

Logic Synthesis for Manufacturability

Alessandra Nardi and
Alberto L. Sangiovanni-Vincentelli
University of California, Berkeley

Editors' note:

Typically, design optimization during synthesis is for area and/or performance, while optimization for yield occurs at the layout level. To obtain more effective yield improvement, this article proposes elevating the abstraction level for yield optimization by introducing an interesting approach to yield-driven logic synthesis.

—Yervant Zorian, *Virage Logic*; and Dimitris Gizopoulos, *University of Piraeus*

■ **IC DESIGN METHODOLOGIES** typically target nominal designs. However, defects and variations in the IC manufacturing process can make a circuit behave substantially different from the nominal design. Particularly in deep-submicron technologies, manufacturing process variations and defects can turn a successful design into a failure. *Design for manufacturability* denotes all techniques designers use to estimate and control yield and robustness during the design phase, prior to manufacturing. DFM is strategically important to decreasing the manufacturing cost of VLSI ICs.

Yield is the percentage of manufactured products that meet all performance and functionality specifications. *Parametric yield loss* usually refers to the effects on circuit performance caused by process variations. *Functional (or catastrophic) yield loss* refers to physical and structural defects that cause the circuit to fail completely.

In deep-submicron technology, yield improvement is as much a design problem as a manufacturing problem. Presently, designers perform optimization for catastrophic yield at a postsynthesis stage: First, synthesis targets area and timing;^{1,2} then the design is optimized for yield at the layout level.

The traditional synthesis approach assumes that larger area leads to lower yield. However, for complex nanometer designs, yield might depend more on design attributes than on total chip area. In fact, the postsynthesis optimization might involve adding redundancy to contacts and vias, modifying the spacing between

existing interconnections, and replacing cell instances with higher-yielding variants, placed in the same position as the original ones. Postsynthesis yield optimization does not perturb placement, thus avoiding costly repetition of optimization aimed at maximizing speed and minimizing area. Traditional postsynthesis yield optimization can reduce manu-

facturing cost by up to 10%.^{3,4}

Researchers have devoted extensive work to understanding circuit sensitivity to manufacturing defects and reducing catastrophic yield loss.^{5,9} Heineken, Khare, and d'Abreu present an attempt to make yield optimization less of an art.³ Their approach is to generate yield-optimized cells and substitute them in the synthesized design. However, their algorithm is limited to an in-place substitution to preserve the original design's footprint.

Following the general rule that the higher the abstraction level, the greater the optimization horizon, we propose a synthesis-for-manufacturability approach that makes manufacturability part of the cost function that drives synthesis. Shaikh, Khare, and Heineken independently pointed out the necessity of including manufacturability (and testability) in the synthesis cost function,¹⁰ but, to the best of our knowledge, we are the first to develop a methodology for doing so. In particular, we optimize for yield on noncritical paths, thus avoiding introducing additional delays. We trade off yield with area in the secondary component of the synthesis cost function. Our approach is completely user transparent, and it leverages the logic synthesis engine itself (rather than an external optimizer running during physical synthesis) to maximize effective yield.

New approach

Figure 1 shows the current, yield-aware design flow and the new design approach we propose. The stan-

standard flow is an optimization process with two main objectives: maximum speed and minimum area.^{1,2} At the end of this process, the design is optimized for yield enhancement at the layout level. Our new design paradigm introduces manufacturability to replace area in the cost function. This approach is transparent to designers, who need only choose whether to minimize area or minimize manufacturing cost.

A possible further development in the synthesis-for-manufacturability paradigm is to incorporate yield during standard-cell design. In addition to variants for high speed and low power, the cell library might also include variants for high yield. Later, we describe a preliminary analysis of how our new technology-mapping heuristic exploits libraries containing yield-optimized variants to reduce manufacturing cost.

