Designing Mega-ASICs in Nanogate Technologies

David E. Lackey
IBM Microelectronics Division
1000 River Street
Essex Junction, VT 05452
delacke@us.ibm.com

Paul S. Zuchowski
IBM Microelectronics Division
1000 River Street
Essex Junction, VT 05452
paulz@us.ibm.com

Juergen Koehl
IBM Microelectronics Division
1000 River Street
Essex Junction, VT 05452
koehl@de.ibm.com

ABSTRACT

This paper discusses challenges the designer faces in integrating entire system product designs, containing tens or even hundreds of millions of logic gates, into single chip solutions now within reach using circuit densities possible in the latest silicon technologies. Managing designs of this size presents a new dimension of issues, and managing the physical and electrical effects of these high density device geometries presents another; solutions in both these areas are presented. Lastly, this paper discusses the integration of multiple functional components (previously organized as systems of multiple chips from multiple design sources and technologies) into a single chip product.

Categories and Subject Descriptors

B.7.1 [**Hardware Types and Design Styles]** Advanced Technologies, VLSI. B7.2 [**Hardware Design Aids**] Layout, Placement and Routing, Verification, Synthesis. B.8.1 [**Reliability, Testing, and Fault Tolerance**]. B.8.2 [**Performance Analysis and Design Aids**].

General Terms: Algorithms, Management, Measurement, Design, Economics, Reliability,

Keywords: Time to Market, Design Productivity, Methodology, Power Management, Signal Integrity, System-on-Chip

1. INTRODUCTION

Silicon technologies are advancing from recent 40 Million gate .13 micron, to current 80 Million gate 90 nanometer technology designs to, in the near future, 65 nanometers enabling chips over 100 Million gates. This enormous growth in gate capacity has led to unprecedented capability for design size, functional integration, and complexity on a single chip. Single chips are now replacing multiple chip packages and entire systems. Managing this functional capacity and complexity drives three major thrusts:

- Design productivity and quality through design methodology
- Design performance and density enablement through circuit techniques and physical architecture
- Functional complexity management, and combining IP from multiple sources and multiple technology platforms, through System On Chip (SoC) integration methods.

DAC 2003, June 2-6, 2003, Anaheim, California, USA.

Time-to-market pressures on the designer prohibits a proportional increase in product schedule with the size of the design, requiring increased design productivity, and decreasing turn-around time (TAT) of a given design size, to keep pace with the increased design capacity of today's chips. The designer balances the *pros* and *cons* of flat and hierarchical design methods to optimize TAT. Similarly prohibitive to achieving product schedules is design redo, and thus chip releases (tape outs) that are First Time Right become increasingly important. Coupling the increasing market demand for First Time Right with far greater design size leads to a requirement for design quality that improves exponentially.

Ever-increasing density has drawn shrinking circuit geometries toward an array of fundamental limits produced by electrical and material effects. In .18 micron technologies, interconnect delay began to overwhelm circuit or gate delays. In .18 and .13 micron technologies, capacitive cross-coupling produced inter-signal delay effects and signal integrity issues that could no longer be ignored. Active power consumption, and the power contribution of static leakage current, adds yet more complexity to high performance design, and to power supply design and distribution. Designing digital logic amidst these issues has led to varied and innovative circuit designs, physical chip architectures, and methodologies, as well as new tool flows for detecting, repairing, and preferably avoiding signal integrity issues. The need to combine gate placement with logical optimization, using placement-based synthesis approaches [1], has gone from industry innovation to business-as-usual in a few short years, as new design has migrated from greater than .25 micron technologies to .18 micron and lower.

Managing functional complexity of 40 Million to over 100 Million gate designs drives the product developer away from new functional design for each new product, toward reuse of existing pre-designed and pre-verified functional "islands," and toward efficient methods of integrating these islands into a functional system. These reused functions may be contributed by designers of multiple schedules, geographies, and companies. Additionally, the individual optimization of each island may vary the physical design points within a single chip through multiple placement, power, and input/output (I/O) structures, multiple circuit families, multiple voltage levels, and varying implementation platforms.

