



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/85856/>

Version: Published Version

Article:

Trefzer, Martin Albrecht, Walker, James Alfred, Bale, Simon Jonathan et al. (2015) Fighting stochastic variability in a D-type flip-flop with transistor-level reconfiguration. IET Computers and Digital Techniques. pp. 190-196. ISSN: 1751-861X

<https://doi.org/10.1049/iet-cdt.2014.0146>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Fighting stochastic variability in a D-type flip-flop with transistor-level reconfiguration

ISSN 1751-8601

Received on 18th July 2014

Revised on 23rd October 2014

Accepted on 31st October 2014

doi: 10.1049/iet-cdt.2014.0146

www.ietdl.org

Martin A. Trefzer , James A. Walker, Simon J. Bale, Andy M. Tyrrell

Department of Electronics, University of York, York YO10 5DD, UK

 E-mail: martin.trefzer@york.ac.uk

Abstract: In this study, the authors present a design optimisation case study of D-type flip-flop timing characteristics that are degraded as a result of intrinsic stochastic variability in a 25 nm technology process. What makes this work unique is that the design is mapped onto a multi-reconfigurable architecture, which is, like a field programmable gate array (FPGA), configurable at the gate level but can then be optimised using transistor level configuration options that are additionally built into the architecture. While a hardware VLSI prototype of this architecture is currently being fabricated, the results presented here are obtained from a virtual prototype implemented in SPICE using statistically enhanced 25 nm high performance metal gate MOSFET compact models from gold standard simulations for pre-fabrication verification. A D-type flip-flop is chosen as a benchmark in this study, and it is shown that timing characteristics that are degraded because of stochastic variability can be recovered and improved. This study highlights significant potential of the programmable analogue and digital array architecture to represent a next-generation FPGA architecture that can recover yield using post-fabrication transistor-level optimisation in addition to adjusting the operating point of mapped designs.

1 Introduction

Over the last 20 years, field programmable gate arrays (FPGAs) have rapidly improved in performance and function density enabled by the continuous shrinking of technology sizes. However, device sizes have now approached atomistic scales where the presence or absence of single doping atoms and structural irregularities become more prevalent. Owing to this, the characteristics and behaviour of single devices, and therefore the circuits built with them, are altered in a random fashion. As a consequence, time-consuming statistical simulation program with integrated circuit emphasis (SPICE) simulations using specific, statistically enhanced device models become necessary in order to accurately model and create reliable electronic designs that behave according to specification. Unfortunately, because of the statistical nature of the variations, the fabrication yield still decreases and failure rates increase significantly, because every physical instance of a design behaves in a stochastically different manner [1–3]. Even when verifying designs using accurate device models and statistical SPICE simulation, this will in the first instance only allow for a more accurate yield prediction, rather than provide a means to overcome the effects of random variability in the physical devices fabricated.

Therefore variability must be addressed at all stages of the design flow: during design, physical implementation and post-fabrication. At the design stage, the availability of good quality/accurate statistically enhanced device models is essential in order to be able to predict the effects of stochastic variations on a design and to take appropriate counter measures. This work focuses on digital reconfigurable devices, particularly on enhancing FPGA architectures. The proposed architecture provides an additional analogue level of reconfiguration that allows on-line performance optimisation of designs mapped at the digital level. In this work, a D-type flip-flop (DFF) with degraded timing characteristics because of intrinsic stochastic variability in a 25 nm technology process is presented as an optimisation case study using the proposed architecture and methodology.

2 Background and rationale

In the case of CMOS transistor designs, optimising device sizes and selecting appropriate topologies are methods used to tackle variability [4–6]. The greatest workload and responsibility remains to date with chip manufacturers, who continuously improve fabrication facilities and feed-back appropriate design rules for creating the physical layout to the designers in order to ensure high yield figures. Moreover, new devices and technologies are continuously being developed and refined in order to further advance technology. For example, silicon-on-insulator and FinFET transistors, which work with undoped channels thereby eliminating one major cause of variability [7]. There are also a number of post-fabrication measures to improve the performance of a device or make it at least usable with reduced performance. For instance, altering power-supply voltages, slowing down clock-speed or disabling (redundant) parts. These are at present the predominantly commercially driven post-fabrication counter measures, known as ‘binning’, which allow more devices with different, appropriately guaranteed performance to be sold.

