System-Level Performance Modeling with BACPAC – Berkeley Advanced Chip Performance Calculator

Dennis Sylvester
University of California, Berkeley
dennis@eecs.berkeley.edu

Kurt Keutzer
University of California, Berkeley
keutzer@eecs.berkeley.edu

Abstract

In this paper, a new system-level performance model is introduced. The model, called BACPAC (Berkeley Advanced Chip Performance Calculator), is comprised of a large set of analytical models that have been newly developed and/or compiled to collectively describe the performance characteristics of future high-performance designs. New deep submicron effects are considered in BACPAC that were ignored in previous system-level performance models, such as noise and enhanced power modeling. BACPAC users can explore the capabilities of future VLSI designs given certain expected input conditions. The model can be used to predict clock frequency, chip area, power consumption, yield and other important characteristics of future systems.

1. Introduction

System-level performance models can be defined as *first-order models,* which attempt to capture the *majority of relevant system design issues* in order to provide *useful predictions or early feedback* to designers. An early example of a system-level performance model is SUSPENS (Stanford University System Performance Simulator) which was developed in the mid-1980's [1]. Since the introduction of SUSPENS, on-chip interconnect has become a primary performance bottleneck. This issue was addressed somewhat in an updated cycle-time model by Sai-Halasz, which accounts for multi-level metallization [2].

The need for new system-level performance models is illustrated by the rise of new deep submicron (DSM) effects that are limiting performance in modern designs. For instance, power dissipation has become a major issue in the design of high-speed microprocessors as well as application-specific integrated circuits (ASIC's). Previous models have only given token consideration to power issues. To be more accurate, a hierarchical approach to power estimation should be employed such as that in [3]. Also, the impact of noise is a relatively new phenomenon in digital designs. Noise can be a reliability problem as well as a timing issue. No previous system-level models have sufficiently addressed noise concerns. Finally, important topics such as wirability, yield, and new packaging technologies (e.g. flip-chip) have a significant impact on chip performance.

