FINAL REPORT (GR/S12036)
Synthesis and Testing of Low-Latency Asynchronous Circuits (STELLA)

A.Yakovlev, A. Bystrov, A.M.Koelmans, G. Russell and D.J.Kinniment
(School of Electrical, Electronic and Computer Engineering, Newcastle University,
Merz Court, Newcastle-upon-Tyne, NE1 7RU, England)
http://async.org.uk/

1 Background/Context

There is a consensus among the international research community that asynchrony is becoming an increasingly important issue at all levels because of rapid advances in complexity and feature-size-reduction of systems on chip (cf. the findings of ITRS-2005, http://www.itrs.net/Common/2005ITRS/Home2005.htm). The state of the art (as per June 2004) in asynchronous design tools is presented in the 3rd edition of ACiD-WG report "Design, Automation and Test for Asynchronous Circuits and Systems" by Edwards and Toms (http://www.scism.sbu.ac.uk/ccsv/ACiD-WG/). Perhaps, the best reflection of today's status of tool support would be the recent announcement of the availability of tools from Handshake Solutions from Europractice.

The research in STELLA was a natural progression from the successful EPSRC projects ASAP (GR/J52327), MOVIE (GR/M94366) and BESST (GR/R16754), along the line of investigating synthesis and testing of high speed control circuits. In BESST the target was to increase the size of controllers and improve productivity of their synthesis by applying direct mapping methods. MOVIE studied the visualization support and interactive synthesis of asynchronous controllers. STELLA's focus was on:

- synthesis of asynchronous control circuits for high-speed, and
- off-line and on-line testing of asynchronous circuits with minimum impact on their speed.

Additionally, bearing in mind the increasing use of globally asynchronous and locally synchronous (GALS) approaches in academia and industry, the project has also addressed issues of analysis and synthesis of latency-aware GALS methods. This research was closely related with other projects being conducted at Newcastle such as one on development of design flow for secure circuit design (SCREEN, GR/S81421) and more recently the NEGUS project (EP/C512812/1) on next generation of interconnection technology for multiprocessor SOC, carried out in collaboration with Southampton University, the former involving studying secure datapath architectures, including those in dual-rail, and the latter studying reliable self-timed on-chip communications.

2 Key advances and supporting methodology

STELLA's target was to develop an overall performance-oriented architecture and design methods for asynchronous control blocks with emphasis on minimization of input-output latency. The goal was to provide a high-productivity flow for synthesis and testing of circuits with low-latency. The methods and tools were supposed to be applicable to a wide range of designs but especially effective in the design of time-critical parts of SoCs such as on-chip network adaptors, interface controllers and synchronization logic, for example interfaces between timing domains of GALS. The project has largely followed the concept of a two-tier controller structure in which the context evaluation logic is decoupled from input-output triggering logic. The other key contribution was expected to be a snooper, i.e. a subsystem for tracking the switching behaviour on interfaces enabling on-line testing.

The project largely followed the original objectives of this project:

(a) Develop implementation architecture for a low latency controller with techniques for decomposition, synthesis and timing analysis.

(b) Develop the main supporting structures for off-line testing, such as internal scanning, for a class of stuck-at, bridging and delay faults with minimum impact on performance.

(c) Develop the detailed architecture for a snooper for on-line testing of verifiable self-timed structures with minimum time, area and power consumption overheads.

(d) Develop a case study with a demonstrator of the testable low-latency methodology; the application area will be an on-chip communication adaptor.

We believe the project has made significant advances in all these objectives. Important foundations for this project were laid in [a1-a3], the papers where synthesis experiments showed the indication of possible gains in terms of speed out of the use of direct-mapping techniques from STGs and Petri nets as opposed to conventional logic synthesis. While large part of this work on automatic synthesis continued under the BESST project (in terms of tool integration with Verilog and VHDL front ends, see: async.org.uk/besst), the emphasis of this project was particularly on developing latency-aware synthesis and testing techniques and applying them in various design examples. The project proceeded according to its plan with a few minor exceptions (indicated below).