Comparing Figure 1a and 1b shows that our approach focuses on choosing the best gate mapping before physical synthesis, rather than manually or semiautomatically performing in-place optimization limited by design placement. Our methodology is a global optimization, which traditional local optimization can effectively complement once the design is laid out. The designer can make additional improvements during routing.⁷ Other circuit techniques, such as adding redundancy to render the circuit robust to failure, are also possible.⁹

In the synthesis-for-manufacturability approach, performance is still the first target during optimization. In fact, whereas the standard flow tries to optimize performance and area on noncritical paths, our approach tries to optimize performance and manufacturability (a combination of area and layout sensitivity to defects) on noncritical paths.

Catastrophic-yield modeling

Now we review the catastrophic-yield-modeling problem. Our focus is not on rigorous statistical modeling for yield prediction; rather, we provide an intuitive view of the new cost function we propose to use during synthesis, along with the related results.

Various types of defects occur in ICs during the manufacturing process. These defects can cause open or short circuits, for example. However, not all defects necessarily cause circuit failures; functional-yield loss depends, in fact, on wafer defectivity and design attributes.

The metric commonly used to quantify layout sensitivity to defects is critical area. $A_i^c(x)$ is the critical area for defects of type i and diameter x ; it is the area within which the center of defects of type i and diameter x must fall to cause a circuit failure. We denote critical

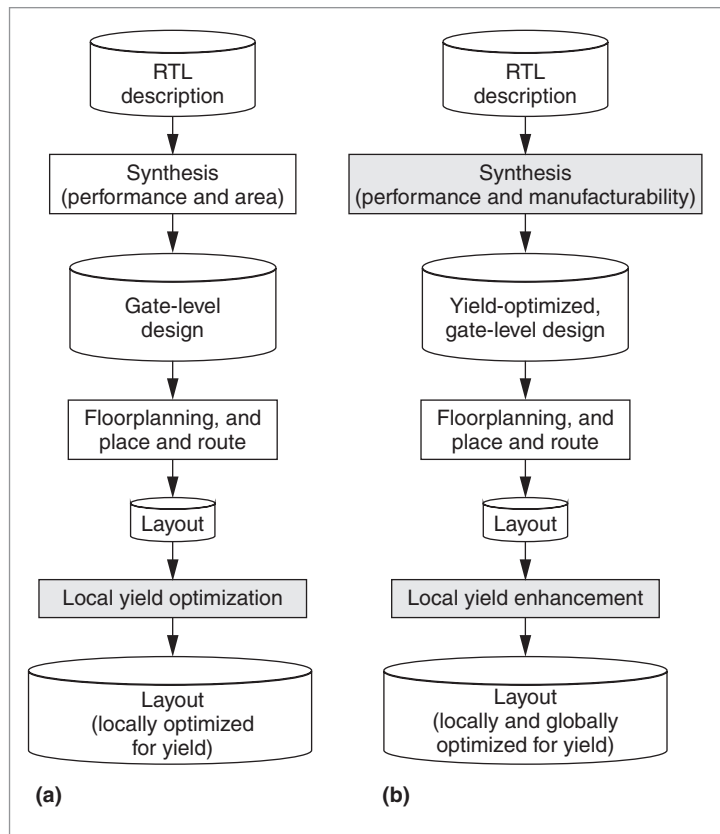


Figure 1. Standard (a) and synthesis-for-manufacturability (b) design flows.

area for defects of type i averaged over all defect diameters x as

$$A_i^c = \int A_i^c(x) f_d(x) dx \quad (1)$$

where $f_d(x)$ is the defect probability density function. With d_i denoting the average number of defects of type i per unit area, the average number of faults on the circuit is

$$\lambda = \sum_i A_i^c \times d_i \quad (2)$$

where the sum is taken over all possible defect types on the circuit.⁹

In general, for given design component b , we can express yield as

$$Y_b = e^{-\lambda_b} \quad (3)$$

where λ_b is the failure rate of component b .