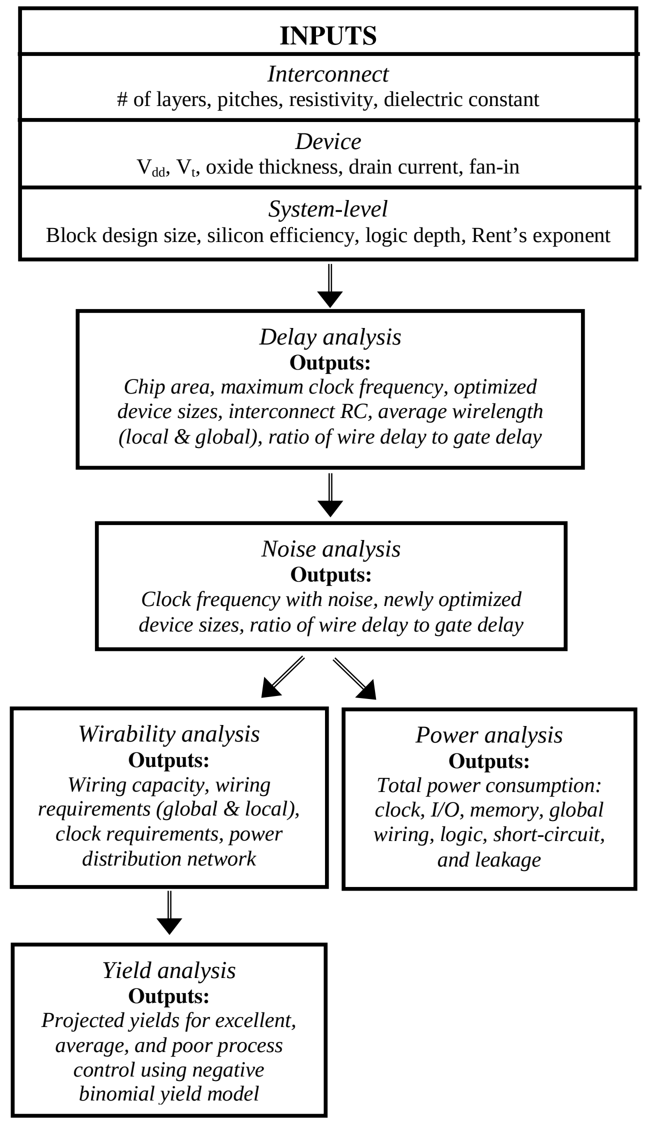

Figure 1. Overall flow of BACPAC

Berkeley Advanced Chip Performance Calculator (BACPAC) has been developed to address the above concerns. The model is available online at http://www-device.eecs.berkeley.edu/~dennis/BACPAC. An overview of BACPAC is shown in figure 1. This paper highlights some of the new models incorporated into BACPAC and presents preliminary results which point to important design considerations in future VLSI systems.

2. Delay Analysis

The drive for higher performance in ASIC's and microprocessors is often quantified by the clock speed. In order to predict operating frequency in BACPAC, a model of a critical path is specified by the user. Key device and interconnect parameters are used in conjunction with analytical wirelength estimations to optimize device sizes in terms of performance. Based on the user-defined critical path, local and global delays are computed. Clock frequency is determined by the sum of these delays as well as process variation, clock skew, and latch delays. This section explains this process in more detail.

2.1 Device and Interconnect Characteristics

The first input to BACPAC is the technology parameters for the process of interest, in particular the feature size, f_s. This value gives the program an idea for later default parameters such as metal linewidths, dielectric constants, etc. At this point, the user inputs expected values for back-end parameters such as contacted pitch, line thickness, ILD thickness, and material properties (resistivity, dielectric constant). BACPAC then proceeds to calculate the resistance and capacitance of each metal layer. The analytical capacitance formulae are from [4] based on the assumption of minimum pitch routing with upper and lower ground planes.

Device characteristics required by BACPAC include the device capacitances, both junction and input, as well as an effective device resistance value to be used in delay equations. While capacitance calculations are straightforward (see [5]), we present a novel means of determining device resistance. Driver resistance is calculated using a new technique that relates a device's current and voltage relationships directly to its effective resistance. It is a physical derivation, as opposed to empirical efforts in the past [6]. We begin by looking at the driver-interconnect network in two fashions, as depicted in figure 2. First, we can view the network as a lumped RC system where the 50% delay is defined as 0.69RC. We can also see the system as a current source charging the load capacitance. In this instance the 50% delay point is set by $C_{load} * (V_{dd} / 2) / I_{supply}$. If we assume that the driver supplies I_{dsat} throughout the voltage

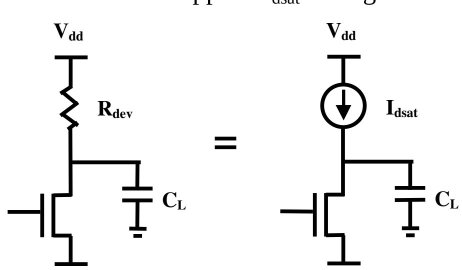

Figure 2. Basis of device resistance derivation.

swing of interest, we can equate these 2 expressions to find the effective device resistance (note that we are neglecting wire resistance, which is a good assumption in local routing with high-resistance drivers):

$$\mathbf{R_{dev}} = \frac{\mathbf{C_L V_{dd}}}{\mathbf{(0.69)2aC_L I_{dsat}}} = \frac{\mathbf{0.725}}{\mathbf{a}}\frac{\mathbf{V_{dd}}}{\mathbf{I_{dsat}}} \qquad (1)$$

In this expression, the parameter a represents the fact that I_{dsat} will not be supplied throughout the voltage swing of V_{dd} / 2. Due to velocity saturation effects, however, a value close to I_{dsat} will be conducting throughout most of the time period, so we expect a to be close to 1. From simulations and published I_d-V_d curves in the literature, we have determined a to be approximately 0.9. Solving (1) with a = 0.9, we find that device resistance for 50% delay is equal to $0.805 * V_{dd} / I_{dsat}$. Accuracy of this expression was confirmed using HSPICE. In addition, the method is extendable to calculate effective device resistance for 10-90% rise time and other arbitrary delay values.

2.2 Wirelength Modeling

Wirelength models are typically based on Rent's Rule [7]. BACPAC uses Donath's model [8], which serves as a good upper bound on expected wirelengths for a given Rent's exponent.