2.1. Low-latency controller architectures and new synthesis tools (objective a)

Tracker-bouncer architecture in direct mapping methods. We developed the high-performance implementation architecture in which control is decomposed into the contextual state-holding unit, called "*tracker*", and output flip-flop unit, called "*bouncer*". The former operates in parallel with the environment thus minimising the critical path delay and the effect of test support components on the output latency. The latter uses only the minimum necessary context input from the tracker, which is pre-calculated in advance and generates new output values immediately in response to the changes on the (trigger) inputs. Initially proposed in [a2] the tracker-bouncer idea was developed further into methods and algorithms for direct mapping from STGs and implemented in a tool called Optimist [2, 20]. This technique allows the 'logical effort' expended on producing output signals to be estimated directly from the behavioural STG specification. The method involves a comprehensive set of latency-oriented optimisations at the level of the tracker model, various designs of state holding David cells, designs of bouncer flip-flops. A detailed comparison of this approach with controller synthesis methods based on Petrify has been made in [2,20] and shows an average latency improvement for most STG benchmarks in the order of 2X. This improvement has been mainly achieved due to the fact that Petrify resolves CSC by inserting internal state signals between inputs and outputs in Petrify while direct mapping avoids CSC solving altogether.

Logical effort analysis in latency-aware synthesis. Realising that the area penalty involved in the use of direct mapping may be too high a price for speed, we also investigated ways of improving the Petrify-like logic synthesis flow by introducing timing analysis into this flow and adding interactive element into synthesis at the stages of state conflict resolution [13]. This was motivated by the fact that Petrify had only a crude estimation of latency in the number of transitions between input and output signals in the controller. This is not acceptable for particularly high-speed controllers, such as glue logic in SoC. We developed a technique based on Logical Effort (normally used for optimal transistor sizing) for a much more accurate estimation of i/o path delays based on delay models of complex gates taking into account, loads, parasitics and signal slew rate. Sufficiently accurate path delays calculated in the new tool (using first-order arithmetic rather than performing complex simulations) allowed ranking to be performed for large sets of solutions within the CSC resolution procedure, thereby helping the selection of the best subsets of solutions taken further to the layout stage. The method was tested on a set of benchmarks and corroborated well with the SPECTRE simulations [19]. There is much scope here for further research in latency-aware synthesis, particularly into the statistical effects, such as i/o-path probabilities and parametric variability of transistors and interconnects.

Low-latency interface with datapath. While the main focus of STELLA was on control logic, techniques for interfacing control to datapath with minimum cost in latency have been investigated. With the applications targeted at communication adaptors, where most of the datapath operations were non-arithmetic manipulations of data received from multiple input ports and sent to multiple output ports, multiplexing and demultiplexing, decoding and encoding, compare and latch. Two examples of investigation of low latency paradigms were studied. One was more theoretical, investigating the effect of OR-causality in pipeline-like structures with slack [11]. This work provides a good modelling underpinning (it is presented in a sufficiently abstract, Petri net, based form) to datapath protocols such as weak-indication or early propagation, which are becoming widespread in the asynchronous design community. Some early propagation techniques were used in the design of a carry-look-ahead adder based on tree architecture (MSc project [26]). The other paradigm was designing logic for a self-timed duplex communication channel, where low latency is achieved in combining direct mapping of the control logic from the Petri net model and organising the push and pull handshakes with the datapath in such a way that the control actions are maximally out of the critical path of the data channel [1].

