A Heuristic Method for Statistical Digital Circuit Sizing

Stephen Boyd[a], Seung-Jean Kim[a], Dinesh Patil[b], Mark Horowitz[b]

[a]Information Systems Laboratory, Department of Electrical Engineering,
Stanford University, CA, USA 94305;
[b]Computer Systems Laboratory, Department of Electrical Engineering,
Stanford University, CA, USA 94305

ABSTRACT

In this paper we give a brief overview of a heuristic method for approximately solving a statistical digital circuit sizing problem, by reducing it to a related deterministic sizing problem that includes extra margins in each of the gate delays to account for the variation. Since the method is based on solving a deterministic sizing problem, it readily handles large-scale problems. Numerical experiments show that the resulting designs are often substantially better than one in which the variation in delay is ignored, and often quite close to the global optimum. Moreover, the designs seem to be good despite the simplicity of the statistical model (which ignores gate distribution shape, correlations, and so on). We illustrate the method on a 32-bit Ladner-Fischer adder, with a simple resistor-capacitor (RC) delay model, and a Pelgrom model of delay variation.

Keywords: Design for manufacturing, design for yield, statistical circuit sizing

1. INTRODUCTION

For current integrated circuit (IC) technologies, statistical uncertainty and process variation can be indirectly handled by incorporating simple margins in the timing and other critical constraints, or by a post-design step that does centering or yield improvement.[1–3] As device dimensions shrink, however, growing (relative) statistical uncertainty and process variation will require an approach where design and yield optimization are combined.[4–7] The increasing importance of process variation explains the growing interest in process variation modeling, and statistical timing analysis and design of of digital circuits.[8–12]

In statistical design we take into account statistical variation in the device, process, and environment parameters *for each gate*. In other words, we consider local variations in the model parameters for gate delay, energy, and leakage current, with each gate having its own set of parameter values, drawn from some distribution. (This is in contrast to the framework for robust design over corners, where the gate model parameters variations are global, i.e., the same for all gates.) In the simplest case, the parameter values for each gate are modeled as independent random variables, but more sophisticated models can include correlation between the parameters associated with different gates. Another approach is to include two unknown terms in the parameters of each gate: one is a systematic one, global for the whole circuit (as in robust design over corners), and the other is the local uncertainty we consider here. This leads to a blend of robust design over corners (to handle global parameter variation) and statistical design (to handle local parameter variation).

While there are many possible choices of objectives in statistical design, we concentrate on a typical one, an upper quantile (e.g., 95%) of the cycle time. Sizing a circuit that meets a timing specification with high probability, despite statistical variation, is called *statistical circuit sizing*. It is closely related to *design for yield* (DFY), *design for manufacture* (DFM), and *design centering*.

There are many exact or direct methods for solving deterministic circuit sizing problems, even for large-scale problems; see Ref. 13 for those methods for digital circuit sizing. There are far fewer methods for solving

Further author information: (Send correspondence to S. Boyd.)
S. Boyd: E-mail: boyd@stanford.edu, Telephone: 1 650 723 0002
S. Kim: E-mail: sjkim@stanford.edu, Telephone: 1 650 724 3640
D. Patil: E-mail: ddpatil@stanford.edu, Telephone: 1 650 725 8811
M. Horowitz: E-mail: horowitz@stanford.edu, Telephone: 1 650 725 3707

statistical circuit sizing problems, and these are computationally far more demanding (or intractable), even for small problems. Moreover, the designs obtained depend on details of the gate delay distributions, which are often not well known in practical applications.

In this paper we describe a simple heuristic method[14, 15] for approximately solving a statistical digital circuit sizing problem, by reducing it to a related deterministic sizing problem that includes extra margins in each of the gate delays to account for the variation. Since the method is based on solving a deterministic sizing problem, it readily handles large-scale problems. Numerical experiments show that the resulting designs are often substantially better than one in which the variation in delay is ignored, and often quite close to the global optimum.[14] Moreover, the designs seem to be good despite the simplicity of the statistical model (which ignores gate distribution shape, correlations, and so on).[14]

2. STATISTICAL DESIGN OF DIGITAL CIRCUITS

We first consider a basic deterministic gate scaling problem for a circuit consisting of n gates, each with a variable scale factor that sets its drive strength. The goal is to choose the scale factors to give minimum delay, subject to limits on the total area and power. This can be expressed as the optimization problem

$$\begin{array}{ll} \text{minimize} & D \\ \text{subject to} & P \leq P^{\max}, \qquad A \leq A^{\max}, \\ & 1 \leq x_i, \quad i = 1, \ldots, n, \end{array} \tag{1}$$

where D is the maximum circuit delay (over all paths), A is the total area, and P is the power. The design variables are the gate scale factors $x_1, \ldots, x_n$, which scale the widths of the devices used to form the gates. Here $P^{\max}$ and $A^{\max}$ are given limits on the total power and area.