In contrast to these methodologies, there are examples where reconfigurable architectures are used for post-fabrication optimisation and fault tolerance [8–13]. Following on from these examples, this work is based on including additional reconfiguration mechanisms in the design of an architecture operating at the analogue level, which allow for alterations of the characteristics of devices and components once they are fabricated and during operation. This provides an access point for optimisation algorithms to find configurations that may improve the circuit’s performance and bring it back into specification. Although introducing (additional) configuration options into a design generates area overhead, there will be an overall benefit allowing continued use of parts of the device that otherwise would have to be disabled because they do not work according to specification, or even worse, not being able to use the whole device.

In particular, this work focuses on enhancing FPGA architectures for the following reasons: first, FPGAs are widely used in

applications where on-line reconfigurable signal processing is required. Current devices feature high logic densities and programmable application-specific macro-blocks, that is, multipliers, ALUs, and can therefore be configured to implement customised digital systems comprising of processors, peripherals and high-density logic, which places them between microprocessors and ASICs. Their versatility and the fact that they incorporate reconfiguration options already make them suitable candidates for the proposed research. Second, the design most affected by intrinsic variability has been SRAM [14, 15]. Since SRAM is mainly used for storing configuration data and look-up tables in reconfigurable devices, and hence can be operated at relatively low speeds compared with the actual applications, FPGA fabric has not been as severely affected by this kind of intrinsic variability issues as other ASICs like memory and processors in the past. However, it is projected that the next ‘victim’ of variability after SRAM will be latches, and this will have a direct impact on FPGA architectures, which consist of a large number of flip-flops of which latches are an essential part. Therefore optimisation of a programmable DFF, implemented on the programmable analogue and digital array architecture (PAnDA) [16, 17], is chosen as a case study in this paper. A multi-objective evolutionary algorithm (MOEA) is used to find configurations that demonstrate the best trade-off performance with regard to delay, setup time, hold time and dynamic power consumption. This allows the selection of the most appropriate setting for a given application. In its role as one of the fundamental sequential building blocks, it is not always only desirable to minimise these performance metrics, but to match, for example, timings of two design parts that feed into a third, which makes optimisation of sequential components a challenging task. Particularly when they are additionally degraded because of stochastic variability.

3 Reconfigurability as a tool

The use of overcoming the effects of intrinsic variability via optimising transistor sizes and how this principle can be manifested as an online optimisation tool using reconfigurable hardware is discussed in the following sections. Examples of previous work where post-fabrication optimisation has been shown to be beneficial in terms of yield, fault tolerance and/or performance can be found in [9–12, 18]. Of course, adding additional reconfiguration options to a hardware architecture always results in increased overhead and the benefits must be weighed against that. However, what makes the proposed PAnDA architecture, briefly described in Section 3.3, unique is that it offers the possibility to combine and optimise two things at the same time that are often considered separately, which is optimising for a desired operating point (the mean) and for variability (the spread of the performance distribution). This paper is mostly concerned about optimising the operating points of DFFs that are fabricated on different dies and are shifted because of intrinsic variability, rather than minimising the spread of the distributions.

3.1 Causes of variability in CMOS design

Intrinsic variability is caused by differences at the atomic scale in devices that could be considered macroscopically identical in terms of their layout [19], construction and local environment. The main sources of intrinsic variability are random dopant fluctuations (RDF) [20, 21], line edge roughness [2], variability of gate oxide thickness [22] and poly-silicon grain boundary variability [23]. In current channel lengths of above 30 nm, RDF has by far the greatest effect on device variability [2]. The impact of other types of variability in future nodes will depend on the specific technology used and improvements that can be made in the lithography and etching processes. For example, advances have been made to reduce the loss of precision caused by the

manufacturing process (e.g. optical proximity correction [24], uniformly dense layout [25]).