2. DESIGN PRODUCTIVITY

Design complexity, measured in unique logic gates per chip, is increasing as shown in Figure 1, and clock frequencies have reached as high as 900MHz for the core logic and more for high-speed interfaces [2]. Time-to-market pressures continue to drive chip design TAT requirements downward, while the design

capacity on a chip increases. The result is an exponential increase in the required design productivity.

Fig 1. Increasing Design Size

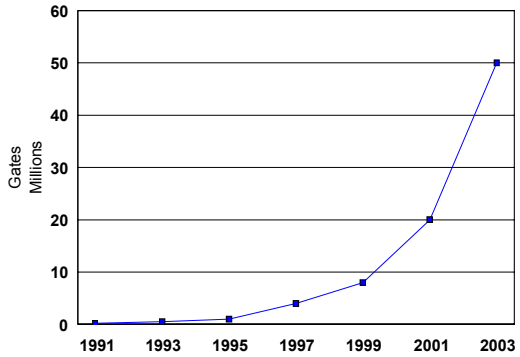

2.1 Design Process and Tools

Reducing TAT is a key focus of design process, or design methodology, and needs to be addressed in two ways simultaneously:

- Design process focus
- Tool integration focus

Design process focus is the realm of continuous incremental improvement. Analysis of every designer-executed step, iteration and redo of steps, the time of and between each step, needs to take place. For example, IBM's ASIC group has been able to reduce the number of steps executed by the designer from around 200 to 130 in one year alone based upon such analysis and subsequent changes in the design methodology [2,3]. These changes include encapsulating sequences of multiple steps into one, and moving the discovery of problems that cause redo to a point earlier in the cycle or eliminating them entirely. Tool development by IBM [4-6] and it's research partners [7-8] has substantially reduced the time required for complex design steps such as layout timing closure, as has the application of faster and highly-parallel CPU's to such performance-intensive steps.

Tool integration focus is the realm of tool and methodology development: the coupling of previously-unique design steps and algorithms into a single algorithm. Additionally, it is the careful selection of designer-driven steps for automation: continuing to leverage the designer's knowledge and decisions in the design process, while automating the sequences that take place between decisions. Finally, integration drives the measurement and fixing of problems discovered late in the design cycle (e.g. a repair action), to incorporating these measurements into the tool processes which create the design initially, thus preventing the problem from ever being introduced (e.g. an avoidance action).

Through the combination of many of these approaches, IBM's TAT for final design version was reduced by 42% in one year [2].

IBM's ASIC Design Methodology [9] has deployed numerous and sweeping examples of successful integrations within the tool flow, including:

1. Placement-based synthesis tool flows for early [1] and late timing closure [4] merged the operation of gate placement (where interconnect timing estimates can be highly accurate) with synthesis (where logic timing optimization is performed). Integration of synthesis and placement was extended to wiring congestion avoidance [10], and timing-driven global routing.

2. Cross-coupling of nets necessitates the detection of timing changes or possible false switching due to activity in a nearby-routed net. IBM is rapidly moving these issues from the realm of final timing analysis and repair, to avoidance methods in the global route and the placement-based synthesis processes.

3. Incremental timing [11], and combining multiple cross-chip process variations into a single path analysis, to reduce the number of timing runs. This is now evolving into statistical timing approaches [12] to account for device-specific process variation.

4. Design Planning methodologies [13] move the designer's decisions for logical and physical partitioning, floorplanning, and timing closure from later in the design flow (when changes require larger TAT) to an earlier timeframe (when changes can be made rapidly). Design Planning further drives increased automation into the final stages of the flow, reducing the schedule's critical path.

2.2 Flat and Hierarchical Design Methods

Managing TAT and design productivity leads, for each chip design, to evaluation of tradeoffs between flat and hierarchical design methods. Differences within the chip design, as well as the organization of the design project, affect how the *pros* and *cons* of flat and hierarchical design apply. In many cases, a combination of both approaches leads to the fastest solution.