Gate pitch is traditionally determined by taking the square root of the chip area divided by the number of gates. According to [9,10], the size of a cell is set by the contacted metal pitch of the lower-level metals in a process. Investigating a 2-input NAND, we see the size of this cell for most standard drive strengths is 4 MP across by 16 MP high, where MP is the contacted metal pitch of bottom-level metals [10]. Other gates are either wider or narrower depending on the number of inputs and drive strength. Thus, in a typical high-performance cell library, the average gate has an area of 64 MP^2.

Gate pitch is then obtained by scaling the ideal value of $(64MP^2)^{1/2}$ using the silicon efficiency (s_e). Silicon efficiency is defined as the percentage of silicon that contains gates (typically ~ 50%). We obtain:

$$\mathbf{P_g} = \frac{\mathbf{8MP}}{\sqrt{\mathbf{s_e}}} \qquad (2)$$

If s_e is one, then the gate pitch becomes simply equal to 8 MP. When s_e is 0.5, the gate pitch is equal to 11.3 MP. From the above, average point-to-point wirelength is determined. Multiple fan-out nets are extrapolated using a modified linear empirical expression from [2].

2.3 Device Sizing

In order to limit the impact of interconnect on performance, driving gates should be sized properly. Given the average wirelengths for a range of fan-outs (typically 1 through 4), BACPAC calculates the size of a gate that no longer yields a pre-determined reduction in delay for a unit-size width increase. This method is defined by a sizing criterion, C, which is typically in the range of 2 to 5%. When an increase in gate width of a single W/L ratio does not yield a reduction in delay of C, we stop increasing the size and fix the gate width as optimal. Although most designs are interconnect-limited as opposed to device-limited [11], arbitrary increases in device sizing will lead to area penalties. These area penalties also translate to increased average wirelengths. Unless extremely small C values are used (e.g. 0.5%), this should not be a problem. Gate size optimization is discussed in more detail in [12].

2.4 Critical Path

Users are now free to create their own critical paths. They may define any distribution of fan-outs within the range of 1 to 4. The critical path logic depth is also user-defined. This is important as it can vary widely from company to company and design to design. Logic depths in modern microprocessors, for example, range from 10 to 25. Overall, however, we expect logic depth to decrease from generation to generation in the interest of getting higher performance out of DSM designs. A buffered global wire is also included in the critical path. BACPAC determines the buffer sizes based on a new area-optimal repeater expression.

The procedure to determine the global wirelength applies Donath's model to the global level. Large functional blocks of gates (e.g. 50,000 gate modules [12]) are modeled as individual gates. The size of these modules is found using the gate pitch of section 2.3 and then used to calculate the global gate pitch. Donath statistics are applied in a straightforward manner to calculate an average global wirelength. It should be noted that it is an *average global wire* and not pathological wires that are being used in BACPAC. Other system-level models typically use corner-to-corner lines, which are increasingly rare due to efficient floorplanning tools.

At this point, total cycle time is found using average wirelengths (local and global), device and repeater sizes, and a delay expression formulated by Sakurai [13]. We add a term to this expression to account for input rise-time dependency based on [14]. This term can account for 30% or more of the total stage delay. Latch set-up times and propagation delays are also included in the overall clock cycle and these times are based on ASIC libraries.

The impact of noise on timing is now considered. The traditional approach to this problem, based on Miller's theorem, doubles the coupling capacitance term of the total load capacitance. This doubling accounts for the negative gain between nets switching in opposite directions. This approach can be quite accurate for nets which do not have significant line resistance. Noise modeling in BACPAC assumes that all stages of the critical path have worst-case noise features. Namely, all gates drive lines which have two neighboring parallel minimum-pitch wires simultaneously switching in the opposite direction as the victim line. These approximations provide pessimistic noise results, which will serve as a lower bound on the speed of the design.

3. Power Analysis

In this section, we outline a hierarchical approach to power modeling based on [3]. This approach isolates the different components of power consumption on a chip and then attempts to model them individually. In contrast to [3], we look at additional forms of power consumption – leakage and short-circuit power.