2.2. Off-line testing structures (objective b)

The fact that the structure of the circuit reflects the structure of its behavioural specification in the synthesis based on direct mapping has helped us to formulate the problem of developing built-in test structures of control logic from Petri net and STG models and solve it minimising the performance losses due to the introduction of circuitry for scan paths. This is presented in [9,10,15]. Our method derives scan architectures with associated tests at the circuit specification level. Circuit physical level faults in the David cells of the tracker (that implement Petri net places) have been analysed and mapped onto the Petri net specification. A "pseudo clock" is used to handle hazards and activate faults allowing the control logic be 100% testable for stuck at faults using appropriate testability features added to latches. An algorithm for inserting testability features has been developed and applied to a range of benchmarks. Tradeoffs between higher testability rates and higher area-time overheads have also been demonstrated. The impact of positioning testability features maximally in the tracker on the overall latency of the systems is minimal because

the tracker works in parallel with the environment. Testability logic delays in the bouncer have thus been minimised. Although the original plan contained research on adding off-line testability features for logic with relative timing, we felt it was more appropriate to focus on time-related testability structures under on-line testing (see below).

2.3. On-line testing techniques (objective c)

Much innovation has been achieved in the on-line testing of asynchronous circuits, the area where at the outset there had been very little known from prior research. Our investigation proceeded in a number of directions. One of them was the development of on-line monitors (snoopers) and fault checkers for handshake protocols. In its simpler form the checker was able to detect violations of 4-phase protocols (cf. refusal sets) and signalling them to a special error handling infrastructure. Different strategies and mechanisms were developed for the latter using early propagation and strong indication signalling methods [5, 14, 16]. In addition to detecting order violations timing errors, violating reasonable delay bounds (min and max) in handshake phases, have been made detectable [3].
Approaching on-line testing of asynchronous control logic from the handshake interfaces rather than by incorporating checkers directly into control logic was probably the right way to go as it had ultimately led us to quite an entirely new area of hardware design, both in synthesis and fault-diagnosis, which can be called *protocol-driven design*. This was a pleasant and unexpected surprise towards the end of the project, which gives plenty of opportunities for future research. The main idea of this approach is based on the use of a "protocol machine" (PM) as an initial specification of a protocol in systems built following the concept of communication-centric design. The PM is a 'hypothetical' automaton (possibly with concurrent actions) which is placed between the real communication entities, such as master(s) and slave(s), to constrain the possible actions of all entities in their communication. The PM technique has been first applied (manually, but in the future we consider its automation!) in designing controllers for the duplex communication channel [1]. But the PM approach has appeared to be more effective in developing non-invasive online checkers for interfaces [4], which is now studied in D. Koppad's PhD project.
The third aspect of the on-line testing research was the advancement of techniques for high-resolution (sub-10-picosecond) time measurement logic so as to facilitate delay fault testing and time-dependent monitoring of interfaces in SoCs [18]. In this, there was close interaction with the COHERENT project.

2.4. Other results (objectives a and d)

The project has involved several directions that were not explicitly planned in the proposal but naturally fell within its scope and led to important advances.
Concurrent error detection. Aspects of on-line and off-line testing were complemented by research on concurrent error detection (CED) techniques, particularly for transient and soft errors, thus forming broader research into methods for strengthening the reliability of asynchronous systems in general. They were particularly relevant to the aspect of low cost in latency for fault-tolerance features. In this work an asynchronous pipelined self-checking RISC processor architecture based on Dong's code has been developed [22,23]. Four pipelined architectures have been compared in [17] – two synchronous and two asynchronous with one of each type using CED. The results indicated an area overhead of 12% in return for a fault coverage of 98% of all unidirectional errors. Use of CED with check symbol prediction logic operating in parallel with the main logic guarantees minimum timing overhead. This research is under further development in the PhD project of M. Marshall [25].
Low-latency (speculative) synchronisation. Synchronization between independently clocked regions in a high performance system is often subject to latencies of more than one clock cycle. We have showed in [8] how the latency can be reduced significantly, typically to half the number of clock cycles required for high reliability, by speculating that a long synchronization time is not required. The small number of synchronization metastability times longer than normal are detected, and subsequent computations re-evaluated, thus maintaining the reliability of the system. These results are now developed further under the SYRINGE project.