Assuming gates are statistically independent com-

ponents and defects are uniformly distributed (the Poisson model), yield for a circuit with N gates is

$$Y_{\text{circuit}} = \prod_{k=1}^N Y_{G,k} \quad (4)$$

where N is the number of circuit gates and $Y_{G,k}$ is a single gate's yield averaged with respect to different defect types and sizes. From Equation 3, circuit yield becomes

$$Y_{\text{circuit}} = e^{-\sum_{k=1}^N \lambda_{G,k}} = e^{-\lambda_{\text{circuit}}} \quad (5)$$

This model is pessimistic for yield value prediction because defects are not uniformly distributed but rather tend to cluster. There is much literature on deriving a model to accurately predict yield with clustered defect distributions for a given circuit layout.^{9,11} Our research aims instead at analyzing how synthesis for manufacturability can lower manufacturing cost. Therefore, we use the simple Poisson model (Equation 4) and the failure rate description (Equation 5) to focus on yield increase and decrease.

We distinguish between a circuit's effective yield $Y_{\text{eff,circuit}}$, defined as the number of good manufactured circuits per wafer, and a circuit's yield Y_{circuit} , defined as the percentage of good circuits. The following equation expresses their relation:

$$Y_{\text{eff,circuit}} = N_{\text{circuit}} \times Y_{\text{circuit}} \quad (6)$$

where N_{circuit} is the number of circuits in the wafer.

To lower a circuit's manufacturing cost, we must improve its effective yield. Recalling Equations 2 and 3, we observe that yield loss is due to the combination of two effects: the presence of a defect and the component's sensitivity to the defect. Assuming we cannot influence defect distribution, we are left to decrease circuit sensitivity and perform a tradeoff with circuit area. This is the goal of synthesis for manufacturability.

Technology-mapping heuristics for manufacturability

To evaluate the usefulness of the synthesis-for-manufacturability approach depicted in Figure 1b, we now analyze various implementation heuristics. We base this preliminary exploration phase on a modified version of the well-established technology-mapping algorithm available in commercial tools such as the Synopsys Design Compiler.¹ Because performance is our primary objective, we allow some flexibility in area along non-

critical paths. The key is to exploit this flexibility to obtain a higher effective yield by manipulating yield and area instead of simply area.

State-of-the-art tools perform technology mapping with a dynamic programming algorithm that relies on the fact that area is an additive function. Things are not so straightforward for effective yield. In fact, using the Poisson model and assuming a given wafer area, the optimization problem for synthesis for manufacturability becomes

$$\max \frac{Y_{\text{circuit}}}{A_{\text{circuit}}} = \frac{\prod_{k=1}^N Y_{G,k}}{\sum_{k=1}^N A_{G,k}} \quad (7)$$

or, equivalently,

$$\min \frac{A_{\text{circuit}}}{Y_{\text{circuit}}} = \frac{\sum_{k=1}^N A_{G,k}}{\prod_{k=1}^N Y_{G,k}} \quad (8)$$

where A_{circuit} and Y_{circuit} are area and yield of the circuit, and $A_{G,k}$ and $Y_{G,k}$ are area and yield of the single gates composing the circuit. The cost function for the optimization problem described in Equation 8 is not additive, so we must devise a good heuristic to approximate it.

Heuristic 1: Yield only

Intuitively, we can try to improve circuit yield by decreasing layout sensitivity to defects, as Equations 2 and 3 describe. Therefore, recalling Equations 4 and 5, the optimization problem becomes

$$\max \prod_{k=1}^N Y_{G,k} = e^{-\sum_{k=1}^N \lambda_{G,k}} \quad (9)$$

or its equivalent additive form

$$\min |\ln(Y)| = \left| \ln \left(\prod_{k=1}^N Y_{G,k} \right) \right| = \sum_{k=1}^N \lambda_{G,k} \quad (10)$$

In practice, this heuristic would minimize the critical area (or equivalently the failure rate), disregarding the actual circuit area.