3.2 Design optimisation via transistor sizing

The results in [4, 26], which have been obtained from statistical SPICE simulations, suggest that optimising the widths of transistors in standard cells can improve their variability tolerance, speed and power consumption. It is also shown that it is possible to design and optimise analogue CMOS circuits in hardware using field programmable transistor arrays (FPTAs) [12, 27], and there are examples where transistor-level reconfiguration is used as a mechanism for design optimisation [5, 6]. Therefore if FPTA-based mechanisms to alter device sizes are incorporated in a hardware architecture, it will be possible to optimise circuit designs post-fabrication, that is, adapt them in such a way that they perform optimally on the silicon die they are fabricated on. This would not only have the advantage of being able to enhance variability tolerance and performance for a specific design, but could also account for variations between different devices. Moreover, because of the large numbers of statistical measurements necessary for characterisation, it will be orders of magnitude faster to perform optimisation directly in hardware – which is what is proposed here – rather than in SPICE simulation.

In other words, the configuration bits for the transistor sizes are user configurable. In practice, there may be a preset during device initialisation (similar to FPGA init). Since every device will be different as a result of intrinsic variability and the fact that the user cannot know which device size and device combination will be optimal for a specific device or design mapped, we propose post-fabrication/post-mapping optimisation that can be performed online.

3.3 PAnDA reconfigurable architecture

PAnDA is a novel FPGA architecture which aims to overcome challenges arising when shrinking device sizes to the nano-scale as well as providing more reliability and performance through optimising built-in configuration settings that allow modifying circuit characteristics. At the post-fabrication stage it is generally no longer possible to modify device sizes or the topology of a design, although these techniques have been proven to be useful at the design stage. With the reconfigurable PAnDA architecture, however, we aim to make this possible by providing reconfiguration options at the transistor level, which effectively allow us to change the sizes of any transistor that is part of a mapped design. This is achieved by replacing each transistor with a number of them connected in parallel, of which any one can be either turned-on or shut-off depending on its associated configuration bit. As a consequence, this group of transistors behaves like a single device of which the size can be altered by turning different subsets on. These configurable transistors (CTs), shown in Fig. 1, are the core low-level building blocks of PAnDA. As can be seen from Fig. 1, device size changes are constrained to transistor widths in the first versions of the PAnDA chip. At higher levels of the design hierarchy the PAnDA architecture CTs form configurable analogue blocks (CABs), configurable logic blocks (CLBs) and logic cells. CLBs, logic elements and logic cells are also present – and have similar structure – in current commercial FPGA architectures, whereas CTs and CABs, which offer reconfigurability at the analogue level, are unique to PAnDA. The inspiration for CTs originates from FPTAs, introduced in [12, 28]. The designs of the PAnDA CTs and CABs are described in detail in [16, 17]. The CT implementation shown in Fig. 1 comprises of seven native transistors, which means that seven configuration bits are necessary to realise all 128 combinations for configuring the width.




Fig. 1 CT from the PAnDA architecture is shown

Effective width of the CT can be changed via enabling different subsets of MOSFETS. The switches are implemented as transmission gates and they are included in the variability simulations performed in this work. The widths of the seven native transistors within the CT are $M_0, \dots, M_6 = 120, 120, 140, 160, 180, 200, 220$ nm

3.4 Optimisation algorithm

In this work, we use a MOEA for design optimisation, which is based on NSGA-II [29]. The algorithm is implemented in C++ and a demo is available at http://www-users.york.ac.uk/~540/downloads/easpice_demo.tar.gz. The non-dominated sorting algorithm, crowding distance and selection scheme are the same as in NSGA-II. More detailed examples for setting-up a MOEA for optimisation of transistor sizes are described in [4, 26, 30]. In this case, a direct encoding of the configuration bits of all CTs required to build the DFF is used as the genome, that is, a bit string of $22 \text{ CTs} \times 7 \text{ bits} = 154$ bits. The objective values are measured in picoseconds and Watts using a SPICE testbench. In every generation (iteration of the algorithm), a netlist is generated for each candidate solution based on a netlist template comprising the measure statements, the circuit and the statistical device models. The CTs are configured using the information stored in the genome resulting in a number of netlists for the same VPT (see Section 4.3), but configured with different CT sizes. A population size of 80 is used, that is, 80 new netlists are generated per generation, which corresponds to the number of cluster nodes currently available to us for parallel SPICE runs in order to make maximum use of computing resources. Since the goal is to optimise a number of configuration bits for the CTs, a mutation rate of 1 bit per candidate solution is chosen. Results shown are

