work in Section 6.

2 Motivation for a new model

As processors become more diverse in terms of timing, the conventional FSM-oriented approach to designing their control logic starts to hit fundamental limitations. An FSM captures behaviour of a system using an explicit set of *states* and *transitions* between these states. There is a direct correspondence between the states of an FSM and those of the modelled system, which leads to problems in terms of concurrency specification. For example, to specify n concurrent events, the designer has to explicitly specify 2^n intermediate states, each of them corresponding to a state, where a subset of these n events have already happened. There are 2^n different subsets of a set of n events, hence the number of the

required states [13]. This brings us to the first requirement for a new model: it should be able to capture concurrency without explicit enumeration of all possible event interleavings.

Models like Petri Nets (PNs) [14] and Signal Transition Graphs (STGs) [15] are able to capture concurrency and choice at a very fine level, and produce more compact and faster control circuits than the methods based on syntax-directed translation from HDLs. However, they are built on the explicit listing of all the event traces and causal relations of a system, what limits their applicability to microcontrollers with a rather small state space. Processors may contain hundreds of encoded instructions, each a microprogram operating on a shared set of datapath components. It is practically impossible to specify such systems using PNs or STGs, as demonstrated in [13].

Another requirement for a new model is *expressiveness*: it should be expressive enough to cover a wide range of solutions for different optimisation criteria, thus enabling a prompt retargeting of the synthesised processor for varying and often conflicting optimisation goals. For example, we can retarget a CPU architecture in order to:

- minimise power consumption – this can be achieved by reducing the width of buses between datapath components and/or by using low-power components; this inevitably affects the performance of the system and necessitates rescheduling of the execution of the instructions.
- maximise performance – this can be achieved by increasing the number of functional units, their bit-width and therefore their performance; the instructions are to be rescheduled in a maximally concurrent way.
- optimise the architecture for a particular software application – this may involve complete re-encoding of the instructions, taking into account their statistical properties.
- etc.

The Conditional Partial Order Graph model introduced recently [13] satisfies all the requirements and provides a formal framework for specification, verification and synthesis of processor microarchitectures. In this paper we present a processor design methodology built around this model. The next section briefly introduces the CPOG model.

3 CPOG essentials

Conditional Partial Order Graph [13][16] is a quintuple $H(V, E, X, \rho, \phi)$, where V is a finite set of *vertices*, $E \subseteq V \times V$ is a set of *arcs* between them, and X is a finite set of *operational variables*. An *opcode* is an assignment $(x_1, x_2, \dots, x_{|X|}) \in \{0, 1\}^{|X|}$ of these variables; X can be assigned only those opcodes which satisfy the *restriction function* ρ of the graph, i.e. $\rho(x_1, x_2, \dots, x_{|X|}) = 1$. Function ϕ assigns a Boolean *condition* $\phi(z)$ to every vertex and arc $z \in V \cup E$ of the graph.

In the context of processor microarchitecture we interpret vertices V as data path components whose order of execution is determined by an instruction code provided by operational variables in X . The order is captured in the precedence relation E which is conditional (every vertex or arc $z \in V \cup E$ has an associated condition $\phi(z)$). The restriction function ρ effectively lists all the allowed instruction codes.

Figure 2(a) shows an example of a CPOG containing $|V| = 5$ vertices and $|E| = 7$ arcs. There is a single operational variable x ; the restriction function is $\rho(x) = 1$, hence both opcodes $x = 0$ and $x = 1$ are allowed. Vertices $\{a, b, d\}$ have constant $\phi = 1$ conditions and are called *unconditional*, while vertices $\{c, e\}$ are *conditional* and have conditions $\phi(c) = x$ and $\phi(e) = \bar{x}$ respectively. Arcs also fall into two classes: *unconditional* (arc (c, d)) and *conditional* (all the rest). As CPOGs tend to have many unconditional vertices and arcs we use a simplified notation in which conditions equal to 1 are not depicted in the graph. This is demonstrated in Figure 2(b).