Heuristic 2: Weighted yield

The previous heuristic doesn't account for area. Introducing area as a weighting factor turns the optimization problem into

$$\min \sum_{k=1}^N |A_{G,k} \ln(Y_{G,k})| = \sum_{k=1}^N A_{G,k} \lambda_{G,k} \quad (11)$$

Heuristic 3: Manufacturability function approximation

A more sophisticated approach is to approximate the original cost function as closely as possible with an additive function. We adopt the following approximation for the optimization problem:

$$\min \frac{A_{\text{circuit}}}{Y_{\text{circuit}}} = \frac{\sum_{k=1}^N A_{G,k}}{\prod_{k=1}^N Y_{G,k}} \approx \min \sum_{k=1}^N \frac{A_{G,k}}{Y_{G,k}^p} \quad (12)$$

where p is a parameter we can choose to tune the heuristic.

An intuitive choice is $p = 1$, but $p = N$ is actually better. The reason becomes evident if we approximate all $Y_{G,k}$ terms with a common average value, Y_a :

$$\sum_{k=1}^N \frac{A_{G,k}}{Y_{G,k}^p} \approx \frac{\sum_{k=1}^N A_{G,k}}{Y_a^p} = \frac{\sum_{k=1}^N A_{G,k}}{\prod_{k=1}^N Y_a} \approx \frac{\sum_{k=1}^N A_{G,k}}{\prod_{k=1}^N Y_{G,k}} \quad (13)$$

This approximation seems reasonable, given the typical values of

$$Y_{G,k} = e^{-FR_{G,k}} \cong 1 - FR_{G,k}$$

with $FR_{G,k} \ll 1$, where FR is the failure rate.

We experimented with different values of p , and we report the results in the next section.

Experimental results

We carry out our evaluation of synthesis-for-manufacturability effectiveness in two steps: First, we show that we can achieve significant improvement with this method, without involving designers and library developers, simply by changing the synthesis cost function. Second, we report preliminary results on further reducing manufacturing cost by adding yield-optimized gates to the library.

Technology mapping for manufacturability

We evaluated our proposed heuristics on test cases from the 1993 International Workshop on Logic Synthesis (IWLS 93) benchmark suite, which we

mapped onto STMicroelectronics' 0.13-micron standard-cell library. For comparison, we used standard synthesis optimization results, based solely on area. Because yield data for the standard-cell library was not available to us, we assigned a yield value to each gate by assuming failure rate FR to be a random variable with uniform distribution (300 ppb, 500 ppb). Based on industrial experience (as described by Enrico Malavasi of PDF Solutions in a personal communication), these values correspond to failure rates for most yield loss mechanisms in typical state-of-the-art industrial processes.

We used the yield information to compute the synthesized circuits' effective yield. Figure 2 shows the heuristics evaluation results. This data confirms our expectations: We obtained the best results for failure rate reduction by optimizing for yield (function Y) or, to a lesser extent, optimizing for yield weighted by area (function $|A \ln(Y)|$). As expected, these heuristics performed poorly for area, and overall the effective yield worsened (Figure 2c). Notice that A/Y gives the same results as the standard optimization for area. The reason is that for a single gate, yield $Y = e^{-FR}$ is almost 1. We also explored giving larger importance to yield by using $p = 100N$ and $p = 1,000N$. The last four columns of the charts in Figure 2 correspond to using different values of p in the heuristic in Equation 12. Table 1 validates our prediction that, assuming yield follows a Poisson model, $p = N$ leads to the best effective-yield results.

Figure 3 (on page 197) shows the percent error resulting from the approximation described in Equation 12 for a set of IWLS 93 benchmark circuits. As we predicted, this error shows a minimum value at approximately $p = N$. The error is less than 0.1% when $p = N$ and increases as p moves away from N ; the error can reach values up to 20% when $p = 100^N$. Therefore, we use this heuristic to drive the technology-mapping algorithm in the following experiments.