GALS modelling and analysis. We have investigated a number of GALS architectures from the point of view of their behavioural correctness, performance and latency. The analysis techniques are based on Petri net modelling [7]. It has become clear that traditional architectures based on plausible and stretchable clocks suffer in terms of latency and throughput because ion one case a two clock period delay is added to the latency of the response in the asynchronous channel in every handshake transaction, whereas in the latter the clock pulse is stretched on every handshake action [29]. We considered other methods of controlling the local clock starting and stopping driven by the events from the asynchronous environment [6]. An automated verification and performance analysis tool is currently under development.

2.5. Chip fabrication

Experimental work has been slightly modified from what was planned initially. Firstly, it was possible to arrange that a demonstrator chip for time measurement was fabricated (using 0.2um CMOS process) through our collaboration with Sun Microsystems. As result of testing this chip at a much earlier stage of the project we were very keen to build on that and, since the COHERENT grant had already finished, we put funds from STELLA into a fabrication of a chip (in Summer of 2005), which included a special phase alignment logic (the absence of which had made it less effective to demonstrate our time-measurement system on the Sun chip). The cost of the chip fabrication was also shared with our SCREEN project, which benefited from having an AES engine to be placed on the same chip, and thereby saving extra funds for another chip experiment under SCREEN later in 2006.

2.6. Prototype tools

Three software tools have been developed: LESTA (Logical-Effort-based Static Timing Analysis) for latency-aware controller synthesis [13,19], ProtoDe for protocol-based decomposition of STGs for online testing [4,20], and TeLLA (offline Testing of Low Latency Asynchronous circuits) for extraction of tracker structures with built-in testability features [9]. These tools were interfaced to Cadence via structural Verilog format.

2.7. Research output

Overall, the project has resulted in contributions to 2 journal papers, 1 book chapter, 19 conference papers (14 of which are peer-refereed), 3 PhD theses by P. Hyde, A. Madalinski and D. Sokolov (plus one PhD thesis, by D.Koppad, is in write-up), 5 technical reports, and 3 tools.

3 Research Impact and Benefits to Society

Our research group at Newcastle considers this research to have been successful. Most of the initially planned tasks have been completed and in many cases the results exceeded the original expectation. For example, the major impact on the practical side has been in developing the high level design flow based on direct mapping of STGs for low latency circuits and built-in testability features for off-line testing and handshake protocol checkers for on-line testing. On the theoretical side, the major advances are in developing a method of synthesis of low latency controllers and on-line checkers from Protocol Machine specifications.

Publications, presentations of academic and visiting staff, RAs, PG students in our weekly seminar of the Asynchronous Systems Laboratory (see http://www.cs.ncl.ac.uk/events/ASL), have played an important role in promoting these results into a number of follow-up projects and external contacts.

The results of the project are relevant to potential users at a number of levels:

(1) The use of new testability features and on-line checkers will enable the designers of systems with asynchronous interfaces and controllers to incorporate facilities for at speed testing,

(2) The use of time measurement circuits will help the designers of FPGAs and SoCs to characterise their systems with synchronisers much more reliably. The speculative synchronisation method would help these designers also to reduce the latency of the conventional two-flop synchroniser.

(3) The use of new OR-causality models will benefit researchers in area of designing complex asynchronous and concurrent systems with higher performance (use of early evaluation techniques) in a wide range of industrial applications.

(4) The use of latency-aware synthesis and test algorithms and tools in asynchronous design flow will benefit CAD developers and vendors working on promoting asynchronous design techniques in microelectronics.

For example, these results will be useful to the UK research in asynchronous design, testing and dependable hardware, supported by EPSRC and EU grants. We are also planning to use these techniques in promoting asynchronous design at Atmel-UK under the SCREEN project.