- *Standard cell logic and local wiring*

We assume a hierarchical layout style like that discussed in [12], where the designer uses blocks of 50,000 to 100,000 gates to limit the impact of interconnect while maintaining reasonable flexibility and design times. Looking carefully at one of these blocks and using the previous discussion of gate pitch, we calculate the physical size of each block. To calculate the interconnect capacitance, we determine the wiring requirements within the module and use the resultant total length with an average capacitance per unit length. The wiring requirements of a block are found by extrapolating the critical path model to the entire module.

Total switching device capacitance consists of oxide, overlap, and junction capacitances. In addition, we consider the impact of internal nodes of a multi-input logic gate. From the delay analysis section, BACPAC has determined the optimal device sizing in local routing. Calculations are made for both with and without noise scenarios, so that comparisons can be made in terms of power vs. delay tradeoffs. All logic gates are assumed to be the optimal size -- this may result in an overestimate of sizing in some non-critical cases but will also underestimate some device sizing when longer wires are present in certain paths. Overall, we expect that this assumption will be slightly pessimistic, leading to somewhat higher power numbers.

- *Global interconnect*

Once a typical global interconnect is found, we seek to determine the number of global wires in a design. Based on Rent's Rule (and defining a global interconnect as a wire that runs between or among blocks), we calculate the number of global nets in a design. The total number of repeaters needed is a function of the ratio between average global wirelength and critical length and is used to determine the total buffer capacitance.

- *Clock distribution*

BACPAC clock distribution models are based on the buffered H-tree. The cluster size is set by the maximum

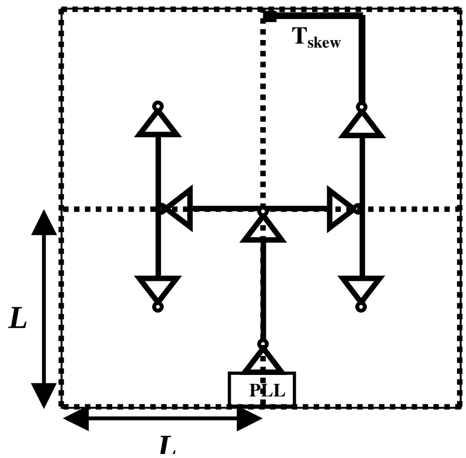

Figure 3. Buffered H-tree with maximum local skew corresponding to line length *L*.

amount of skew allowable in the design. The total distance that must be traversed in this instance is L, where L is the side-length of a cluster. In order to find the maximum L that meets the T_{skew} requirement, we use Sakurai's delay expression [13] to find the delay between a point just at the output of the buffer and the point in the corner of the cluster (see Figure 3). In this case, we expect the dominant source of delay to be the charging of latch capacitances along the wire resistance. Since these latches will not be lumped at the end of the line, but distributed along the length L of the wire, we model this RC component as distributed, which differs from Sakurai's initial formulation. In addition, 10% process variation is assumed for both device and interconnect characteristics, which serves to reduce L. The analytical formulation gives results within 10% of simulations.

Once a maximum cluster size is determined, the required number of drivers and total wirelength can be estimated based on the regularity of the H-tree structure, leading to a total clock load capacitance. Since the activity factor for clock networks is 1, the power consumption from this component can be large.

- *I/O's*

The power modeling of I/O drivers is relatively straightforward. It consists of first determining the number of signal pads for a design based on Rent's rule. Once this value is determined, the pad is optimally driven by a cascaded chain of inverters where each stage is larger by a constant factor. We make an important enhancement to the previous approach by developing a more accurate model for the total driver chain capacitance which accounts for the intrinsic output capacitance of the buffers. HSPICE simulations have confirmed the validity of this new model.

- *Memory (on-chip caches)*

Memory power is estimated by examining the capacitances being charged in read and write operations. Detailed expressions similar to those in [3] are developed which take into account the lower signal swing typically seen for bit lines in on-chip memories.

- *Leakage*

In order to track the increasing importance of leakage current in future designs, we have implemented leakage power models into BACPAC. The calculations are based on an expression for static current in an MOS device [15]. The total chip-level static power consumption is found by determining the total device width in a design. Leakage current from both NMOS and PMOS devices are of interest here, so the total device width is calculated. Ignoring memory, we estimate device width by including repeaters, logic gates, and clock drivers.