Flat design allows the complete chip design to be solved as a single placement and routing problem. The ability to globally optimize placement and logic for the entire design allows for paths between synthesis partitions to be optimized. Avoidance of hard physical partition boundaries can lead to higher utilization of the chip.

Hierarchical design requires partitioning of the design, and can constrain optimization of the physical design. It can, however, be a powerful technique for design architectures with natural functional "islands," and can be particularly leveraged when different design teams work on different islands at the same time, running these smaller designs in parallel. Partitioning can localize the problems of timing closure and wirability, and minimize the issues of global timing and wiring congestion. If the final design change only affects one or a few partitions, the entire design may not have to be reprocessed. Hierarchical design, on the other hand, requires additional design steps including partitioning, partition pin management, planning wiring resource between the partitions and the top level, integrating the partitions, and resolving global timing and wiring issues at the top-level.

IBM has been able to leverage its tool capabilities in applying a number of combined flat and hierarchical techniques for optimal benefit on many designs:

Partial Hierarchy: A critical logic partitions is designed as a hierarchical block, but the remainder of the chip is easily closed as a flat design.

Hybrid Hierarchy: Placement and timing closure is partitioned, but routing is done flat. This allows flat versus hierarchical tradeoffs to be made separately for placement and routing. Further, this allows the additional wiring-related design steps for hierarchical routing to be avoided. With incremental wiring, a final change for only a single block has required only localized rewiring despite the fact that routing was initially flat.

Soft Hierarchy (region constraints): Localized placement objectives, such as timing closure and placement density, are enabled by logical partitioning. However, the partitioning is soft, avoiding most of the hierarchical design steps.

Figure 2 shows a chip designed with a mixture of flat techniques, and hard and soft hierarchical techniques.

Fig 2. Integrating hard and soft hierarchy with flat design

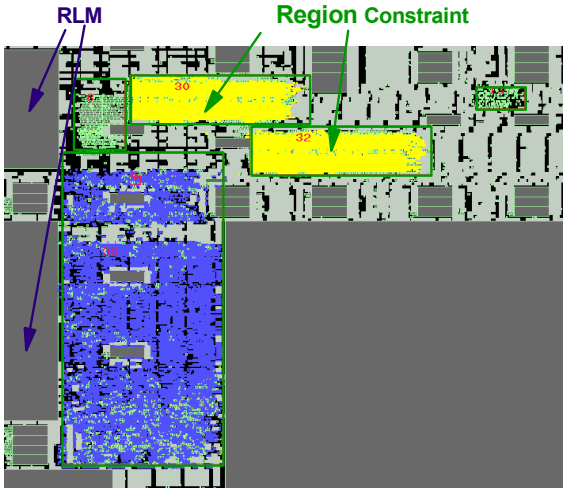

3. DESIGN QUALITY

IBM's First Time Right ASIC strength has fueled its leadership in design quality. The foremost benefit of design quality is elimination of design redo of the manufactured chip, where schedule impact of redo is most costly. Design quality is partly based on error-free execution of the design process; however, the key technical aspects of the design process itself have provided IBM with this quality leadership:

- Static timing analysis and timing modeling characterized to the logic circuit and physical chip implementations, and tuned to the target manufacturing processes [14].
- Race-free full-scan Design-For-Test structures, with full boundary scan, enabling completely automated test and diagnostic pattern generation [15].
- Correct-by-construction physical templates (images) of the chip that provide robust power distribution, signal and power I/O locations, and locations for logic placement [14].
- Technology- and manufacturing-specific checking of the logical and physical implementations [9].
- Equivalency checking to ensure the final logical implementation is the same as that provided originally by the ASIC designer [16].
- Broadened timing analysis to detect and eliminate issues due to cross-coupled [17] noise and power supply drop. Now, noise-avoidance methods are applied to the global routing and placement steps.
- Extended Design-For-Test techniques able to provide the increased test data volume of huge gate counts, and able to identify delay-based defects, the need for which increases with decreasing circuit and process geometries [9].
- Automated image generation to allow specific permutations of the predefined image types, including chip size, I/O types, power structure, signal pre-wiring, and multiple placement terrains [14].