Because the circuits reported so far are quite small, the effective-yield improvement was relatively small. To evaluate the impact of synthesis for manufacturability on larger circuits, we created various 10-mm² circuits, each composed of the repetition of a single benchmark circuit, and we projected the effective-yield value for each large circuit. As Figure 4 shows, effective-yield improvement was as large as 5%. To assess our methodology's validity, we repeated the synthesis experiments using different yield assignments to the library—that is, different seeds for the random-number generator. The seeds correspond to the stripes in Figure 4 (on page 198). Effective-yield improvement varies across the dif-




Figure 2. Comparison of synthesis-for-manufacturability heuristics on IWLS 93 benchmarks: failure rate (a), area (b), and effective yield (c). Data is normalized with respect to results obtained from the standard area optimization (indicated as heuristic function A).

ferent seeds because the first synthesis target is timing performance; hence, if most of the fast cells (typically on the critical path) are randomly assigned low yield values, the tool has little room for manufacturing optimization. We reiterate that designers and library developers obtain this improvement at absolutely no cost, assuming yield information is available.

Further improvement: Yield-optimized gates

Using libraries containing yield-optimized variants in addition to cells optimized for speed and area can further improve our results. To show how our synthesis flow can exploit these variants, we developed a virtual variant for each of the original standard cells in the STMicroelectronics 0.13-micron library. We modeled yield improvement by dividing the original cell failure rate, FR , by a random variable, FR_factor , with uniform distribution (1, 10). In other words, for each cell, the new variant's failure rate decreases randomly from a value in uniform distribution (300 ppb, 500 ppb) to a value in uniform distribution (30 ppb, 500 ppb). According to Heineken, Khare, and d'Abreu, manufacturability can be improved with small or no penalty in area (or even area reduction).³ We assumed that 80% of cell variants would suffer from area increase, while the remaining 20% would benefit from yield increase with no area penalty. Clearly, this is a conservative assumption, and more-advantageous tradeoffs could give us better results. We randomly chose the cells to suffer from area penalty, and we increased area by a discrete quantity (technology pitch multiplied by cell height).

We synthesized IWLS 93 benchmark circuits using the enhanced library. We used both the standard flow (area optimization) and synthesis for manufacturability (implementing heuristic A/Y^N). To project the improvement on large circuits, we again replicated the IWLS 93 circuits to form 10-mm² and 1-cm² areas. Figure 5 shows the results. The light-colored columns indicate effective-yield ratios taken from the plot in Figure 4 and correspond to using the base library. This data shows that synthesis for manufacturability effectively exploits the extended library to substantially improve effective yield.

All our experiments were based on the Poisson yield model. As discussed earlier, this model typically produces a lower bound on effective yield. Our goal in this article is to highlight the manufacturability improvement trend we expect from the new heuristics for the technology-mapping algorithm. Although we do not compute exact yield values, we expect the increase in

yield lower bounds to decrease defect sensitivity and increase exact yield values. Hence, we expect a more accurate yield model to show a similar trend.

Finally, with different yield and process models, we can use a different heuristic to better match the data, or we can simply tune the value of p in Equation 12. In fact, we've shown that increasing p increased the weight of yield in the manufacturability cost function.

Production costs

Although we can obtain effective-yield improvement at no additional cost simply by changing the synthesis cost function, we expect a larger advantage when additional cells with higher yield are available. Clearly, we should include this development cost when considering a product's return on investment. On the basis of industrial experience, we can assume extending a standard-cell library will require a few staff-months (again, from personal communication with Enrico Malavasi). To estimate the return on investment for a large-volume product, we must compare this cost with the cost reduction associated with even a small effective-yield improvement. In addition, if the same library is used on multiple products, the return on investment will skyrocket. Another factor to consider is that improving yield at the layout stage in the traditional flow would incur a similar development cost, but the benefits would not be directly reusable on multiple products.