Figure 2: Graphical representation of CPOGs




Figure 3: Multiple CPOG projections

The purpose of conditions ϕ is to ‘switch off’ some vertices and/or arcs in the graph according to the given opcode. This makes CPOGs capable of specifying multiple partial orders or instructions (a partial order is a form of behavioural description of an instruction). Figure 3 shows a graph and its two *projections*. The leftmost projection is obtained by keeping only those vertices and arcs whose conditions evaluate to 1 after substitution of the operational variable x with 1. Hence, vertex e disappears, because its condition evaluates to 0: $\phi(e) = \bar{x} = \bar{1} = 0$. Arcs $\{(a, d), (a, e), (b, d), (b, e)\}$ disappear for the same reason. The rightmost projection is obtained in the same way, with the only difference that variable x is set to 0. Note also that although the condition of arc (c, d) evaluates to 1 (in fact it is constant 1), the arc is still excluded from the resultant graph because one of the vertices it connects (vertex c) is excluded, and obviously an arc cannot appear in a graph without one of its vertices. Each of the obtained projections can be treated as a specification of a particular behavioural scenario of the modelled system. Potentially, a CPOG $H(V, E, X, \rho, \phi)$ can specify an exponential number of different partial orders of events in V according to one of $2^{|X|}$ different possible opcodes.

To conclude, a CPOG is a structure to represent a set of encoded partial orders in a compact form. The synthesis and optimisation methods presented in [16] provide a way to obtain such a representation given a set of partial orders and their opcodes. For example, the CPOG in Figure 3 can be synthesised automatically from the two partial orders below it and the corresponding opcodes $x = 1$ and $x = 0$.

4 Proposed design flow

Figure 1 shows an overview of the process that can be used to generate control logic of a CPU.

Specification of such a complex system as a processor usually starts at the architectural level [17]. This helps designers to deal with the system complexity by structural abstraction: the whole system is divided into several subsystems in a such way that each of them can be designed individually, thus significantly simplifying the whole design flow and reducing the solution search space. As stated in the introduction, the next particularly important step is refining the architecture to the level of scenarios or instructions, each of which corresponds to a schedule of primitive actions. In this aspect, the whole design flow can be separated into 3 subsections: system specification, synthesis and physical implementation.

4.1 System specification

First of all, we need to define the set of instructions (behavioural scenarios) and datapath components of the microprocessor. There are several ways to do this:

- Extraction from legacy software

Components and instructions can be extracted from some legacy software (could be written in C/C++, Assembly language, etc.).

- Using Architecture Description Languages (ADLs)

The ADL specification is used to help in performing different design automation tasks, e.g. hardware generation and functional verification of processors. There are several categories of ADLs, based on the nature of the available information, such as: structural ADLs (e.g. MIMOLA [18], UDL/I [19]), behavioural ADLs (e.g. nML [20], ISDL [21]), mixed ADLs (e.g. LISA [22], EXPRESSION [23]) and partial ADLs (e.g. AIDL [24]).

- Other instruction specifications

There can be some other sources, such as a list of instructions obtained from a microprocessor specification (e.g. a CPU manual). This type of specification is not formal and therefore should be processed manually.

In order to distinguish scenarios, they need to be encoded with Boolean vectors. These encodings can be assigned arbitrarily or can be provided as part of system specification. Importantly, the size and latency of the final microcontroller circuit depends significantly on the chosen encoding of the scenarios. Section 5 presents an example of a binary encoding scheme for a given set of scenarios. However, other types of encoding can also be used, such as one hot, matrix, balanced and others [25].

4.2 Synthesis

Once the behavioural scenarios and datapath components of the system are specified and encoded, it is possible to synthesise a CPOG which contains all of them by using algebraic techniques described in [13]. At this stage the obtained CPOG can be viewed as an intermediate model for a compact representation of the system behaviour, which will be further mapped onto a gate-level circuit. Therefore it is not necessary for the designer to be familiar with the CPOG-based theory in order to use the proposed design flow.