4. POWER MANAGEMENT

Power consumed by CMOS circuitry is driven by active power, whose primary component is dynamic signal switching, and static power which is produced mainly by leakage current.

Active power can be expressed as

$$\mathbf{Pactive = C * Vdd2 * F} \quad [18].$$

Whereas each successive technology generation decreased the Vdd requirement by around 30%, this has been offset by a corresponding 30% increase in capacitance per unit area. But given the increasing required frequency of product design by generation, the overall effect is increased active power [19].

The most dominant component of leakage current is the circuit's sub-threshold transistor current [20]. Transistor performance has been increased through reduced oxide thickness (Tox), which for reliability requires a drop in Vdd and a corresponding drop in threshold voltage (Vt) to provide performance. The combined reduction in Tox and Vt increases leakage current, which 90nm technologies has emerged to equal importance with active power [21], as shown in Figure 3.

Fig.3 Active/Leakage Power Density by Lpoly width [21]

(reprinted by permission)

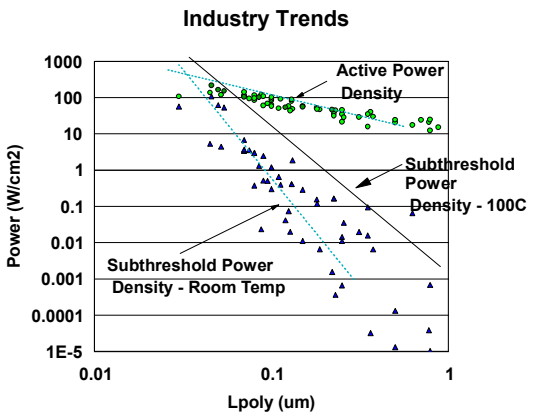

Managing power consumption in an ASIC design can be addressed at multiple levels:

- Circuit or library level
- Logic design level, including design optimization
- Architectural level

Multiple circuit libraries providing options for Vt can provide logic design and optimization options for increasing performance in a given logic path at the expense of increased static power (low Vt), or conversely reducing power in a path that meets, with extra margin, the required performance (high Vt). Taking advantage of multiple Vt options is in the realm of the logic designer and the synthesis and layout optimization tool flows. Further, Vt libraries can be assigned architecturally to entire functional partitions, providing high performance yet high leakage for the highest performance applications, while reducing power (with high Vt) for logic blocks that can operate at lower performance. When leveraging multiple Vt, optimization tools must consider the maximum allowable leakage current for the chip in the test and product environments. Further, integrating multiple transistor design points (multiple Vt libraries) within a single chip extends the need to consider device-specific process variation in the timing signoff tools [12].

Logic design techniques for reducing active power include drive strength reduction for non-timing-critical logic paths, glitch- free combinational logic, disabling unobserved combinational blocks [22], gating the clock locally for registers that retain logic state across several cycles, allowing clock skew to reduce simultaneous switching, and double-edged clocking [23].

Managing power at the architectural level can provide significant leverage in reducing chip power, and will be discussed in more depth later in this paper.

5. NOISE AND SIGNAL INTEGRITY

Coupled noise was the most problematic form of noise in digital designs using the 0.18 and 0.13 micron technology nodes, and design methodologies have been developed for avoiding, detecting, and fixing coupled noise problems [17][24]. IR drop (both AC and DC) will become the predominant noise problem in 90 nm. While power densities have increased or remained the same due to thermal considerations [19], the supply voltage has continued to scale. This results in more current-per-unit-area on the chip. Metal lines have also scaled, raising resistance in the on-chip power distribution. Additionally, transistor threshold voltage has not scaled due to the exponential increase in leakage current that would result, thereby resulting in circuits that are more sensitive to IR drop, due to decreased $(V_{GS} - Vt)$.

The analysis of the ASIC power supply system requires knowledge of the power distribution design of the card, the package, and the chip. Unfortunately, the power supply response of the system incorporating the ASIC is dependent on functional patterns, and a representative pattern set is rarely available. This forces the ASIC designer to develop a robust power distribution to minimize IR drop due to power consumption. The physical design methodology needs to consider not only the design of the power grid, but also the placement of high-current cells, and the number and location of decoupling capacitors needed for reducing the effects on the power distribution.

