Microelectronics Reliability 54 (2014) 464–474

Contents lists available at ScienceDirect

Microelectronics Reliability

journal homepage: www.elsevier.com/locate/microrel

Monte Carlo Static Timing Analysis with statistical sampling

Michael Merrett, Mark Zwolinski *

Electronics and Computer Science, University of Southampton, Southampton SO17 IBJ, UK

ARTICLE INFO

Article history:
Received 2 July 2013
Received in revised form 12 October 2013
Accepted 14 October 2013
Available online 7 November 2013

ABSTRACT

With aggressive scaling of CMOS technologies, MOSFET devices are subject to increasing amounts of independent local statistical variability. The causes of these statistical variations and their effects on device performance have been extensively studied, but their impact on circuit performance is still difficult to predict. This paper proposes a method for modeling the impact of random intra-die statistical variations on digital circuit timing. The method allows the variation modeled by large-scale statistical transistor simulations to be propagated up the design flow to the circuit level, by making use of commercial STA and standard cell characterization tools. By using statistical sampling techniques, we achieve close to the accuracy of full SPICE simulation, but with a computational effort similar to that of Statistical Static Timing Analysis, while removing some of the inaccurate assumptions of Statistical Static Timing Analysis.

1. Introduction

The International Technology Roadmap for Semiconductors (ITRS) has repeatedly highlighted that improvements in design productivity are not consistent with the scaling of manufacturing technologies, specifically in areas such as designing for performance and power variability [1]. The accuracy of timing validation and power estimation methods are being challenged by the scaling of process dimensions down to the nanometer level, where the impact of statistical process variations has become increasingly significant. Traditionally, process variations have been modeled at the circuit level by performing static timing analysis (STA), which makes use of calibrated lookup tables for standard cells at multiple technology-dependent corners.

The shortfalls of corner-based analysis have been known for many years, [2]. STA can be both overly pessimistic and optimistic, and hence Statistical Static Timing Analysis (SSTA) was proposed. Since then much effort has been placed into the development of practical and accurate SSTA tools but modern SSTA algorithms are still unable to address issues that are overcome by commonly-used STA methods, such as interconnect analysis, latch based designs and clock-skew analysis [3]. There remain large obstacles to the widespread use of SSTA in the industry, including a limited amount of statistical foundry data available in a standardized format.

The effects of localized statistical process variations such as random discrete dopants (RDD) [4] and line edge roughness (LER) [5,6] have been investigated and are understood at the device level, but these effects must be modeled at the circuit level in order to improve design for manufacturability methods. The challenge is to provide circuit designers with a transparent method for modeling the impact of MOSFET variability to allow trade off design decisions between power consumption, performance and manufacturing yield. We are not concerned in this work with systematic variations, due to, for example, the position of a device on a wafer.

This paper extends the use of Monte Carlo Static Timing Analysis (MCSTA), [7], a transparent method of modeling the impact of statistical transistor variations. MCSTA makes use of existing industry standard timing analysis and cell characterization tools, providing a low cost addition to existing design flows, and takes advantage of the maturity of established STA methods. MCSTA can be performed without assumptions of underlying transistor model parameter distributions, providing greater accuracy than SSTA, while executing considerably faster than Monte Carlo SPICE simulations. Nevertheless, MCSTA is still considerably slower than SSTA. This paper describes how statistical sampling may be used to limit the Monte Carlo simulations to the (interesting) tails of the distributions.

2. Simulation methodology

Previous research into the statistical variation of individual transistors has found that Random Discrete Dopants (RDD), Poly Silicon Granularity [8,9] and Line Edge Roughness (LER) [10] are the three main sources of fluctuations in threshold voltage (V_T). Simulations of three dimensional atomistic transistor models have provided distributions of I–V curves for transistors which traditionally would have been represented by a single continuous charge model. These I–V curves have been converted to a library of BSIM models [11], allowing the range of effects of statistical variations on a transistor to be modeled using SPICE.

* Corresponding author. Tel.: +44 2380593528.
E-mail address: mz@ecs.soton.ac.uk (M. Zwolinski).

0026-2714/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.microrel.2013.10.016

From the point of view of a circuit designer, the effect of variability on individual transistors is not necessarily of interest. The aim of this work is to predict the variability of cells and hence of a particular implementation of the overall circuit. In the case of digital circuits, the main characteristics of interest are the delays through combinational logic paths (which determine the maximum clock speed) and the power dissipated by the circuit.

Thus, the effects of these independent local statistical variations on the performance of a number of test circuits were measured and compared using three methods: Monte Carlo transistor-level SPICE simulations, SSTA, and our proposed MCSTA. The large scale statistical SPICE simulations provide a reference.