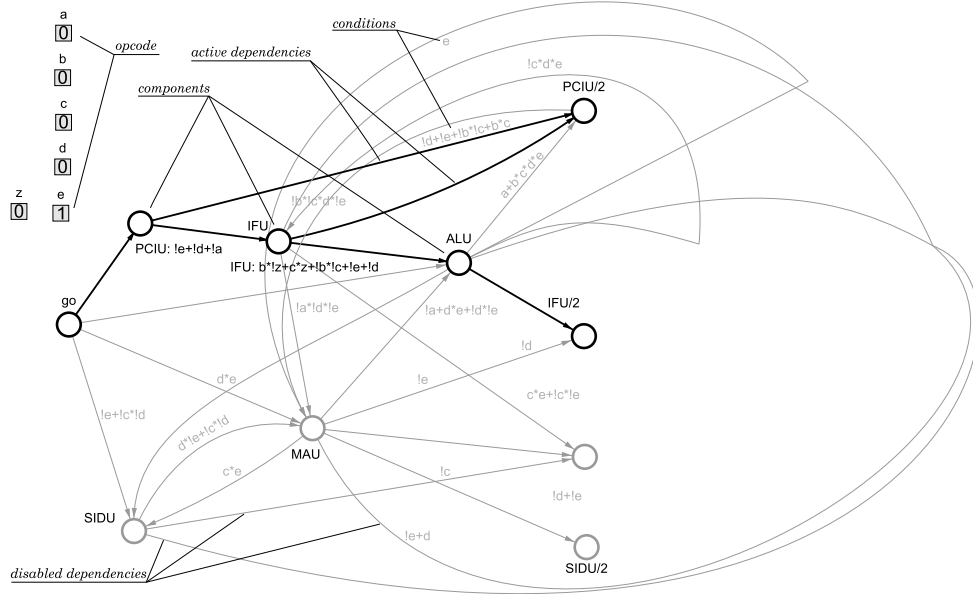


Figure 4: Synthesised CPOG

It is important to mention the flexibility of the CPOG-based methodology, in the sense that CPOG synthesis is possible not only from partial orders, it can also be a mixture of partial orders and CPOGs. This allows a complete reuse of the existing CPOG specifications without their decomposition into separate partial orders, so that one can add a new instruction into the specification of a system and avoid its complete resynthesis.

There are several possible optimisation techniques at this stage, which can reduce the size of a given CPOG using logic minimisation and by exploiting the structural graph properties. It is particularly important that the size and latency of the obtained microcontroller strongly correlates with the complexity of the original graph, therefore it is crucial to optimise the CPOG specification as far as possible by the application of CPOG optimisation techniques [16].

4.3 Physical implementation

The final stage of the our automated generation of control logic is mapping the CPOG representation into a set of logic gates. As soon as the CPOG specification of a system is synthesised it can be mapped onto Boolean equations in order to produce a physical implementation (gate-level netlist) of the specified microcontroller. The size of the resulting Boolean equations is linear in the size of the CPOG. The details of translating of CPOGs into Boolean equations can be found in [13]. Finally, we can translate them to VHDL, Verilog or other HDL (Hardware Description Language) and/or input these equations into the technology mapping and Place and Route (P&R) tools. See the example in the next Section.

Along with hardware mapping we have to perform software mapping, i.e. the compilation of the program code from a given legacy software. The compilation result is stored in the program memory.

Our design process, defined in Figure 1, shows that the interface between control and operation components is based on a handshake protocol. This allows significant flexibility in developing or reusing the datapath of the controller. Due to the handshake protocol, the full power of partial orders can be exploited, because the timing of control events is not bounded to particular delay constraints. The advantages of such an approach have been recently applied to the designs in [26][27][28]. All information about the number and type of the CPU components, such as registers, program counter, ALU (Arithmetic logic unit), etc., can also be extracted from the legacy software and later on added to the design.

In the next section we present an example of control logic synthesis.