Now we consider the effects of parameter variation. Statistical variation in the gate parameter values induces statistical variation in the performance objectives D and P, which are described by probability distributions, that depend on the choice of scaling factors, i.e., x. We can modify the basic deterministic problem (1) by requiring that the power constraint should hold with some minimum probability (or reliability) such as 95%, and taking as objective the 95% quantile (say) of delay:

$$\begin{array}{ll} \text{minimize} & \mathbf{Q}^{.95}(\mathbf{D}) \\ \text{subject to} & \mathbf{Q}^{.95}(\mathbf{P}) \leq P^{\max}, \qquad A \leq A^{\max}, \\ & 1 \leq x_i, \quad i = 1, \ldots, n. \end{array} \tag{2}$$

Here $\mathbf{D}$ and $\mathbf{P}$ are random variables, and $\mathbf{Q}^{.95}(\mathbf{X})$ denotes the 95% quantile of the random variable $\mathbf{X}$. In the basic statistical design problem formulation (2), we insist that 95% of the circuits meet the power constraint, and we judge a design by the 95% quantile of its delay. The formulation (2) is a *stochastic optimization problem*. The design variables are still just the gate scale factors $x_1, \ldots, x_n$, but the objective $\mathbf{Q}^{.95}(\mathbf{D})$ and constraint function $\mathbf{Q}^{.95}(\mathbf{P})$ are now very complicated functions of x. The stochastic optimization problem (2) is very difficult or impossible to solve exactly, so most approaches are heuristic, and attempt to solve the problem approximately.

The statistical *analysis* problem, however, is more tractable than the statistical design problem (2). If we are given a proposed design x, we can find the associated probability distributions of the gate delays and powers (including any correlations or dependencies among them), so we can estimate $\mathbf{Q}^{.95}(\mathbf{D})$ and $\mathbf{Q}^{.95}(\mathbf{P})$ via Monte Carlo analysis. Monte Carlo analysis reveals not only these quantiles of the random variables, but gives estimates of their probability density functions (PDFs) as well.

We will illustrate our method using simple models for power, delay, and their statistical variation, as a function of gate scale factor and load capacitance. With these models, the final optimization problem has a special form: it is a (generalized) geometric program (GP), a special type of optimization problem that can be solved globally and efficiently.[16] We refer the reader to Ref. 17 for an introduction to geometric programming, some of the basic tricks used to formulate problems in GP form, a number of examples, and an extensive list of references. An implementation of a primal-dual interior-point method for GP is available in GGPLAB, a Matlab-based toolbox for GP[18] and in the MOSEK software package.[19] GP-based circuit sizing is by no means new; it has been used for digital circuits since the 1980s,[13, 20] as well as analog circuits[21, 22] and RF circuits.[23–26]

2.1. Statistical power constraint

The statistical power constraint $\mathbf{Q}^{.95}(\mathbf{P}) \leq P^{\max}$ is not very difficult to handle, at least approximately, since the power is the *sum* of the individual gate powers $\mathbf{P}_i$. Assuming the parameter variations are independent (enough) and the circuit contains a large number of gates (which holds in problems of interest) the power $\mathbf{P}$ has a standard deviation that is small relative to its mean $\mathbf{EP}$. Thus, we can use $\mathbf{EP} \leq P^{\max}$ as a reasonable approximation of $\mathbf{Q}^{.95}(\mathbf{P}) \leq P^{\max}$. Since $\mathbf{P} = \sum_{i=1}^{n} \mathbf{P}_i$, our approximation of the statistical power constraint $\mathbf{Q}^{.95}(\mathbf{P}) \leq P^{\max}$ reduces to

$$\sum_{i=1}^{n} \mathbf{EP}_i \leq P^{\max}.$$

Thus, we can take into account the effect of statistical variation on total circuit power by replacing the nominal gate power model with a mean gate power model (where the mean is over the parameter variation). The mean gate power can be well approximated in a form compatible with geometric programming.[13]

This simple treatment of the power constraint relies on the fact that in a sum of a large number of random variables, which are uncorrelated (or just not too correlated), the random variations tend to cancel each other out, leading to (relatively) lower variation in the sum.