4 Explanation of Expenditure

The expenditure plans in the original proposal have been followed without significant changes. The biggest change was the relocation of some of the consumables and salaries budget in order to cover the shortfall on travel and student maintenance. The latter was made possible thanks to some savings due to replacement of a senior RA (D. Shang) by a junior RA (A. Madalinski) at the last stage of the project. The project provided excellent opportunity for active involvement of several PG students (D. Sokolov, who was employed on another EPSRC grant, but helped STELLA project in the area of low latency synthesis, tools OptiMist, ProtoDE and modelling structures with OR-causality), D. Koppad – main work on testing, S. Dasgupta –PhD work on GALS and A. Mokhov – PhD work on timing analysis, P.Wang – MSc project on low latency CLA adders) and even UG students (J. Murphy, who is currently on our PhD list). With the involvement of a large number of research manpower in this project, it was inevitable that the demand for conference travel (DATE, IOLTS, ASYNC, ACSD, ACNS, ATPN, VLSI Design, PATMOS, DDECS, FMGALS, TestCom, ACiD-WG, UK Async Forum) was greater than it was anticipated in the proposal.

5 Further research and dissemination activities

Further research is being carried out on developing software tools for generation of BIST structures, online testing and concurrent error detection methods for self-timed circuits and circuits with synchronisers. The research outcome from this grant is directly applicable for exploitation in our three EPSRC grants, Secure circuit design (SCREEN, GR/S81421), in collaboration with Atmel Smart Cards UK, Next Generation of Interconnection Technology For Multiprocessor SoC (EP/C512812/1) in collaboration with Southampton University and MBDA UK Ltd, Synchronizer Reliability in the Next Generation of SoC with Multiple Clocks (SYRINGE, EP/C007298/1) in collaboration with Intel, and most recently Self-timed Datapath Synthesis (SEDATE, EP/D053064/1) in collaboration with Universities of Manchester and Edinburgh and two companies, FTL Systems and Silistix. The latter contains a significant portion of research programme on built-in testability.

The ideas of latency-aware synthesis and testing of asynchronous logic as well as analysis of timing conditions and testability and software error are now relevant to the PhD research of D. Koppad, J. Murphy, S. Dasgupta, C. D'Alessandro, A. Mokhov, P. Wang, N. Minas and M. Marshall.

Besides publications, this research has had its impact on Adv. Tutorial on Hardware Design and Petri Nets at ATPN in Miami, June 2005, invited lectures at the fourth advanced course on Petri nets in Eichstaett in September 2003, ACiD-WG Winter school in Cambridge in Jan. 2005, presentation of design tools at DATE 2005 in Munich, series of invited lectures at the Xidian University, Xi'an, China and Institute of Electronics, Chinese Academy of Sciences, Beijing in March-April 2006 It has also been put forward in our presentations at the recent conferences on Applications of Concurrency in Systems Design (ACSD'06) and ATPN'06, DATE'06, ASYNC'06. The team members have also co-organised two UK Embedded Forums [27,28], and Newcastle will host the 18th UK Async Forum in September and 3rd UK Embedded Forum in October 2006.