obtained after running the MOEA for about 100 generations, which corresponds to ~ 8000 SPICE runs.

4 Optimisation of a DFF

A DFF is one of the fundamental building blocks of digital design, in particular they are widely used in FPGAs, for pipelining in digital signal processing and register files in microprocessors. The timing characteristics of DFFs have a major impact on the maximum achievable clock speed of a digital design. Since DFFs comprise two latches, which are highly sensitive to the effects of variability, they are chosen as a case study for the post-fabrication optimisation approach proposed here. As mentioned, the reconfigurable PAnDA architecture in combination with the MOEA offers the possibility to optimise the operating point (mean) as well as robustness against variability (distribution spread). This paper investigates the potential use of the PAnDA architecture for post-fabrication yield recovery in the case of degraded circuit performance because of the variability present on a specific die. Hence, the focus is on optimising the operating points of specific physical instances affected by intrinsic variability rather than on producing a generic design that exhibits increased resilience against variability on all virtual dies. Experiments are carried out in SPICE simulation using statistically enhanced 25 nm high-performance metal gate bulk MOSFET compact models from gold standard simulations (GSS) in order to assess the effects of stochastic variability present in deep sub-micron technology nodes. PAnDA prototype chips are currently being fabricated. Once they become available later in 2014, it will be possible to verify the results presented in hardware.




Fig. 2 Performance characteristics of a DFF are shown and the grey area highlights the time period used for integration of dynamic power consumption, t_{delay} corresponds to the clock-to-q delay whereas t_{setup} and t_{hold} represent setup time and hold time, respectively

4.1 DFF performance measures

DFFs are characterised by four performance characteristics: clock-to-q delay, setup time, hold time and (dynamic) power consumption. The functional behaviour of a DFF and its performance characteristics are pictured in Fig. 2. All measurements are taken at the points where the signals cross the 50% supply voltage (V_{DD}) mark, in this case $V_{DD} = 1$ V. Clock-to-q delay is measured between the rising edge of the clock and the next change of the output signal. Setup time is the point in time when the correct data need to be present (and stable) at the




Fig. 3 DFF implemented using the CTs from Fig. 1 is shown

Note that the symbol for CTs is a square with an arrow denoting PMOS and NMOS, respectively, in order to highlight the fact that they are not single transistors. The DFF design used in this work is a standard design that has been adapted for using CTs from [31], Chapter 10

input before the next rising clock edge (the DFF is positive edge sensitive) in order to be successfully latched. Hold time is the period of time where the data must not change in order to guarantee that the DFF latches and stores the correct value. Note that hold times may be negative (which means good performance).

4.2 Mapping to the PAnDA architecture

Since the goal is to investigate the use of the transistor-level reconfiguration capabilities of the PAnDA architecture in order to optimise the DFF, the first step required is to map a standard DFF design onto the PAnDA architecture. This is shown in Fig. 3, where each transistor of the DFF is represented by a CT, shown in Fig. 1. This allows the transistor widths to be altered within the range and granularity given by the configuration options of the PAnDA CTs.

4.3 Virtual physical instances (VPIs)

This work uses SPICE simulation to investigate the potential use of the PAnDA architecture for post-fabrication optimisation, hence, we assume that the compact models used capture the effects of stochastic variability with sufficient accuracy to make realistic performance predictions of designs when fabricated. The GSS model generator processes SPICE netlists in such a way that each device is assigned its own, statistically enhanced model card, that is, the netlist reflects effects of one specific scenario of intrinsic stochastic variability after processing. In order to fully assess how a design is affected by variability, this process needs to be repeated many times in order to generate a large number (ideally thousands) of different netlists representing different points of a distribution of resulting design characteristics.