With these considerations in mind, we're confident that our approach can increase effective yield, and in most cases, this corresponds to lower manufacturing cost. Apart from the development of an extended library, our methodology has basically no effect on the design cycle; the overhead is the time needed for running a logic synthesis step to roughly estimate the number of cells in the design. Finally, our approach lets us estimate product yield improvement as a function of process maturity. In fact, if methods (such as the Arima model¹²) of estimating the evolution of a process are available, we can repeat the optimization

Table 1. Effective-yield ratio with respect to standard synthesis optimization on IWLS 93 benchmark circuits.

| Function | Benchmark circuit | | | | |
|-----------------------|-------------------|------|--------|--------|-------|
| | s38584 | clma | C7552 | apex2 | seq |
| A | 1.0000 | 1.00 | 1.0000 | 1.0000 | 1.000 |
| $ \ln(Y) $ | 0.6100 | 0.37 | 0.3600 | 0.4600 | 0.460 |
| $ \ln(Y) $ | 0.9000 | 0.95 | 0.9300 | 0.9000 | 0.910 |
| A/Y | 1.0000 | 1.00 | 1.000 | 1.0000 | 1.000 |
| A/Y ^N | 1.0001 | 0.99 | 1.0010 | 1.0002 | 1.006 |
| A/Y ^{100N} | 0.9900 | 0.99 | 1.0020 | 0.9900 | 1.005 |
| A/Y ^{1,000N} | 0.9000 | 0.94 | 1.0009 | 0.9800 | 0.990 |