Project publications

[1] A. Yakovlev, S. Furber, R. Krenz, A. Bystrov. Design and Analysis of a Self-timed Duplex Communication System, IEEE Transactions on Computers, Vol. 53, No.7, pp. 798-814, July 2004.
[2] D. Sokolov and A. Yakovlev, Clockless circuits and system synthesis, IEE Proceedings, Computers and Digital Techniques, vol.152, No.3, May 2005, pp. 298-316 (this work was also presented as a book chapter: D. Sokolov and A. Yakovlev, Clock-less circuits and system synthesis, Chapter 16 in B. Al-Hashimi (Ed.), System On Chip: Next Generation Electronics, IEE Circuits Devices and Systems Book Series, 2006, pp. 541-586**)**
[3] D. Shang, A. Yakovlev, F. Burns, F. Xia and A.Bystrov, Low-cost online testing of asynchronous handshakes, accepted by European Test Symposium, Southampton, May 2006, IEEE CS Press, pp. 225-230.
[4] D. Koppad, D. Sokolov, A.Bystrov, A. Yakovlev, Online testing by protocol decomposition, accepted by IOLTS'06, Como, Italy, July 2006.
[5] D. Shang, A. Bystrov, A. Yakovlev and D. Koppad, On-line testing of globally asynchronous circuits, Proc. 11th International Online Testing Symposium (IOLTS'05), St. Rafael, France, July 2005, IEEE CS Press, pp. 135-140.
[6] S. Dasgupta, D. Potop-Butucaru, B. Caillaud, A. Yakovlev. Moving from Weakly Endochronous Systems to Delay-Insensitive Circuits, Proceedings of the Second Workshop on Globally Asynchronous, Locally Synchronous Design (FMGALS 2005), Electronic Notes in Theor. Comp. Sci., Vol. 146, No.2, January 2006, pp 81-103
[7] S. Dasgupta and A. Yakovlev, Modelling and verification of globally asynchronous and locally synchronous ring architectures, DATE 2005, Munich, March 2005, pp. 568-569.
[8] D.J.Kinniment and A.V.Yakovlev. Low latency synchronisation through speculation, PATMOS 2004, Santorini, September 2004, LNCS 3254, Springer, pp. 278-288.
[9] D. Koppad, A. Bystrov and A. Yakovlev, Off-line testing of asynchronous circuits, 18th Int. Conf. on VLSI Design, IEEE CS Press, Kolkata, Jan. 3-7, 2005, pp. 730-735
[10] D. Koppad, A. Bystrov and A. Yakovlev. Algorithm for Testing Asynchronous circuits, 16th IFIP Int. Conf. on Testing of Communicating Systems (TestCom 2004), St Anne's College, Oxford, UK, 17-19 March 2004.
[11] A. Bystrov, D. Sokolov, and A. Yakovlev. Low latency control structures with slack, Proc. of Int. Symp. on Adv. Res. in Asynchronous Systems and Circuits (ASYNC'03), May 2003, Vancouver, IEEE CS Press, pp. 164-173.
[12] D. Koppad, A. Bystrov and A. Yakovlev, On-line Monitoring of Asynchronous Interfaces, 17th UK Asynchronous Forum, Southampton, September 2005.
[13] A. Madalinski, C. D'Alessandro, P. Wang and A. Yakovlev, Latency aware Conflict Resolution in Asynchronous Control Logic Synthesis, 17th UK Asynchronous Forum, Southampton, September 2005.