In this case, netlists generated – each representing a different scenario of a DFF design affected by variability – are regarded as possible physical implementations of a design. This allows us to infer post-fabrication performance from the simulation results and make predictions as to whether the optimisation process can recover (or improve) the performance of a design. Each unique randomised netlist of the DFF is therefore referred to as ‘virtual physical instance (VPI)’.

As a starting point for the experiments, 20 000 VPIs of the DFF are generated and delay, setup time, hold time and dynamic power consumption of all of them are measured using the same SPICE testbench. The results are shown as scatter plot point clouds in Figs. 4a and b. The density, shape and spread of the point clouds illustrate the resulting variability distributions from the 20 000 simulation runs with regard to delay and power, and with regard to setup and hold times, respectively. The statistical outliers that exhibit worst performance both in terms of power consumption and clock-to-q delay are highlighted in the figures (5656, 10 270, 16 294) and are selected as candidates for optimisation. Note that all four performance measures are measured and subsequently optimised at the same time, but it is easier to visualise the results in separate two-dimensional plots.

4.4 Optimisation results

Three VPIs with an overall degraded performance are selected in order to investigate whether it is possible to recover and/or improve performance through optimising CT configurations, that is, changing the effective widths of the CTs through reconfiguration. Note that during the optimisation process only the configuration of the CTs is optimised, no other design parameters such as, for instance, model parameters are altered since this




Fig. 4 Resulting performance characteristics of 20 000 virtual physical implementations of a DFF are shown in the figure

Statistical outliers are highlighted with red circles. The dashed lines indicate the mean values

a Delay against dynamic power consumption

b Setup against hold time




Fig. 5 Designs resulting from multi-objective optimisation of all three virtual prototypes are shown

Results highlighted with a square exhibit better performance in all metrics (delay, setup, hold time) except dynamic power consumption. The straight dashed lines highlight the mean values and the ones connecting measuring points help visualising the Pareto-fronts
a Delay against dynamic power consumption
b Setup against hold time

Table 1 Summary of resulting performance characteristics taken from the solutions, marked with a square in Figs. 5*a* and *b*, with the best overall trade-off performances are listed

Characteristic	5656		10 270		16 294	
	Orig.	Opt.	Orig.	Opt.	Orig.	Opt.
clock-to-q, ps	185	126	169	109	150	103
setup time, ps	40	0.5	41	25	45	35
hold time, ps	15	5	11	-11	11.5	-19.5
dynamic power, μ W	295	319	308	384	323	326

would invalidate a virtual physical instance. All three designs chosen exhibit worst-case performance in terms of power consumption and/or delay (statistical outliers).

The optimisation results are shown in Figs. 5*a* and *b*; and the results are also summarised in Tables 1 and 2. As can be seen from the figures, the multi-objective optimisation algorithm yields a population of optimised designs with different performance trade-offs. The subset of solutions where each solution is better than any other in at least one objective is denoted as the first non-dominated front (NDF). The NDF is also called a Pareto-optimal set (Pareto-front) of solutions, because there is no overall best solution, but rather the designer can make trade-off choices within this set. Since, in practice, the designs with the best overall performance are the most generally useful ones, those featuring shorter delay and improved setup and hold times at the least additional expense of power are highlighted with squares. As can be seen from the table and from the figures, multi-objective

optimisation yields designs that are significantly improved in three objectives (delay, setup and hold times) in the case of all three VPIs chosen at the expense of slightly increased power consumption. In two cases (5656, 16 294), the increase in power consumption is <10% and in one case (10 270) it is about 25%. Note that there are solutions present in the resulting Pareto-optimal sets that consume less power, but at a considerable expense of speed. This reflects the expected trade-off between power consumption and performance, illustrated in Figs. 5*a* and *b*: in the case of VPT 5656, where speed is most severely affected by variability, this can generally only be improved by increasing the size of certain devices or by enabling different subsets of devices within a CT that result in a similar effective size, thereby increasing power consumption or keeping it constant. In the case of VPT 16 294, this extreme situation is reversed. There is room for making the devices smaller, hence reducing power and improving speed, until again a certain trade-off point is reached where the speed can only be further improved at the cost of increased power consumption.