5 Example of control logic synthesis

Regarding our design flow, first of all we need to specify the instructions of the microprocessor. We chose the instruction set from the core of the Intel 8051 microcontroller, as it is a well-known, popular CPU and used in various applications even nowadays.

Extractions of components. Functional components, such as PCIU, ALU, IFU, etc. can be extracted from the given instruction set (see Figure 4), moreover each functional block can be use several times in one partial order, what is shown as PCIU/2, PCIU/3, etc.

| Components | Description |
|------------|----------------------------------------|
| PCIU | Program Counter Increment Unit |
| IFU | Instruction Fetch Unit |
| ALU | Arithmetic Logic Unit |
| MAU | Memory Access Unit |
| SIDU | Stack pointer Increment Decrement Unit |

Extractions of instructions. Intel 8051 microcontroller's instruction set contains 110 instructions. Many of them have the same partial order representation, so we group them into 20 classes. For instance, one of the classes corresponds to a group of ALU operations (namely, *ALU op. #data to Rn*). In this set of instructions one of the operands is a register and the other is an immediate constant (e.g. MOV A, #1 – move 1 to accumulator, ANL A, #1 – bitwise "AND" operation between accumulator and 1). Figure 5 shows the partial order of the actions for such an instruction. First, the constant has to be fetched into the Instruction Register (IR) – execution of components PCIU and IFU. Then ALU is executed concurrently with an increment of Program Counter (PC). Finally, it is possible to fetch the next instruction into IR.

Other instructions are described in the same way.

Encoding of partial orders. In order to distinguish between the synthesised partial orders, we need to encode them. As stated in the proposed flow, the chosen type of encoding has significant influence on the size and complexity of their CPOG composition. In our example we used a binary encoding scheme. As we grouped all the instructions into 20 classes, we need at least 5 bits to encode them, i.e. assign opcodes {00000, 00001, ...10011}. Figure 5 shows that opcode 00000 was given to our example class. This encoding scheme has been chosen for the sake of simplicity; it might not be optimal in terms of area, performance or latency of the final microcontroller circuit. The decision as to which encoding has to be used in the final design is governed by the overall optimisation criteria, as pointed out in Section 2.




Figure 5: ALU op. #data to Rn

CPOG generation. Now, we have specified and encoded all the partial orders of the microcontroller. We can proceed to synthesis of a CPOG containing all of them. Figure 4 shows the obtained CPOG, which contains 15 vertexes and 46 arcs between them, with 5 opcode variables we received 225 literals. This result shows that it is a compact implementation and it would be impossible to describe this control circuit in such a compact form by using other techniques, such as FSMs. Most of the vertices and arcs have conditions, depending on which a particular partial order can be active or disabled. Our example

partial order from Figure 5 is highlighted in Figure 4. Its elements are activated in accordance with the evaluation of the conditions on the operational code 00000.

Mapping. We can translate the obtained CPOG into Boolean equations using algebraic techniques described [13]. Then these equations are imported into Altera's Quartus II tool for technology mapping into FPGA. The mapping result as a Register Transfer Level (RTL) netlist is shown in Figure 6. As a result our control circuit placed on the chip EP1S10F780I6 from the Stratix family contains 106 logic elements, which is only about 1% from the total chip area.

The synthesised control circuit was simulated using the same EDA tool. The results confirmed the correct implementation of the handshake protocol between the controller and data path components, and the expected orders of component activation according to given instructions.

The future work intends to validate the new CPU design flow by implementing both control and datapath for the 8051 microprocessor in hardware, first as FPGA and as ASIC later on. This will include experiments with different types of operational units (e.g. bundled data, dual-rail), and different encodings to best fit with the given program benchmarks and hardware requirements.