[14] A. Bystrov, D. Koppad, A. Yakovlev, On-line testing of asynchronous circuits, Workshop of ACiD-WG Working Group, Turku, Finland, June 2004.
[15] D. Koppad, A. Bystrov and A.Yakovlev, Algorithm for testing asynchronous circuits, Proc. 15th UK Asynchronous Forum, Cambridge Jan. 2004.
[16] D. Koppad, D. Shang, A. Bystrov, A. Yakovlev, Asynchronous Checker Designs for Monitoring Handshake Interfaces, Tech Report NCL-EECE-MSD-TR-2005-104, Microelectronic System Design Group, School of EECE, University of Newcastle upon Tyne, March 2005.
[17] P. D. Hyde and G. Russell. A comparative study of the design of synchronous and asynchronous self-checking RISC processors. Proc. 10th IEEE Int. On-Line Testing Symposium (IOLTS'04). IEEE CS Press 2004, pp. 89-94
[18] M. A. Abas, G. Russell and D. J. Kinniment. Design of sub-10-picoseconds on-chip time measurement circuit. In: Proceedings DATE'04, IEEE Comp. Soc. Part Vol.2, 2004.
[19] C. D'Alessandro, A. Madalinski, P. Wang and A. Yakovlev, Analysis of Logical Effort in Latency-Aware Synthesis of Asynchronous Control Circuits, submitted to ICCD 2006 (a preliminary version of this paper was presented by C. D'Alessandro at Cadence Labs UC Berkeley in March 2006, see: http://www.cadence.com/company/cadence_labs/upcoming_talks.aspx).
[20] D. Sokolov, A. Bystrov, A. Yakovlev, Direct mapping of low-latency asynchronous controllers from STGs, Tech Report NCL-EECE-MSD-TR-2006-110, Microelectronic System Design Group, School of EECE, University of Newcastle upon Tyne, January 2006 (submitted to IEEE Trans.on CAD).
[21] A. Bystrov and J.P. Murphy, On-line IDDQ testing of security circuits, NCL-EECE-MSD-TR-2004-103, Microelectronic System Design Group, School of EECE, University of Newcastle upon Tyne, September 2004.
[22] P. D. Hyde and G. Russell. ASSEC: an asynchronous self-checking RISC-based processor, Proceedings of the Euromicro Symposium on Digital System Design. IEEE Comput. Soc. 2004, pp. 104-111.
[23] P. Hyde. A Pipelined Asynchronous Self-Checking RISC-Based Processor, PhD Th., Univ. of Newcastle, 2004.
[24] D. Sokolov. Automated synthesis of asynchronous circuits using direct mapping for control and data paths, PhD thesis, Univ. of Newcastle, 2006 (see also Tech Report NCL-EECE-MSD-TR-2005-111 at async.org.uk).
[25] M.J Marshall and G. Russell, An information-redundant scheme for on-line testing of an asynchronous ALU operation, ETS'06 Poster presentation, Southampton, May 2006, pp. 141-145.
[26] P. Wang, A 16-bit self-timed look-ahead adder, Proc. DDECS'04, Stara Lesna, Slovakia, April 2004.
[27] A. Koelmans, A. Bystrov, M. Pont (Eds). First UK Embedded Forum, Proceedings, NEC, Birmingham, October 2004. University of Newcastle upon Tyne, ISBN 0-7017-0180-3 (e-book, http://www.staff.ncl.ac.uk/albert.koelmans/books/firstukembforum.pdf).
[28] A. Koelmans, A. Bystrov, M. Pont, R. Ong, A. Brown (Eds). Second UK Embedded Forum, Proceedings, NEC, Birmingham, October 2005, University of Newcastle upon Tyne, ISBN 0-7017-0191-9 (e-book, http://www.staff.ncl.ac.uk/albert.koelmans/books/secondukembforum.pdf).
[29] S. Dasgupta, A. Yakovlev, Modelling and Performance Analysis of GALS Architectures, Tech. Report NCL-EECE-MSD-TR-2005-114 (see async.org.uk), MSD Group, School of EECE, University of Newcastle, April 2006.
[30] A. Madalinski. Interactive Synthesis of Asynchronous Systems based on Partial Order Semantics, PhD Thesis, Tech Report NCL- EECE-MSD-TR-2005-112 (see async.org.uk), MSD Group, School of EECE, University of Newcastle, February 2006.

Other relevant papers

[a1] A. Bystrov, A. Yakovlev, Synthesis of Asynchronous Circuits with Predictable Latency, IEEE/ACM 11th International Workshop on Logic and Synthesis, New Orleans, June 2002, IEEE Computer Society and ACM SIGDA, pp. 239-243.
[a2] A. Bystrov and A. Yakovlev. Asynchronous Circuit Synthesis by Direct Mapping: Interfacing to Environment, Proc ASYNC'02, Manchester, April 2002, IEEE CS Press, 127-136.
[a3] D. Shang, F. Xia and A. Yakovlev, Asynchronous Circuit Synthesis via Direct Translation, Proc. ISCAS'02, Scottsdale, Arizona, May 2002, IEEE, Volume III, pp. 369-372