2.2. Statistical delay analysis

The effect of statistical variation in the gate delays is far more difficult to analyze than the effect of variation in the gate powers, and is what makes the statistical gate sizing problem (2) challenging. Assuming that the variations in gate delays are independent, the delay of the paths, which are sums of gate delays, have reduced variance, by a factor on the order of the number of gates on the path (with is typically not large, on the order of 10). Thus, the paths still exhibit substantial statistical variation in delay.

The overall delay of a circuit is the *maximum* of the path delays. The *maximum* of a set of random variables behaves very differently from a *sum*. The maximum of a set of random variables can have a distribution with a strong right skew, and a variance substantially *larger* than the variance of the individual variables. In a sum of random variables, the individual random variations tend to cancel each other out; for the sum to be large, it is required that many of the individual variable should be large, which is very unlikely. But in a maximum, no such cancellation occurs; all that is required for the maximum to be large is that just *one* of the individual variables should be large. With a large number of random variables, with enough independence, it is quite likely that one of them is large.

We use the *alpha-power law model*[27] and *Pelgrom's model*,[28] to carry out an analysis of gate delay variation, induced by variation in the threshold voltage V_{th}, a critical electrical parameter of the devices in a gate. Gate delay varies with threshold voltage according to the alpha-power law model

$$D \propto \frac{V_{\text{dd}}}{(V_{\text{dd}} - V_{\text{th}})^{\alpha}},$$

where α is a parameter typically between 1.3 and 2, and V_{dd} is the supply voltage. One simple and commonly used model for threshold voltage variation is Pelgrom's model, which predicts that the threshold voltage of the devices in a gate has variance

$$\sigma_{V_{\text{th}}}^2 = \overline{\sigma}_{V_{\text{th}}}^2 x^{-1},$$

where $\overline{\sigma}_{V_{\text{th}}}^2$ is the variance for a unit scaled gate. This model predicts that larger gates have smaller variation in threshold voltage than smaller ones, due to spatial averaging. Assuming the statistical variation in V_{th} is not too large, we have

$$\sigma \approx \left| \frac{\partial D}{\partial V_{\text{th}}} \right| \sigma_{V_{\text{th}}} = \frac{\alpha \overline{\sigma}_{V_{\text{th}}}}{(V_{\text{dd}} - V_{\text{th}})} x^{-1/2} D, \tag{3}$$

where $\overline{\sigma}_{V_{\text{th}}}^2$ is the variance for a unit scaled gate and σ is the standard deviation of the gate delay D. Thus the relative delay variation decreases with increasing device area. Moreover, the gate delay standard deviation is a generalized posynomial of the scale factors, whenever D is, and so is compatible with geometric programming.[13]

The distribution of the overall circuit delay is found from the gate delay distributions, propagated through the function that maps gate delays into overall circuit delay, which is a maximum of a large number of sums. There is no simple analysis or description of this distribution, which, as mentioned above, can have a large right skew. For a fixed design (i.e., choice of x), however, we can compute a good estimate of this distribution by Monte Carlo analysis.

3. A HEURISTIC FOR STATISTICAL DESIGN

In this section we describe a simple heuristic for the statistical design problem (2). We start with a (deterministic) model for each gate, which gives the *mean* delay D_i for gate i. To this mean gate delay we add a zero-mean random variable with variance σ_i^2, which represents the statistical fluctuation or uncertainty in the gate delay. We assume that σ_i is, like D_i, a function of the scale factor, load capacitance, and input signal transition times. The ratio σ_i/D_i is a measure of the relative or percentage variation in the gate delay.