- *Short-circuit*

Based on [16], BACPAC calculates short-circuit power consumption by determining the length of time that the short-circuit current flows (t_{base}) as well as the peak value of the current (I_{peak}). Combining delay expressions from [13] with the device resistance method explained earlier, t_{base} is found for each fan-out configuration. I_{peak} is found using the alpha-power law model for I_{ds} to estimate drain current at $V_{gs} = V_{dd} / 2$ [14]. Short-circuit power dissipation per gate is found which is then extrapolated to the entire design by focusing on logic gates only.

4. Wirability Analysis

This section presents models to calculate available routing resources. In this discussion, we will look only at the logic portion of a design as the wiring requirements for memory are generally much lower than that for logic. Ideal routing capacity is reduced by several factors which have been ignored or roughly approximated in previous system-level performance models, including power distribution, via blockage, and clock routing.

The power distribution network uses significant routing resources, especially at the top layers. BACPAC first determines power grid dimensions by fixing the peak IR voltage drop at 4% of V_{dd}. Analytical expressions are derived to describe the IR drop of an arbitrary layer as a function of metal linewidth. When the reliability constraints are met, the percentage of routing resources used is calculated. BACPAC allows for both wirebonding and C4 flip-chip packaging in constructing a power grid. Results indicate that wirebonding is a non-scalable technology due to rising die sizes, dropping noise margins, and increasing current levels.

It has been estimated empirically that for metal layers with equivalent pitches, an upper layer blocks 12-15% of an underlying layer due to its need to connect to the substrate using vias [2]. However, when larger metal pitches are used on higher levels, the amount of blockage can be reduced if we are using relatively fixed-size vias.

The relationship between via blockage and metal pitch can be shown to be linear. With a multilevel interconnect system, vias connecting the top layer to the substrate necessarily block *all* underlying levels. Thus, metal 1 is blocked by all subsequent layers resulting in a sizable loss in metal 1 routing capacity. Fortunately upper layers usually have significantly larger pitches than bottom levels, reducing the via penalties associated with multilevel interconnect. BACPAC calculates exact via blockages for each layer according to the user-specified metal pitches.

Clock distribution also serves to reduce available routing resources. Due to the regularity of H-tree structures, the total wiring required for such a network can be found fairly accurately. Given the number of clusters in the H-tree (see section 3), the total wirelength is given by a simple analytical expression in terms of the chip-side length. Additional "within-cluster" routing is approximated using a heuristic that allocates wires from the central driver to the perimeter of the cluster in all directions.

Routing tools cannot fully utilize all the available routing resources for a given design. This is mainly due to the algorithms used within the routing tools. BACPAC models this effect by first calculating the available routing resources after clock routing, power/ground routing, and via blockages. At this point, the routing area is multiplied by a routing efficiency factor to give the estimated available routing resources. This routing efficiency factor is set at 0.5 in this work. This value is based on discussions with CAD engineers from industry.

5. Results and Discussion

BACPAC has been used to investigate future VLSI systems, with an emphasis on microprocessors. Default parameters are used in many instances – defaults use conservative values for parameters such as drain current, threshold voltage, dielectric constant, etc. Looking at Figure 4, one can see that BACPAC default predictions are very close to those from the 1997 NTRS for global clock frequency [17]. Logic depth is 15 throughout technologies in the default case. When the logic depth is reduced (by 1 per generation), NTRS expectations can be exceeded with significant improvements seen at sub-0.1 μm technologies. The impact of noise is seen to be about an 8 to 10% drop in critical path speed. The main effect of noise on system design is the need for enhanced current drive to offset the higher overall capacitances. The larger devices lead to more power dissipation in addition to a delay penalty.

Figure 5 looks at performance in terms of MHz / W and finds that increases in transistor current and module design size (from 50,000 to 100,000 gates / module) do not appreciably alter performance. However, the use of silicon-on-insulator (SOI) technology does yield significantly higher performance due to its lower device capacitances and improved leakage characteristics.