Reliability wear-out mechanisms that were safely guard-banded against in the past, such as negative bias temperature instability (NBTI) and hot carrier effects, must be considered during the design of the ASIC to ensure proper function over the life of the part. Both NBTI and hot carrier effects result in degraded transistor performance over time. The effects, as seen in .13 micron and 90 nm silicon processes, can result in significant delay changes. Unfortunately, the delay change is not uniform for every path on the design due to differences in the path delay components (wire-dominated vs. circuit-dominated, rise-delay vs. fall-delay, etc). This implies that the designer needs to worry not only about logic cycle time but also the differential variation between two paths with common dependencies (setup and hold checks). The amount of margin in the design will vary as a function of time as the paths degrade at different rates, and this needs to be accounted for in the timing analysis.

Unlike NBTI, the impacts of hot carrier effects are a function of individual node switching activity. Clocks degrade more than logic because they switch more often, and clock gating can actually result in the creation of additional clock skew as the design ages.

6. SoC INTEGRATION

A System-on-Chip (SoC) can be characterized both by

- Large design size as measured in gates
- Integration of functional blocks

The issues of designing chips of large design content have been the subject of the paper up to this point. Integrating a mix of functional blocks into a correctly-functioning SoC, on the other hand, may likewise be faced with integrating a mix of design attributes and development status:

- Design flow state (RTL, netlist, placed gates, full layout)
- Sources (organization, geography) and implementation schedule of each function
- Performance, power, clocking, connectivity, and physical area requirements.

6.1 Functional SoC

In the past, it made sense to deliver libraried designs or IP (intellectual property) in a synthesized netlist form. This provided a reasonable assurance that when incorporated into the chip design, the function would operate at the performance determined in the original logic synthesis. For digital IP with higher performance requirements, a fully implemented *hard core* is pre-defined and integrated as a physical block in the target SoC.

However, the increased dominance of interconnect delay (as described earlier) has made it far more difficult to optimize a function to a given performance level outside the context of the intended chip design. Further, the success of placement-based synthesis methodologies generally leads to best results by integrating logic synthesis and placement [1] within the context of the target chip. With these factors in mind, the following levels of librarying and integrating predefined digital functions emerge:

- *Standard performance*: Provide RTL that can be synthesized and placed together with the SoC's remaining RTL
- *Higher performance*: Provide synthesized and placed gates that can be flattened into, or integrated hierarchically into, the SoC's floorplan.

- *Highest performance*: Provide fully laid-out hard cores, to be integrated into the SoC's floorplan.

With the need to implement an SoC comprised of functional blocks from multiple sources with possibly different development schedules, comes the need for project management and methods to integrate design data and flow that are able to deal with this complexity. IBM's eDesign ™ and TheGuide™ are used and will further evolve to support increasingly collaborative and distributed design [9], and multi-threaded hierarchical design flows whose required design steps vary by functional block.

System-level design planning tools and methodologies will need to extend from their existing physical floorplanning features, into the realm of functional architecture, integrating both [25]. The needs include:

- Partitioning for power and path performance
- Determination of clock domains, including frequency and physical distribution
- Architectural performance modeling and functional/physical pipeline planning
- Characterization and/or abstraction of the above attributes, and their application in high-level SoC design.

6.2 Technology "Islands"

Latest technology generations have provided far greater functional integration on a single chip. As shown in Figure 4, this integration brings together functional components implemented in varying circuit families and/or physical layout architectures, varying voltage operating points for specified operation, analog and digital designs, and varying design platforms such as standard cell and FPGA. Diverse design requirements, different optimal design points, varying flexibility, and product schedule preclude redesigning all functions into a common homogeneous physical structure. And of course, power management is of ever- increasing importance. Further, the diverse manufacturing test requirements of the functional components must be integrated into a single test process for the chip.