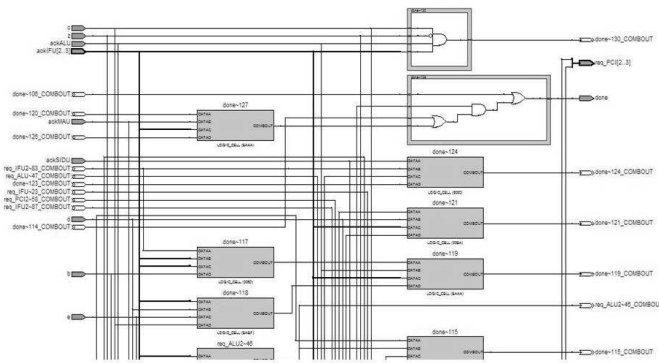


Figure 6: Synthesis of circuit implementation

6 Conclusion and Future Work

We have presented a new design flow for specification and synthesis of processor control logic. We used a new formal model for the description of instruction sets, to deal with concurrency and prompt retargeting of the CPU.

The presented example of control logic synthesis shows the capability of proposed design flow and the potential benefits of CPOGs-based methodology.

Our future work will focus on optimisation techniques for the current flow, such as implementing other types of instructions encodings and further development of self-timed pipelining within the represented methodology.

Acknowledgement

This work was supported by EPSRC grants EP/G037809/1.

References

- [1] Mishra and Kejariwal. Rapid exploration of pipelined processors through automatic generation of synthesizable rtl models. In *RSP '03: Proceedings of the 14th IEEE International Workshop on Rapid System Prototyping (RSP'03)*, page 226, Washington, DC, USA, 2003. IEEE Computer Society.
- [2] Scheffer and Grant Martin. *EDA for IC System Design, Verification, and Testing (Electronic Design Automation for Integrated Circuits Handbook)*. CRC Press, Inc., Boca Raton, FL, USA, 2006.

- [3] Makiko Itoh and Shigeaki Higaki. Peas-iii: An asip design environment. In *ICCD '00: Proceedings of the 2000 IEEE International Conference on Computer Design*, page 430, Washington, DC, USA, 2000. IEEE Computer Society.
- [4] Jin-Hyuk Yang and Chong-Min Kim. Metacore: an application specific dsp development system. In *DAC '98: Proceedings of the 35th annual Design Automation Conference*, pages 800–803, New York, NY, USA, 1998. ACM.
- [5] Ricardo E. Gonzalez. Xtensa: A configurable and extensible processor. *IEEE Micro*, 20(2):60–70, 2000.
- [6] R. Woudsma and R. A. M. Beltman. Epics, a flexible approach to embedded dsp cores. *Int'l Conference on Signal Processing Applications & Technology (ICSPAT)*, 651:506–511, Oct. 1994.
- [7] Nikolaos Andrikos, Luciano Lavagno, Davide Pandini, and Christos P. Sotiriou. A fully-automated desynchronization flow for synchronous circuits. In *DAC '07: Proceedings of the 44th annual Design Automation Conference*, pages 982–985, New York, NY, USA, 2007. ACM.
- [8] J. Cortadella, A. Kondratyev, L. Lavagno, K. Lwin, and C. Sotiriou. From synchronous to asynchronous: An automatic approach. In *DATE '04: Proceedings of the conference on Design, automation and test in Europe*, page 21368, Washington, DC, USA, 2004. IEEE Computer Society.
- [9] Jordi Cortadella, Alex Kondratyev, Senior Member, Luciano Lavagno, and Christos P. Sotiriou. Desynchronization: Synthesis of asynchronous circuits from synchronous specifications. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25:2006, 1994.
- [10] K. L. Chang and B. H. Gwee. A low-energy low-voltage asynchronous 8051 microcontroller cores. *Proc. ISCAS*, page 4, 2006.
- [11] L. A. Plana and Riocreux. Spa—a secure amulet core for smartcard applications. *Microprocessors and Microsystems*, 27:15, 2003/10.
- [12] Hans van Gageldonk, Kees van Berkel, Ad Peeters, Daniel Baumann, Daniel Gloor, and Gerhard Stegmann. An asynchronous low-power 80c51 microcontroller. *Asynchronous Circuits and Systems, International Symposium on*, 0:0096, 1998.
- [13] Andrey Mokhov. *Conditional Partial Order Graphs*. PhD thesis, School of EECE, Newcastle University, 2009.
- [14] Tadao Murata. Petri Nets: Properties, Analysis and Applications. 77(4), 1989.
- [15] J. Cortadella, M. Kishinevsky, A. Kondratyev, L. Lavagno, and A. Yakovlev. *Logic synthesis of asynchronous controllers and interfaces*. Springer-Verlag, ISBN: 3-540-43152-7, 2002. InternalNote: submitted by: hr.
- [16] Andrey Mokhov and Alex Yakovlev. Conditional partial order graphs and dynamically reconfigurable control synthesis. In *DATE '08: Proceedings of the conference on Design, automation and test in Europe*, pages 1142–1147, New York, NY, USA, 2008. ACM.
- [17] J. Silc and B. Robic. A survey of new research directions in microprocessors. *Microprocessors and Microsystems*, pages 175–190, 2000.
- [18] G. Zimmermann. The mimola design system: a computer aided digital processor design method. In *25 years of DAC: Papers on Twenty-five years of electronic design automation*, pages 525–530, New York, NY, USA, 1988. ACM.