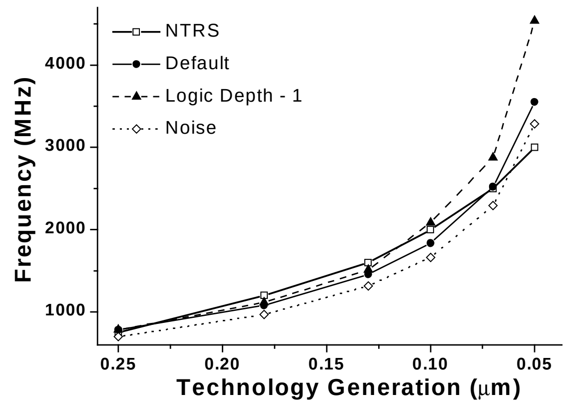

Figure 4. Anticipated microprocessor operating frequencies.

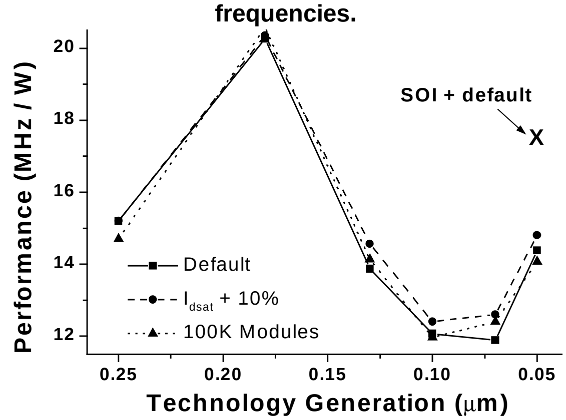

Figure 5. Speed-power metric demonstrating only a minor advantage for higher I_{dsat}.

Global wires are examined and the use of scaled global wires (reducing pitch each generation) is found to have detrimental effects on both speed and power (due to excessive buffering). As explained in [18], the use of "fat" [2] or unscaled global wires will be necessary to achieve performance goals without heavy power and delay penalties.

Total power consumption is seen to rise quickly in Figure 6, exceeding NTRS predictions beyond 0.1 μm. Power calculations are based on the assumption that 70% of the devices in a design are allotted for memory. With billion transistor designs by 50-nm, a larger portion of the device count may be memory-related. This would drop BACPAC power estimations. In addition, leakage is seen to become significant with scaling V_t values, even at a conservative V_t of 150 mV at the 50-nm technology node. For instance, at 50-nm, a generic microprocessor could see ~ 6% of its total power go to static consumption. Clock distribution uses a buffered H-tree, which generally lead to lower power consumption than grid-based designs.

A wirability analysis of a generic 50-nm microprocessor has found an optimal wiring architecture to consist of nine metal levels, six of which are dedicated to global wiring where global wiring is defined as routing

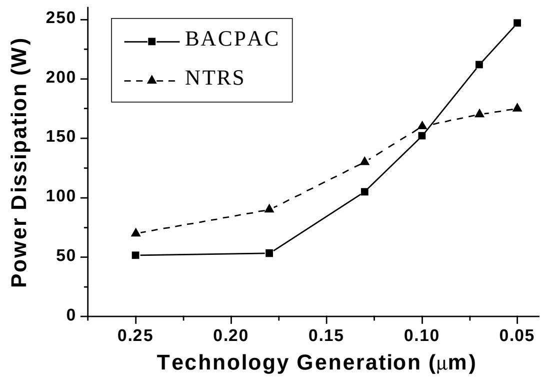

Figure 6. Low-power design techniques must be used in DSM to achieve manageable power.

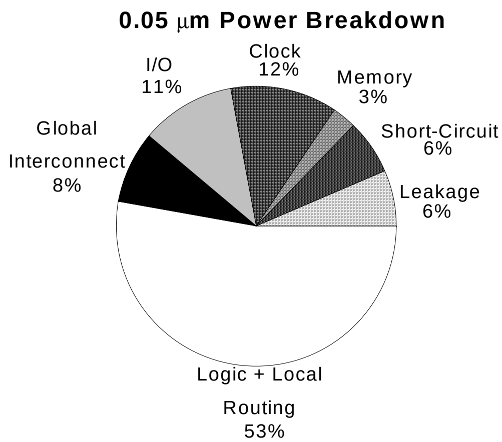

Figure 7. Logic and local wiring dominates power consumption. Leakage becomes appreciable due to dropping V_t.

among 50,000 gate modules. While local wiring dimensions will remain to be set by lithography limits, global wires will need to be carefully optimized to allow for pathological routes while supplying sufficient routing resources for the entire design.