Now we can describe the heuristic method. We form a 'surrogate delay' model

$$\tilde{D}_i = D_i + \kappa_i \sigma_i,$$

where κ_i are constants (often all the same), typically between 1 and 3, that are used to trade off mean and variance of the overall delay. The extra term $\kappa_i\sigma_i$ adds an extra margin in the gate delay model, that scales with increasing gate delay variance. Now we optimize the circuit as usual, using the surrogate delays $\tilde{D}_i$ in place of the (mean) delays D_i. Note that when $\kappa_i = 0$, this method is the same as the standard (non-statistical) design method. We can analyze the resulting design using Monte Carlo analysis, coupled with static timing analysis. We can then adjust the constants κ_i for optimum performance (for example, estimated yield). For example, designs can be carried out with all κ_i constant, and equal to the values $\kappa_i = 2$, $\kappa_i = 2.5$, $\kappa_i = 3$; the best of the resulting three designs is taken as the final design.

This simple heuristic method is similar in spirit to the general method of *regularization*,[29] in which an extra penalty term is added to a problem to approximately account for some variation in the problem data. It is also related to *robust optimization*,[30,31] a more sophisticated and recent approach to handling uncertainty and variation in optimization problem data, in which an explicit model of data uncertainty is used, and the objective is taken to be the average or worst-case value of the objective, over the parameter variation set.

This simple heuristic method looks simple, but is more sophisticated than it appears, since the extra margins are added on a gate-by-gate basis, and not on a path-by-path basis. The quality of the resulting suboptimal designs can be assessed using widely applicable lower bounds on achievable performance in optimal statistical design.[14] In some cases, the method yields a design that is *provably* close to the global optimum of the (difficult) stochastic optimization problem.[14]

4. STATISTICAL DESIGN EXAMPLE

We illustrate the heuristic method for statistical design on a 32-bit Ladner-Fisher adder,[32] consisting of 459 gates, including 64 input gates and 32 output gates, and 1714 arcs. The circuit has a total of 3214 paths from input gates to output gates. The maximum path length is 12.

We use a simplified static timing model, with a single delay for each gate (ignoring differing rise and fall times, different delays for different gate transitions, and the effects of signal slopes). The scale factor $x_i \geq 1$ scales the widths of the devices used to form the gate and therefore affects its drive strength, input capacitance, and area. (The same method can be applied to a full custom design, in which each device is sized individually.[15]) The scale factor $x_i = 1$ corresponds to a minimum sized gate, and a scale factor $x_i = 16$ (say) corresponds to a version of the gate in which all devices have width 16 times the widths of the devices in the minimum sized gate.

Gate i has three parameters: an input capacitance $\bar{C}_i^{\mathrm{in}}$, an intrinsic or internal capacitance $\bar{C}_i^{\mathrm{int}}$, and driving resistance $\bar{R}_i^{\mathrm{int}}$. The input and intrinsic capacitances are modeled as linear functions of the scale factor,

$$C_i^{\mathrm{in}} = \bar{C}_i^{\mathrm{in}} x_i, \qquad C_i^{\mathrm{int}} = \bar{C}_i^{\mathrm{int}} x_i,$$

Table 1. The 5 gate types used in the Ladner-Fisher adder. The first column gives the gate name; the second column gives the logic function the gate implements, and the remaining 4 columns give the model parameters.

gate type	function	$\bar{C}^{\text{in}}$	$\bar{C}^{\text{int}}$	$\bar{R}$	$\bar{A}$
INV	$\overline{A}$	3	3	0.48	3
NAND2	$\overline{AB}$	4	6	0.48	8
NOR2	$\overline{A+B}$	5	6	0.48	10
AOI21	$\overline{AB+C}$	6	7	0.48	17
OAI21	$\overline{(A+B)C}$	6	7	0.48	16

where $\bar{C}_i^{\text{in}}$ and $\bar{C}_i^{\text{int}}$ are the input capacitance and intrinsic capacitance of gate i with unit scaling. The driving resistance R_i is inversely proportional to the scale factor:

$$R_i = \bar{R}_i/x_i,$$

where $\bar{R}_i$ is the driving resistance of gate i with unit scale factor. Let C_i^{L} be the load capacitance that gate i drives. Then, for an output gate,

$$C_i^{\text{L}} = \sum_{j \in \mathbf{FO}(i)} C_j^{\text{in}},$$

where $\mathbf{FO}(i)$ is the set of fan-out gates of i. Using the simple resistor-capacitor model of a gate and its load, we approximate the gate delay as

$$D_i(x) = 0.69 R_i (C_i^{\text{int}} + C_i^{\text{L}}), \tag{4}$$

which is the time required for the output voltage of an RC circuit to reach the midpoint between the logic voltage levels.