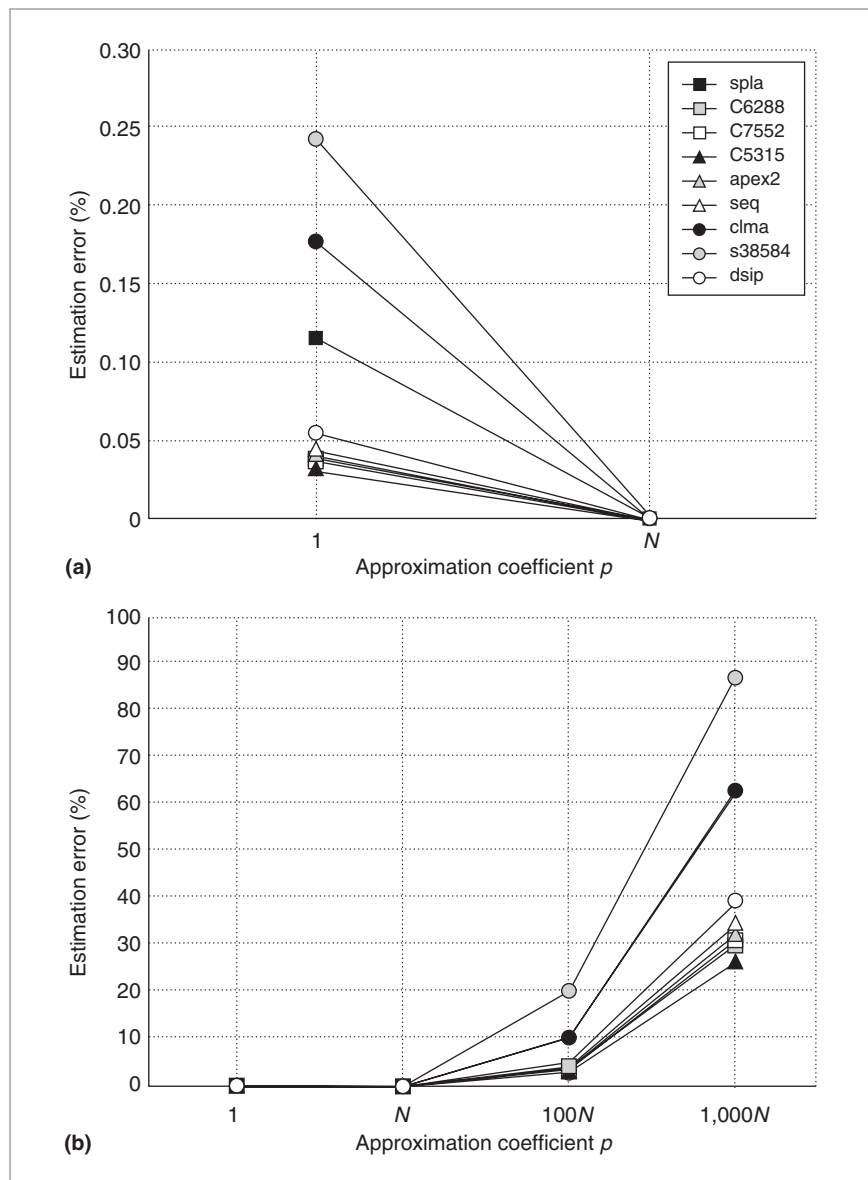


Figure 3. Percent estimation error of additive heuristic as a function of coefficient p for $p \leq N$ (a) and $p \geq N$ (b).




Figure 4. Projection of synthesis-for-manufacturability results on repetition of circuits to constitute a 10-mm² area. Different stripes correspond to different assigned random-yield values.




Figure 5. Projection of effective-yield ratio using IWLS 93 benchmarks to obtain 10-mm² (a) and 1-cm² (b) circuits. Data is normalized with respect to results obtained by using the standard synthesis flow on the extended library.

with different failure rates to project the expected savings from future process improvements.

GIVEN THE PROMISING RESULTS we obtained with our synthesis-for-manufacturability approach, we plan to further investigate this research direction and apply the methodology to industrially relevant examples. To characterize cells for yield and estimate catastrophic yield loss of synthesized circuits, we plan to use yield extraction tools. ■

Acknowledgments

We thank Claudio Pinello for many useful suggestions, and we thank Enrico Malavasi of PDF Solutions, Michael Orshansky, and Farhana Sheikh for discussions during the preliminary phase of research and during manuscript preparation. The reviewers provided very valuable feedback. This work was partially supported by the Microelectronics Innovation and Computer Research Opportunities program and PDF Solutions.

References

1. *Design Compiler User's Manual*, Synopsys, 1998.
2. *Logic Synthesis and Verification*, S. Hassoun and T. Sasao, eds., Kluwer Academic, 2002.
3. H.T. Heineken, J. Khare, and M. d'Abreu, "Manufacturability Analysis of Standard Cell Libraries," *Proc. IEEE Custom Integrated Circuits Conf. (CICC 98)*, IEEE Press, 1998, pp. 321-324.
4. C. Ouyang et al., "Maximizing Wafer Productivity through Layout Optimizations," *Proc. 13th Int'l Conf. VLSI Design*, IEEE CS Press, 2000, pp. 192-197.
5. V.K.R. Chiluvuri and I. Koren, "Layout-Synthesis Techniques for Yield Enhancement," *IEEE Trans. Semiconductor Manufacturing*, vol. 8, no. 2, May 1995, pp. 178-187.
6. V.K.R. Chiluvuri and I. Koren, "Yield Enhancement vs. Performance Improvement in VLSI Circuits," *Proc. Int'l Symp. Semiconductor Manufacturing (ISSM 95)*, IEEE Press, 1995, pp. 28-31.
7. H.T. Heineken and W. Maly, "Interconnect Yield Model for Manufacturability Prediction in Synthesis of Standard-Cell-Based Designs," *Proc. Int'l Conf. Computer-Aided Design (ICCAD 96)*, IEEE CS Press, 1996, pp. 368-373.
8. H.T. Heineken and W. Maly, "Performance-Manufacturability Tradeoffs in IC Design," *Proc. Design, Automation and Test in Europe (DATE 98)*, IEEE CS Press, 1998, pp. 563-567.