5. Conclusions

BACPAC was introduced to account for new DSM effects within the framework of a system-level performance model. It can be used to investigate trends with respect to various input parameters, whether they are device, interconnect, or system-level properties. Results from BACPAC indicate that power consumption will be a primary concern in future microprocessor designs. The global wiring hierarchy will be a key to achieving future performance goals. In the interests of brevity, many features of BACPAC were not included in this paper. For instance, yield modeling is included in BACPAC, based on a negative binomial yield model as well as a new method of estimating critical area. For further information on these topics and relevant equations, readers are referred to http://www-inst.eecs.berkeley.edu/~dennis/bacpac/bacpac_models.html.

Acknowledgements

The authors wish to acknowledge William Jiang of UC-Berkeley for his work on implementing BACPAC online.

References

[1] H.B. Bakoglu and J.D. Meindl, "A system-level circuit model for multi- and single-chip CPU's," *Proc. of International Solid-State Circuits Conference*, pp. 308-309, 1987.

[2] G.A. Sai-Halasz, "Performance trends in high-end processors," *Proc. of the IEEE*, vol. 83, pp. 20-36, Jan. 1995.

[3] D. Liu and C. Svensson, "Power consumption estimation in CMOS VLSI chips," *IEEE Journal of Solid-State Circuits*, vol. 29, pp. 663-670, June 1994.

[4] J.H. Chern, J. Huang, L. Arledge, P.C. Li, and P. Yang, "Multilevel metal capacitance models for CAD design synthesis systems," *IEEE Electron Device Letters*, vol. 13, pp. 32-34, Jan. 1992.

[5] N. Weste and K. Eshraghian, *Principles of CMOS VLSI Design*, New York, Addison-Wesley, 1993.

[6] P. Fisher and R. Nesbitt, "The test of time: Clock cycle estimation and test challenges for future microprocessors," *IEEE Circuits and Devices Magazine*, pp. 37-44, Mar. 1998.

[7] B.S. Landman and R.L. Russo, "On a pin versus block relationship for partitions of logic paths," *IEEE Transactions on Computers*, vol. 20, pp. 1469-1479, Dec. 1971.

[8] W.E. Donath, "Placement and average interconnection lengths of computer logic," IEEE Transactions on Circuits and Systems, vol. 26, pp. 272-277, April 1979.

[9] R. Payne, "Metal pitch effects in deep submicron IC design," *Electronic Engineering*, pp. 45-47, July 1996.

[10] IBM SA-12 ASIC family databook, revision 5, June 1998.

[11] R.W. Keyes, "The wire-limited logic chip," *IEEE Journal of Solid-State Circuits*, vol. 17, pp. 1232-1233, Dec. 1982.

[12] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron," *Proc. of International Conference on Computer-Aided Design*, pp. 203-211, 1998.

[13] T. Sakurai, "Closed-form expressions for interconnection delay, coupling, and crosstalk in VLSI's," *IEEE Transactions on Electron Devices*, vol. 40, pp. 118-124, Jan. 1993.

[14] T. Sakurai and A.R. Newton, "Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas," *IEEE Journal of Solid-State Circuits*, vol. 25, pp. 584-594, April 1990.

[15] C. Hu, "Device and technology impact on low power electronics," in *Low Power Design Methodologies*, ed. Jan Rabaey, Kluwer Publishing, pp. 21-35, 1996.

[16] H.J.M. Veendrick, "Short-circuit dissipation of static CMOS circuitry and its impact on the design of buffer circuits," *IEEE Journal of Solid-State Circuits*, vol. 19, pp. 468-473, Aug. 1984.

[17] National Technology Roadmap for Semiconductors, Semiconductor Industry Association, 1997.

[18] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron II: The global wiring paradigm," *Proc. of International Symposium on Physical Design*, submitted, 1999.