Another parameter of gate i is the area. We approximate the (physical) area of gate i as proportional to the scale factor x_i, so the total area of the circuit has the form

$$A = \sum_{i=1}^{n} x_i \bar{A}_i,$$

where $\bar{A}_i$ is the area of gate i with unit scaling. The gate area is the total width of the devices in the gate (since the gate lengths are always chosen to be the smallest value allowed in the technology.)

The Ladner-Fisher adder contains 5 types of gates, with associated functions and model parameters listed in Table 1. The capacitance unit is the capacitance of the NMOS device in a unit scaled inverter, and the area unit is the width of the NMOS device in a unit scaled inverter. The drive strength value $\bar{R} = 0.48$ is chosen so that the delay of a unit size inverter with no load is $0.69 \cdot 0.48 \cdot 3 = 1$. In other words, the time unit is normalized to the delay of a unit scale inverter, with no load, denoted by τ. The model parameters come from the *logical effort model*.[33]

The expression (4) gives the mean delay of a gate. We take the standard deviation of the gate delay to be

$$\sigma_i(x) = \gamma x_i^{-1/2} D_i(x),$$

which is motivated by Pelgrom's model as discussed above. The parameter γ gives the relative variation for a minimum sized gate (i.e., $x_i = 1$). We used the model parameter $\gamma = 0.15$, which means that for a minimum sized gate, the delay standard deviation is 15% of its mean, and that this ratio decreases with increasing gate size.

We assume that the delay distributions are Gaussian and independent. The actual gate delay distribution cannot be exactly Gaussian, of course, since a Gaussian variable has a positive probability of being negative. But

Table 2. Comparison of nominal and statistical designs.

	nominal delay	**ED**	$\sigma_{\mathbf{D}}$	$\mathbf{Q}^{.95}(\mathbf{D})$
nominal design	45.1	48.9	0.88	50.4
statistical design	45.7	47.5	0.47	48.2

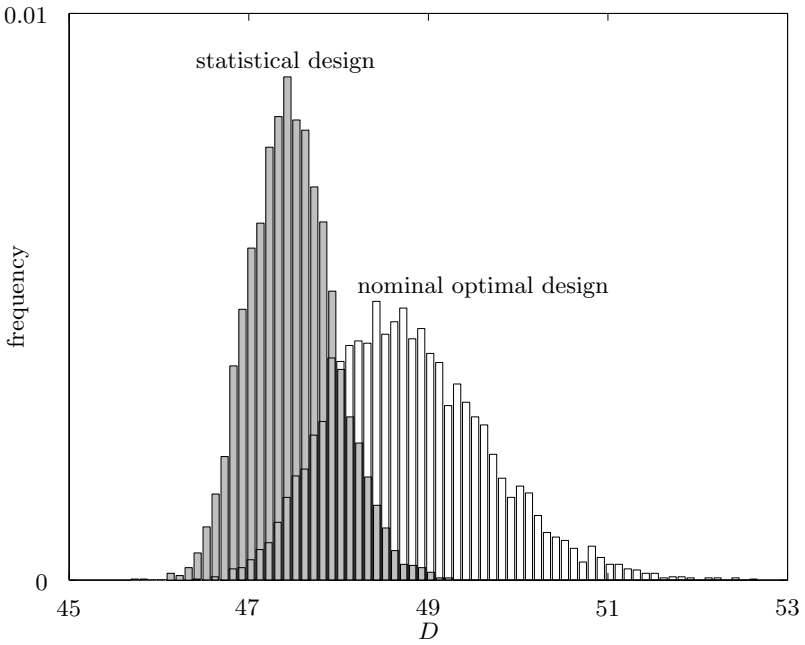

Figure 1. Distributions of circuit delay for nominal and statistical designs.

since each gate delay satisfies $\sigma/\mu \leq 0.15$, the probability of a negative gate delay is vanishingly small. In fact, modeling the statistics of gate delay is an area of active research, with many open issues.[34–37] In any case, we have found that the designs produced by our method are not sensitive to details of the gate distribution shape.[14]

For the optimization problems, we impose a constraint on the area, as well as lower bounds on the scale factors:

$$A \leq 15000, \qquad 1 \leq x_i, \quad i = 1, \ldots, n. \tag{5}$$

The load capacitance of each primary output is taken as $C_i^{\mathrm{L}} = 6$. Monte Carlo analysis reveals that the heuristic statistical design method gives good results with $\kappa_i = 2$. The mean delay and variance model described above are both posynomial functions of x, as are the constraint functions, so the surrogate circuit sizing problem can be formulated as a (generalized) geometric program.