2.1. Monte Carlo SPICE simulations

We used a statistical circuit generation tool [12] to replace each MOSFET model instance within a SPICE netlist with BSIM models selected randomly, with a uniform distribution, from a process specific statistical library. Each individual transistor within a circuit was therefore modeled by a separate atomistic model. The tool was used to generate 10,000 randomized transistor level netlists of each circuit under test, and SPICE simulations were then performed on each circuit instance. The input vectors to the circuits included the stimulation of critical paths that were reported during STA. This allowed for a direct comparison of delays through fixed paths, as well as comparisons of the longest path delay through the circuit. The distributions of static and dynamic power consumption were also recorded, as reported elsewhere [13].

2.2. Statistical Static Timing Analysis

Traditional STA allows each delay element within a circuit to be replaced with a single delay value. This delay value is interpolated from a look-up table within a standard cell library (SCL) in which the delay elements (cells) have been characterized for a fixed set of process corners. The SCL is generated by passing SPICE-level netlists of standard cells through a cell characterization tool, Fig. 1(a). Timing analysis tools model the delay and output transition time for the element to be obtained for a given input transition time and output load capacitance, Fig. 1(b). Using this method every instance of a cell within a circuit references an identical model in the SCL; the three instances of NAND gates in the example of Fig. 1(b) all refer to the same NAND model. Traditional STA is therefore unable to represent statistical process variations.

A commercial statistical cell characterization tool was used to characterize a selection of standard cells. The same transistor model parameters and distributions used within the Monte Carlo SPICE simulations were used during this process. A statistical standard cell library (SSCL) was created using a 'Transistor Mismatch' model, which allows the performance of the cell to be established for multiple σ levels of each transistor parameter. These results were combined using a method specific to the commercial tool, allowing the variation to be modeled by a single synthetic variation parameter. This process is based on the observation that if each variable has an independent Gaussian distribution then the impact on a given timing parameter (delay, slew or constraints) resembles a Gaussian distribution.

The SSCL was then used within a commercially available SSTA tool, providing a distribution of delays for paths within the test circuits in a single run. Statistical power analysis was not possible using this method.

2.3. Monte Carlo Static Timing Analysis

The Monte Carlo Static Timing Analysis method was introduced in a previous paper [7]. It requires no assumptions about the underlying statistics of process parameters, instead using Monte Carlo SPICE simulations of standard cells. Multiple SPICE netlists were generated for selected standard cells, where each transistor referred to different BSIM models selected randomly from the process specific statistical library using RandomSpice. The multiple randomized cell netlists were passed into the same commercial cell characterization tool as with SSTA, generating a Variation Cell Library (VCL) where every standard cell has multiple instances, Fig. 2(a).

Each VCL contained 500 randomized instances of each of the standards cells. Each characterization involves a SPICE simulation and the extraction of cell parameters. The complexity of each simulation depends on the nature of the cell – the number of inputs and outputs and whether there is an internal state. Although we did not characterize every cell in the library for this exercise, only those that were instanced in the circuits, a typical cell library might contain 1000 different cells. This implies up to half a million unique cell characterizations and hence hundreds of hours of CPU time. Fortunately, this exercise only needs to be done once per cell library (and is trivially parallel).

The statistical process variations within a VCL were modeled at the circuit level by making simple modifications to the gate level netlist of a design. Each reference to a standard cell within the original netlist was altered to represent a variation cell reference by the addition of the suffix _x, where *x* refers to the chosen variation model. Each individual cell was therefore modeled by a separate set of atomistic transistor models, rather than the same model as in traditional STA.

Multiple copies of the original netlist were created, where a randomized suffix was added to each SCL reference. This randomization process does not alter the structure or behavior of the circuit, only the statistical process variations of the cells being modeled. Each of these randomized netlists can then be passed through the STA tool, producing individual timing reports for each netlist. These results can then be combined and analyzed to produce a distribution of the timing of a design, measuring any slack within the critical paths and providing the probability of any paths through the design failing to meet timing requirements. The process of generating and performing STA on randomized gate level netlists is referred to as Monte Carlo Static Timing Analysis and is illustrated in Fig. 2(b).

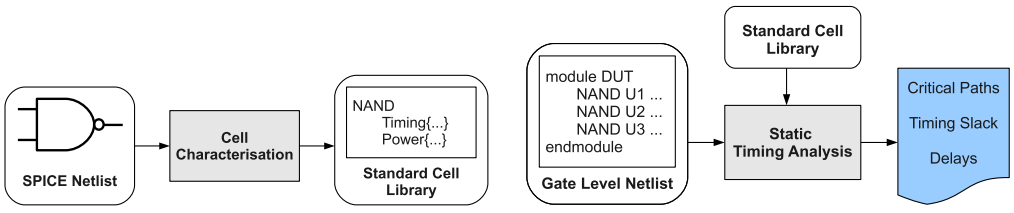

(a) Generation of a Standard Cell Library (SCL) (b) Static Timing Analysis (STA)

Fig. 1. Generating a standard cell library and performing static timing analysis.