Fig 4. Integrating Technology Islands into an SoC

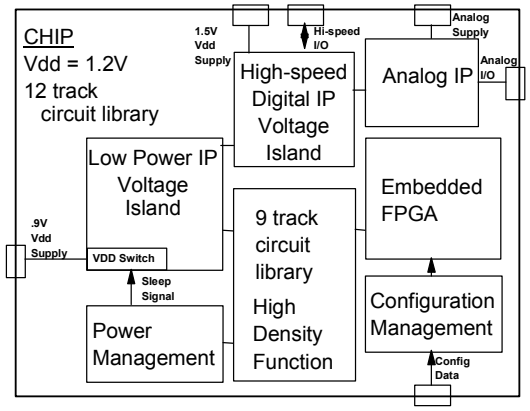

Voltage Island [26] techniques provide a functional block with a voltage source that can be unique from other functional blocks of the chip. An SoC comprised of multiple Voltage Islands can provide each functional block the specific voltage level needed to meet required performance. Therefore, substantial power savings can be realized for functional blocks of lower performance and thus lower voltage requirement, even when there are other, much higher performance functions on the chip. The power to a Voltage Island can be uniquely switched, whereby an SoC of multiple Voltage Islands need only provide power to active functions, a capability valuable to low power or battery-powered applications.

Mixed Terrain techniques allow each functional block to use a circuit library and a corresponding circuit placement row pattern optimized to the performance and wirability requirements of the function. For example, a lower performance function with low wiring congestion can be implemented in a physically-smaller block by using a lower-track circuit library and placing these circuits in a high density structure of corresponding circuit row sizes.

While First Time Right methodology takes much of the risk out of the physical and electrical implementation of an SOC design, there is still a chance that a logic error may be introduced by the logic designer. While gate array backfill can reduce a logic fix to a simple wire change, the use of an embedded programmable FPGA block can further mitigate the risk of error. By implementing "risky" logic in a programmable FPGA, a logic error can be repaired without the cost and schedule impact of a chip re-spin. Embedded FPGA logic is slower and less area-efficient than standard cell logic, so architectural planning is required to leverage this capability [27].

Hierarchical physical design approaches become necessary for integrating a design of functional/technology "islands" [9]. Uniquely by island, a physical architecture can be defined including circuit row topology, power distribution structure, Vdd supply, transistor threshold and/or transistor voltage bias supply, and island-specific circuitry for voltage level shifting, voltage regulation and switching [26], capacitive decoupling, and electrostatic discharge.

To provide these approaches, a highly-flexible methodology for detailed design planning [25] becomes extremely important for managing and automating the implementation and verification of chip structures needed to integrate all these physical design methods into a single SoC. Further, early design planning methods such as functional partitioning and architecture-level timing and power analysis [13], must be extended to assist the SoC designer in making effective use of these integrated approaches. The traditional chip-level design tradeoff mix of circuit density, wirability, performance, and power becomes more complex through the mix of applications across the SoC.

Today's dominant approach to testing a manufactured ASIC uses full-scan design for test (DFT) structures and automatic test pattern generation [15]. This approach derives much of its benefit in productivity, test quality, and ability to diagnose failures, from consistent and predictable DFT structures across the entire chip. A functional/technology "island" often brings with it a unique DFT design and pattern application requirement that differs from the overall ASIC. Such is the case when integrating analog IP [14] and embedded FPGAs, for example. Creating test data for diverse components may require unique test data development or characterization, and the resulting data stream must be integrated into the data stream of the overall chip, which may continue to be based at least in part on traditional full-scan methods.

7. CONCLUSION

Designing hundred million gate chips, made possible by nanometer silicon technologies, has first presented the challenge of managing massive design size and complexity, while product performance requirements continue to increase and time-to-market requirements continue to shrink. Design productivity gains and design schedule reductions are being realized through design process improvement and tool integration. A comprehensive strategy for design closure of large flat designs, hierarchical designs, and combinations of both, can provide the path to the earliest design closure solution. First-Time-Right design provides the greatest benefit to time-to-market, and thus design quality methods continue to rise in importance with increased design content and complexity.