The initial transistor widths of the DFF are compared with those of the optimised designs in Table 2. In most cases, the transistors have become bigger, but there are some that have been made smaller as a result of the optimisation process. However, a quantitative analysis of how transistor sizes change subject to optimisation requires a significant amount of additional experiment runs beyond the scope of this paper, hence, will be subject to future work.

An additional experiment has been conducted where all 20 000 VPIs are configured with the optimised configurations of the three example solutions chosen in order to investigate how generally applicable they are. Note that this was not explicitly included as

Table 2 Initial transistor size settings of all CTs of the DFF before optimisation are compared with those of the three selected designs (5656, 10 270, 16 294) after optimisation

Transistor	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11
initial setting	200	160	160	160	160	160	160	160	160	160	160
optimised 5656	280	440	560	460	600	420	240	140	120	660	400
optimised 10 270	620	600	260	300	380	120	400	120	320	140	540
optimised 16 294	600	160	140	440	560	340	120	300	280	260	160
Transistor	M12	M13	M14	M15	M16	M17	M18	M19	M20	M21	M22
initial setting	160	160	160	160	160	160	160	160	160	280	220
optimised 5656	320	460	300	340	780	340	340	600	620	240	280
optimised 10 270	120	180	360	320	240	700	640	660	140	120	380
optimised 16 294	440	300	600	120	520	760	460	300	340	140	580




Fig. 6 Performance of all 20 000 virtual prototypes configured with the optimised configurations from Figs. 5a and b, highlighted with a square, are shown

Straight dashed lines highlight the mean values of the respective clouds

a Delay against dynamic power consumption

b Setup against hold time

an optimisation objective. The resulting point clouds are shown in Figs. 6a and b. It is observed that the entire cloud is shifted in the same direction. In the case of delay and power the shape remains similar, whereas for setup and hold times the clouds resulting from VPIs 5656 and 16 294 are highly skewed and only the one from VPI 10 270 retains its shape. Note that VPIs 5656 and 16 294 are worst-case outliers in delay and power, respectively, which may explain this. However, in the case of VPI 16 294, a large number of simulations failed because of a limit set in the SPICE testbench. In some cases, the spread of the entire cloud is smaller, although the reason for this is likely to be the increased power consumption when larger devices are enabled.

5 Conclusions

This paper has investigated the application of multi-objective evolutionary optimisation on reconfigurable hardware for recovering/improving the performance of a DFF mapped onto it that is degraded because of stochastic variability. There are two novel aspects to this work: first, the novel reconfigurable PANDA architecture has been used to implement the DFF design. PANDA is a hierarchical architecture, comprising of CTs, CABs, CLBs and interconnect. At the CLB and interconnect level, PANDA is compatible with commercial FPGA architectures. However, PANDA offers additional lower levels of reconfiguration (CAB and CT levels), which allows the optimisation of electronic designs at a smaller granularity. The lowest analogue level is represented by the CTs, which are used in the case study presented here. Second, multi-objective evolutionary optimisation has been successfully applied – working at the analogue reconfiguration level of PANDA – to recover and optimise the performance of a DFF where the performance was degraded because of stochastic variability. It has been shown that timing can be significantly improved in exchange of a relatively small increase in power consumption. The results suggest that this kind of multi-reconfigurable architecture, which allows the optimisation performance at both the analogue and the digital level, has great potential to enhance current standard field-programmable digital devices, such as FPGAs, with post-fabrication optimisation capabilities. Note that optimisation goals can be defined by the user and thus include manipulating a circuit for a desired operating point and recovering yield as well as increasing robustness against silicon substrate variations, which makes the approach highly flexible.