9. I. Koren and Z. Koren, "Defect Tolerance in VLSI Circuits: Techniques and Yield Analysis," *Proc. IEEE*, vol. 86, no. 9, Sept. 1998, pp. 1819-1838.
10. S.A. Shaikh, J. Khare, and H.T. Heineken, "Manufacturability- and Testability-Oriented Synthesis," *Proc. 13th Int'l Conf. VLSI Design*, IEEE CS Press, 2000, pp. 185-191.
11. D.M.H. Walker, *Yield Simulation for Integrated Circuits*, Kluwer Academic, 1987.
12. *NIST/SEMATECH e-Handbook of Statistical Methods*, <http://www.itl.nist.gov/div898/handbook/>.



Alessandra Nardi is a product engineer at Magma Design Automation. She participated in the work described in this article as a visiting postdoctoral researcher at the University of California, Berkeley. Her research interests include CAD methodologies for VLSI circuits such as design for manufacturability, modeling for signal integrity, and timing verification. Nardi has a Laurea degree and a PhD, both in electrical engineering, from the University of Padova, Italy.



Alberto L. Sangiovanni-Vincentelli holds the Edgar L. and Harold H. Buttner Chair of Electrical Engineering and Computer Sciences at the University of California, Berkeley. His research interests include design tools and methodologies, large-scale systems, embedded controllers, and hybrid systems. Sangiovanni-Vincentelli has a DrIng in electrical engineering and computer science from Politecnico di Milano, Italy. He is cofounder of Cadence and Synopsys, an IEEE Fellow, and a member of the National Academy of Engineering.

■ Direct questions and comments about this article to Alessandra Nardi, Magma Design Automation, 5460 Bayfront Plaza, Santa Clara, CA 95054; anardi@magma-da.com.

For more information on this or any other computing topic, visit our Digital Library at <http://www.computer.org/publications/dlib>.

IEEE Design & Test Call for Papers

IEEE Design & Test, a bimonthly publication of the IEEE Computer Society and the IEEE Circuits and Systems Society, seeks original manuscripts for publication. *D&T* publishes articles on current and near-future practice in the design and test of electronic-products hardware and supportive software. Tutorials, how-to articles, and real-world case studies are also welcome. Readers include users, developers, and researchers concerned with the design and test of chips, assemblies, and integrated systems. Topics of interest include

- Analog and RF design,
- Board and system test,
- Circuit testing,
- Deep-submicron technology,
- Design verification and validation,
- Electronic design automation,
- Embedded systems,
- Fault diagnosis,
- Hardware/software codesign,
- IC design and test,
- Logic design and test,
- Microprocessor chips,
- Power consumption,
- Reconfigurable systems,
- Systems on chips (SoCs),
- VLSI; and
- Related areas.

To submit a manuscript to *D&T*, access Manuscript Central, <http://cs-ieee.manuscriptcentral.com>. Acceptable file formats include MS Word, PDF, ASCII or plain text, and PostScript. Manuscripts should not exceed 5,000 words (with each average-size figure counting as 150 words toward this limit), including references and biographies; this amounts to about 4,200 words of text and five figures. Manuscripts must be doubled-spaced, on A4 or 8.5-by-11 inch pages, and type size must be at least 11 points. Please include all figures and tables, as well as a cover page with author contact information (name, postal address, phone, fax, and e-mail address) and a 150-word abstract. Submitted manuscripts must not have been previously published or currently submitted for publication elsewhere, and all manuscripts must be cleared for publication.

To ensure that articles maintain technical accuracy and reflect current practice, *D&T* places each manuscript in a peer-review process. At least three reviewers, each with expertise on the given topic, will review your manuscript. Reviewers may recommend modifications or suggest additional areas for discussion. Accepted articles will be edited for structure, style, clarity, and readability. Please read our author guidelines (including important style information) at <http://www.computer.org/dt/author.htm>.

Submit your manuscript to IEEE Design & Test today!

D&T will strive to reach decisions on all manuscripts within six months of submission.

**IEEE
Design & Test**
of Computers