The performance of the statistical design, found by Monte Carlo analysis with 5000 samples, is compared to the nominal design (i.e., the one obtained by ignoring the statistical variation) in Table 2. The robust design has a nominal circuit delay (i.e., a circuit delay ignoring statistical variation) of 45.7, only 1.3% more than the nominal design. When we take statistical variation into account, however, the two designs differ. The 95% circuit delay for the nominal optimal design is 50.4, which is 11.8% more than the nominal circuit delay. For the robust design, the 95% circuit delay is 48.2, which is only 5.6% more than the nominal optimal circuit delay. Thus, the statistical design has reduced the effect of statistical delay variation by a factor of around 2, compared to the nominal optimal design. The standard deviation of the statistical design is also reduced by a factor around 2, compared to the nominal design. The distributions of the delay for the nominal and statistical designs (estimated by Monte Carlo simulation) are shown in Fig. 1.

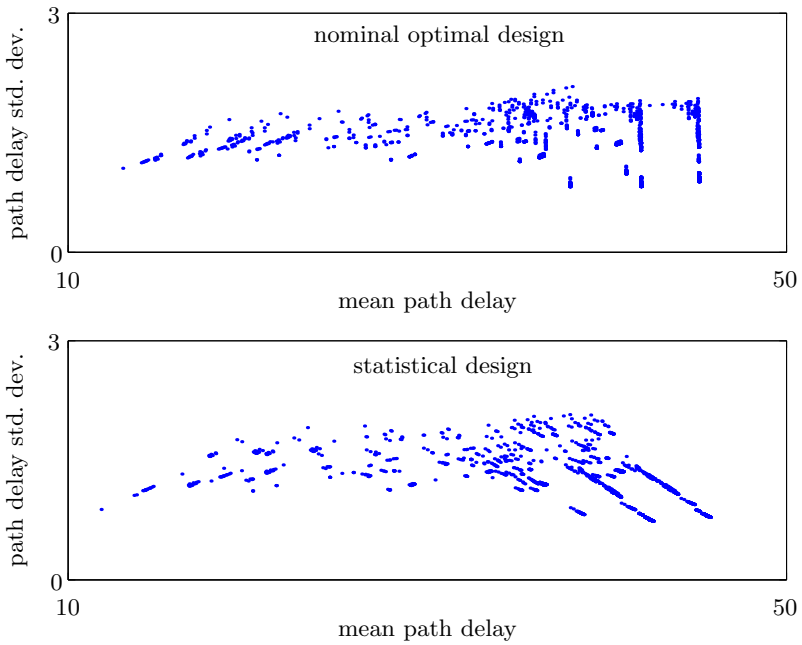

Figure 2. Scatter plots of path delay mean versus path delay standard deviation, for the nominal design (top) and statistical design (bottom).

Some insight into why the statistical design performs better than the nominal design can be found in Fig. 2, which shows scatter plots of mean delay versus standard deviation for all 3214 paths in the Ladner-Fischer adder, for the nominal and statistical designs. The problematic paths are the ones at the upper right, which represent paths that are near critical, and in addition have large standard deviation. The plots show that in the nominal design, a number of paths with large expected delays have large standard deviation; in the statistical design, however, the variances of the paths with large expected delays are smaller; paths with relatively small expected delays, however, have relatively larger variances.

5. CONCLUSIONS

We have described a heuristic method for statistical sizing of a digital circuit, by relating it to a related deterministic sizing problem which includes extra margins in each of the gate delays to account for the variation, using Monte Carlo analysis to verify the performance of the designs, and choosing the best one. Our computational experience[14, 15] with the method so far suggests that the heuristic method produces designs that

- are often far superior to the nominal optimal design (obtained by ignoring statistical variation),
- are often provably close to the global optimum,
- are not very sensitive to the details of the individual gate delay distributions or correlations among them.