The sequential case study presented here compliments our previous work on optimising combinational circuits [16]. The DFF chosen represents one of the fundamental building blocks of

digital design. Therefore the results shown are generally relevant to digital design with FPGAs and the results obtained will feed into subsequent PANDA hardware prototypes that are currently being designed and fabricated. The PANDA architecture will close the gap between the analogue design of standard cells and the design of reconfigurable digital systems based on standard cell libraries, by providing a design platform that allows the mapping of logic designs and then optimise them in multiple stages at runtime through reconfiguration at the different levels. This is currently not possible with any existing commercial FPGA. Future work will verify the results of this paper in hardware, once the PANDA silicon is available.

6 Acknowledgments

This work is part of the PANDA project that is funded by EPSRC (EP/I005838/1) and is the subject of a UK patent application (GB1119099.8).

7 References

- Bernstein, K., Frank, D., Gattiker, A., *et al.*: 'High-performance CMOS variability in the 65-nm regime and beyond', *IBM J. Res. Dev.*, 2010, **50**, (4.5), pp. 433–449
- Asenov, A.: 'Variability in the next generation CMOS technologies and impact on design'. Proc. of the First Int. Conf. of CMOS Variability, 2007
- Borkar, S., Karnik, T., Narendra, S., Tschanz, J., Keshavarzi, A., De, V.: 'Parameter variations and impact on circuits and microarchitecture'. Proc. of the 40th Annual Design Automation Conf. (DAC), 2003, pp. 338–342
- Walker, J.A., Hilder, J.A., Reid, D., *et al.*: 'The evolution of standard cell libraries for future technology nodes', *Genet. Program. Evol. Mach.*, 2011, **12**, (3), pp. 235–256
- Ali, S., Ke, L., Wilcock, R., Wilson, P.: 'Improved performance and variation modelling for hierarchical-based optimisation of analogue integrated circuits'. Design Automation and Test in Europe (DATE), April 2009, pp. 712–717
- Zheng, R., Suh, J., Xu, C., Hakim, N., Bakaloglu, B., Cao, Y.: 'Programmable analog device array (PANDA): a platform for transistor-level analog reconfigurability'. Design Automation Conf. (DAC), 2011
- Cheng, B., Wang, X., Brown, A.R., Kuang, J.B., Nassif, S., Asenov, A.: 'Transistor and SRAM co-design considerations in a 14 nm SOI FinFET technology node'. Proc. of the Int. Electron Devices Meeting (IEDM), San Francisco, CA, USA, 2012
- Stott, E., Sedcole, P., Cheung, P.: 'Fault tolerance and reliability in field-programmable gate arrays', *IET Comput. Digit. Tech.*, 2010, **4**, (3), pp. 196–210
- Takahashi, E., Kasai, Y., Murakawa, M., Higuchi, T.: 'A post-silicon clock timing adjustment using genetic algorithms'. Symp. on VLSI Circuits, 2003, pp. 13–16
- Murakawa, M., Adachi, T., Niino, Y., *et al.*: 'An AI-calibrated IF filter: a yield enhancement method with area and power dissipation reductions', *IEEE J. Solid-State Circuits*, 2003, **38**, (3), pp. 495–502