- [19] Tamio Hoshino. Udl/i version two: A new horizon of hdl standards. In *CHDL '93: Proceedings of the 11th IFIP WG10.2 International Conference sponsored by IFIP WG10.2 and in cooperation with IEEE COMPSOC on Computer Hardware Description Languages and their Applications*, pages 437–452, Amsterdam, The Netherlands, The Netherlands, 1993. North-Holland Publishing Co.
- [20] A. Fauth and M. Freericks. Describing instruction set processors using nml. In *EDTC '95: Proceedings of the 1995 European conference on Design and Test*, page 503, Washington, DC, USA, 1995. IEEE Computer Society.
- [21] George Hadjiyiannis, Silvina Hanono, and Srinivas Devadas. Isdl: an instruction set description language for retargetability. In *DAC '97: Proceedings of the 34th annual Design Automation Conference*, pages 299–302, New York, NY, USA, 1997. ACM.
- [22] Andreas Hoffmann, Oliver Schliebusch, Achim Nohl, Gunnar Braun, Oliver Wahlen, and Heinrich Meyr. A methodology for the design of application specific instruction set processors (asip) using the machine description language lisa. In *ICCAD '01: Proceedings of the 2001 IEEE/ACM international conference on Computer-aided design*, pages 625–630, Piscataway, NJ, USA, 2001. IEEE Press.
- [23] Ashok Halambi, Peter Grun, Vijay Ganesh, Asheesh Khare, Nikil Dutt, and Alex Nicolau. Expression: a language for architecture exploration through compiler/simulator retargetability. In *DATE '99: Proceedings of the conference on Design, automation and test in Europe*, page 100, New York, NY, USA, 1999. ACM.
- [24] H. Morimoto and K. Yamazaki. Superscalar processor design with hardware description language aidl. In *2nd Asia Pacific Conference on Hardware Description*, volume 75, pages 51–58, Oct. 1994.
- [25] Alex Yakovlev Andrey Mokhov, Arseniy Alekseyev. Automated synthesis of instruction codes in the context of micro-architecture design. In *In Proc. of 10th Int. Conf. on Application of Concurrency to System Design (ACSD 2010)*, 2010.
- [26] Basit Riaz Sheikh and Rajit Manohar. An operand-optimized asynchronous ieee-754 double-precision floating-point adder. In *IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*, 2010.
- [27] V. Varshavsky (Ed.). Self-timed control of concurrent processes. *Kluwer Academic Publishers*, 1990.
- [28] Bah-Hwee Gwee; Chang; Yiqiong Shi; Chien-Chung Chua; Kwen-Siong Chong;. A low-voltage micropower asynchronous multiplier with shift add multiplication approach. *Circuits and Systems I: Regular Papers, IEEE Transactionson*, 56 Issue:7:1349 – 1359, July 2009.