We have extended the method described here to problems with more accurate delay models, with different delay models for rising and falling signals, different input/output pairs for each gate, and effects of signal slope.[15, 38] In the design problem we size individual devices (as opposed to whole gates as in the example considered here), and take into account power as well as area.

Acknowledgments

This material is supported in part by the National Science Foundation under grants #0423905 and (through October 2005) #0140700, by the Air Force Office of Scientific Research under grant #F49620-01-1-0365, by MARCO Focus center for Circuit & System Solutions contract #2003-CT-888, and by MIT DARPA contract #N00014-05-1-0700.

REFERENCES

1. H. Abdel-Malek and J. Bandler, "Yield optimization for arbitrary statistical distributions: Part I-Theory," *IEEE Transactions on Circuits and Systems* **27**, pp. 245–253, Apr. 1980.
2. H. Abdel-Malek and J. Bandler, "Yield optimization for arbitrary statistical distributions: Part II-Implementation," *IEEE Transactions on Circuits and Systems* **27**, pp. 253–262, Apr. 1980.
3. M. Styblinski and L. Opalski, "Algorithms and software tools for IC yield optimization based on fundamental fabrication parameters," *IEEE Transactions on Computer-Aided Design* **5**, pp. 79–89, Jan. 1986.
4. J.-T. Kong, "CAD for nanometer silicon design challenges and success," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **12**(11), pp. 1132–1147, 2004.
5. J. Jess, K. Kalafala, S. Naidu, R. Otten, and C. Visweswariah, "Statistical timing for parametric yield prediction of digital integrated circuits," in *Proc. of 40th Proc. IEEE/ACM Design Automation Conference*, pp. 343–347, (Anaheim, CA), June 2003.
6. C. Visweswariah, "Death, taxes and failing chips," in *Proc. of 40th IEEE/ACM Design Automation Conference*, pp. 343–347, (Anaheim, CA), June 2003.
7. K. White, W. Trybula, and R. Athay, "Design for semiconductor manufacturing-perspective," *IEEE Transactions on Components, Packaging, and Manufacturing Technology-Part C* **20**, pp. 58–72, Jan. 1997.
8. A. Agarwal, V. Zolotov, and D. Blaauw, "Statistical timing analysis using bounds and selective enumeration," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **22**(9), pp. 1243–1260, 2003.
9. S. Bhardwaj, S. Vrudhula, and D. Blaauw, "TAU: Timing analysis under uncertainty," in *International Conference on Computer-Aided Design*, pp. 615–620, (San Jose, CA, USA), Nov. 2003.
10. H.-F. Jyu, S. Malik, S. Devadas, and K. Keutzer, "Statistical timing analysis of combinational logic circuits," *IEEE Transactions on VLSI Systems* **1**(2), pp. 126–137, 1993.
11. A. Brambilla and P. Maffezzoni, "Statistical method for the analysis of interconnects delay in submicrometer layouts," *IEEE Transactions on Compute-Aided Design of Integrated Circuits and Systems* **20**(8), pp. 957–966, 2001.
12. H. Hashimoto and H. Onodera, "A performance optimization method by gate sizing using statistical static timing analysis," in *Proc. of ACM International Symposium on Physical Design*, pp. 111–116, (San Diego, CA), Apr. 2000.
13. S. Boyd, S.-J. Kim, D. Patil, and M. Horowitz, "Digital circuit optimization via geometric programming," *Operations Research* **53**(6), pp. 899–932, 2005.
14. S.-J. Kim, S. Boyd, S. Yoon, D. Patil, and M. Horowitz, "A heuristic for optimizing stochastic activity networks with applications to statistical digital circuit sizing," 2006. To appear in *Optimization and Engineering*. Available from `www.stanford.edu/boyd/~heur_san_opt.html`.
15. D. Patil, Y. Yun, S.-J. Kim, S. Boyd, and M. Horowitz, "A new method for robust design of digital circuits," in *Proccedings of the Sixth International Symposium on Quality Electronic Design (ISQED) 2005*, pp. 676–681, IEEE Computer Society Press, 2004.
16. S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
17. S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," 2006. To appear in *Optimization and Engineering*. Available from `www.stanford.edu/boyd/~gp_tutorial.html`.
18. A. Mutapcic, K. Koh, S.-J. Kim, L. Vandenberghe, and S. Boyd, *GGPLAB: A Simple Matlab Toolbox for Geometric Programming*, 2005. Available from `www.stanford.edu/~boyd/ggplab/`.
19. MOSEK ApS, *The MOSEK Optimization Tools Version 2.5. User's Manual and Reference*, 2002. Available from `www.mosek.com`.