- 11 Takahashi, E., Murakawa, M., Kasai, Y., Higuchi, T.: 'Power dissipation reductions with genetic algorithms'. Proc. of the NASA/DOD Conf. on Evolvable Hardware, 2003, p. 111
- 12 Stoica, A., Zebulum, R., Keymeulen, D., Tawel, R., Daud, T., Thakoor, A.: 'Reconfigurable VLSI architectures for evolvable hardware: from experimental field programmable transistor arrays to evolution-oriented chips', *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, 2001, **9**, (1), pp. 227–232
- 13 Langeheine, J., Trefzer, M., Brüderle, D., Meier, K., Schemmel, J.: 'On the evolution of analog electronic circuits using building blocks on a CMOS FPTA'. Proc. of the Genetic and Evolutionary Computation Conf. (GECCO), June 2004, pp. 1316–1327
- 14 Asenov, A.: 'Statistical nano CMOS variability and its impact on SRAM'. Extreme Statistics in Nanoscale Memory Design, 2010, pp. 17–50
- 15 Ghibaudo, G., Ioannidis, E., Dimitriadis, C., Manceau, J.-P., Haendler, S.: 'Impact of dynamic variability on the operation of CMOS inverter', *Electron. Lett.*, 2013, **49**, (19), pp. 1214–1216
- 16 Walker, J.A., Trefzer, M.A., Bale, S.J., Tyrrell, A.M.: 'PANDA: a reconfigurable architecture that adapts to physical substrate variations', *IEEE Trans. Comput.*, 2013, **62**, (8), pp. 1584–1596
- 17 Trefzer, M.A., Walker, J.A., Tyrrell, A.M.: 'A programmable analog and digital array for bio-inspired electronic design optimization at nano-scale silicon technology nodes'. IEEE Asilomar Conf. on Signals, Systems, and Computers, Asilomar, CA, November 2011
- 18 Langeheine, J., Trefzer, M., Schemmel, J., Meier, K.: 'Intrinsic evolution of digital-to-analog converters using a CMOS FPTA chip'. Proc. of the NASA/DoD Conf. on Evolvable Hardware, June 2004, pp. 18–25
- 19 Paluchowski, S., Cheng, B., Roy, S., Asenov, A., Cumming, D.: 'Investigation into effects of device variability on CMOS layout motifs', *Electron. Lett.*, 2008, **44**, (10), p. 626
- 20 Millar, C., Reid, D., Roy, G., Roy, S., Asenov, A.: 'Accurate statistical description of random dopant induced threshold voltage variability', *IEEE Electron Device Lett.*, 2008, **29**, (8)
- 21 Asenov, A.: 'Random dopant induced threshold voltage lowering and fluctuations in sub 50 nm MOSFETs: a statistical 3D 'atomistic' simulation study', *Nanotechnology*, 1999, **10**, pp. 153–158
- 22 Moroz, V.: 'Design for manufacturability: OPC and stress variations'. Proc. of the First Int. Conf. on CMOS Variability, 2007
- 23 Eccleston, W.: 'The effect of polysilicon grain boundaries on MOS based devices', *Microelectron. Eng.*, 1999, **48**, pp. 105–108
- 24 Matsunawa, T., Nosato, H., Sakanashi, H., *et al.*: 'Adaptive optical proximity correction using an optimization method'. Proc. of the Seventh IEEE Int. Conf. on Computer and Information Technology (CIT), 2007, pp. 853–860
- 25 Kheterpal, V., Rovner, V., Hersan, T.G., *et al.*: 'Design methodology for IC manufacturability based on regular logic-bricks'. Proc. of the 42nd Annual Design Automation Conf., 2005, pp. 353–358
- 26 Hilder, J.A., Walker, J.A., Tyrrell, A.M.: 'Optimising variability tolerant standard cell libraries'. 2009 IEEE Congress on Evolutionary Computation, May 2009, pp. 2273–2280
- 27 Langeheine, J., Trefzer, M.A., Schemmel, J., Meier, K.: 'Intrinsic evolution of analog electronic circuits using a CMOS FPTA chip'. Fifth Conf. on Evolutionary Methods for Design, Optimization and Control with Applications to Industrial and Societal Problems (EUROGEN), 2003
- 28 Langeheine, J.: 'Intrinsic hardware evolution on the transistor level'. PhD dissertation, Rupertus Carola University of Heidelberg, Heidelberg, July 2005
- 29 Deb, K., Agrawal, S., Pratap, A., Meyarivan, T.: 'A fast elitist non-dominated sorting genetic algorithm for multi-objective optimisation: NSGA-II'. Proc. of the Conf. on Parallel Problem Solving from Nature, 2000, pp. 849–858
- 30 Trefzer, M.A.: 'Evolution of transistor circuits'. PhD dissertation, Rupertus Carola University of Heidelberg, Heidelberg, December 2006
- 31 Weste, N., Harris, D.: 'CMOS VLSI design: a circuits and systems perspective' (Addison-Wesley, 2011, 4th edn.)