20. J. Fishburn and A. Dunlop, "TILOS: A posynomial programming approach to transistor sizing," in *IEEE International Conference on Computer-Aided Design: ICCAD-85. Digest of Technical Papers*, pp. 326–328, IEEE Computer Society Press, 1985.
21. J. Dawson, S. Boyd, M. Hershenson, and T. Lee, "Optimal allocation of local feedback in multistage amplifiers via geometric programming," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* **48**, pp. 1–11, January 2001.
22. P. Mandal and V. Visvanathan, "CMOS op-amp sizing using a geometric programming formulation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **20**, pp. 22–38, Jan. 2001.
23. M. Hershenson, A. Hajimiri, S. Mohan, S. Boyd, and T. Lee, "Design and optimization of LC oscillators," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 65–69, 1999.
24. S. Mohan, M. Hershenson, S. Boyd, and T. Lee, "Bandwidth extension in CMOS with optimized on-chip inductors," *IEEE Journal of Solid-State Circuits* **35**, pp. 346–355, Mar. 2000.
25. S. Mohan, M. Hershenson, S. Boyd, and T. Lee, "Simple accurate expressions for planar spiral inductances," *IEEE Journal of Solid-State Circuit* **34**, pp. 1419–1424, Oct. 1999.
26. Y. Xu, L. Pileggi, and S. Boyd, "ORACLE: Optimization with recourse of analog circuits including layout extraction," in *Proceedings of the 41st IEEE/ACM Design Automation Conference*, pp. 151–154, (San Diego, CA), June 2004.
27. T. Sakurai and A. Newton, "Alpha-power law MOSFET model and its application to CMOS inverter delay and other formulas," *IEEE Journal of Solid-State Circuits* **25**(2), pp. 584–593, 1990.
28. M. Pelgrom, "Matching properties of MOS transistors," *IEEE Journal of Solid State Circuits* **24**, pp. 1433–1439, Oct. 1989.
29. A. Neumaier, "Solving ill-conditioned and singular linear systems: A tutorial on regularization," *SIAM Review* **40**(3), pp. 636–666, 1998.
30. A. Ben-Tal and A. Nemirovski, "Robust convex optimization," *Mathematics of Operations Research* **23**(4), pp. 769–805, 1998.
31. D. Bertsimas and A. Thiele, "A robust optimization approach to supply chain management." To appear in *Operations Research*, 2006.
32. S. Knowles, "A family of adders," in *Proceedings of 14th IEEE Symposium on Computer Arithmetic*, pp. 30–34, IEEE Computer Society Press, 1999.
33. I. Sutherland, B. Sproul, and D. Harris, *Logical Effort: Designing Fast CMOS Circuits*, Morgan Kaufmann Publishers, San Francisco, CA, 1999.
34. J.-J. Liou, A. Krstić, Y.-M. Jiang, and K.-T. Cheng, "Modeling, testing, and analysis for delay defects and noise effects in deep submicron devices," *IEEE Transactions on Compute-Aided Design of Integrated Circuits and Systems* **22**(6), pp. 756–769, 2003.
35. K. Okada, K. Yamaoka, and H. Onodera, "A statistical gate delay model for intra-chip and inter-chip variabilities," in *Proc. of 40th IEEE/ACM Proc. Design Automation Conference*, pp. 31–36, (Anaheim, CA), June 2003.
36. M. Orshansky, L. Milor, P. Chen, K. Keutzer, and C. Hu, "Impact of spatial intrachip gate length variability on the performance of high-speed digital circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **21**, pp. 544–553, May 2002.
37. M. Orshansky, L. Milor, and C. Hu, "Characterization of spatial intrafield gate CD variability, its impact on circuit performance and spatial mask-level correction," *IEEE Transactions on Semiconductor Manufacturing* **17**, pp. 2–11, Feb. 2004.
38. M. Horowitz, E. Alon, D. Patil, S. Naffziger, R. Kumar, and K. Bernstein, "Scaling, power, and the future of CMOS," in *Proceedings of IEEE International Electron Devices Meeting*, Dec. 2005.