Silicon density at the 90 nm level has increased the need to manage active and static power within the design, at the circuit, logic, and architecture levels. Signal integrity issues, and their avoidance through design techniques, were presented including coupled noise, IR drop, and reliability wear-out mechanisms.

Leveraging massive design capabilities in a single SoC leads to the integration of diverse functional components. Functional organization and chip organization must be combined into a single design planning solution. The functional components comprising the SoC can be diverse in terms of optimum library and technology, operating point, implementation platform, and test methodology. This has led to design integration methodologies including Voltage Islands, mixed library/placement terrains, and embedded FPGA's.

8. REFERENCES

[1] D.E. Lackey, "Applying Placement-based Synthesis for On-Time System-on-a-Chip Design," *Proceedings of the IEEE Custom Integrated Circuits Conference*," 2000, pp.121-124.

[2] J. Koehl, D.E. Lackey, G.W. Doerre, "IBM's 50 Million Gate ASICs," *Proceedings of the IEEE ASP Design Automation Conference*, 2003, pp. 628-634.

[3] L. Marshall, T.Wagner, J.Koehl: "A New ASIC Timing Signoff Methodology", *IBM MicroNews*, 2/2002.

[4] W. Donath, P. Kudva, L. Stok, P. Villarrubia, L. Reddy, A. Sullivan, "Transformational Placement and Synthesis," *Proceedings of the Conference and Exhibition on Design, Automation, and Test in Europe*, 2000, pp.194-201.

[5] C.J.Alpert, G.Gandham, J.Hu, J.L. Neves, S.T. Quay, S.S. Sapatneker, "Steiner Tree Optimization for Buffers, Blockages, and Bays," *Proceedings of the IEEE International Symposium on Circuits and Systems*, 2001, pp.399-402.

[6] C. J. Alpert, G.J. Nam, and P. G. Villarrubia, "Free Space Management for Cut-Based Placement", *Proceedings of the IEEE ICCAD*, November, 2002, pp.746-751.

[7] J. Koehl, J. Schietke, "In-place Timing Optimization", *Proceedings of SAME, Sophia Antipolis*, 2000.

[8] J. Vygen, "Algorithms for Large-Scale Flat Placement", *Proceedings of the 34th Design Automation Conference,* 1997, pp.746-751.

[9] G.W. Doerre, D.E. Lackey, "The IBM ASIC/SoC methodology – A recipe for first-time success*," IBM Journal of Research and Development*, Vol.46, No.6, pp.649-660, Nov. 2002.

[10] U. Brenner, A. Rohe, "An Effective Congestion Driven Placement Framework", *ISPD 2002*, pp.6-11.

[11] J.A Darringer, *et al.*, "EDA in IBM: Past, Present, and Future," *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, Vol.19, No.12, pp. 1476-1496, December, 2000.

[12] J.A.G. Jess, K, Kalafala, S.R. Naidu, R.H.J.M. Otten, C.Visweswariah, "Statistical timing for parametric yield prediction of digital integrated circuits," *Proceedings of the 2003 Design Automation Conference.*

[13] D.E. Lackey, "Design Planning Methodology for Rapid Chip Deployment," *The Eighth IEEE/DATC Electronics Design Processes Workshop*, April, 2001.

[14] T.R Bednar, P.H. Buffet, R.J. Darden, S.W. Gould, P.S. Zuchowski*,* "Issues and strategies for the physical design of system-on-a-chip ASICs," *IBM Journal of Research and Development*, Vol.46, No.6, pp. 661-673, Nov. 2002.

[15] S.Oakland, J.Monzel, R.Bassett, P.Gillis, "An ASIC Foundry View of Design for Test," Proceedings of the IEEE International Test Conference, 1994.

[16] A.Kuelmann *et al.,* "Verity – a Formal Verification Program for custom CMOS circuits," *IBM Journal of Research and Development,* Vol.39 No.1/2, pp.149-165, Jan/March 1995.

[17] M.R. Becer, D.Blaauw, R. Panda, I.H. Hajj, "Early Probabalistic Noise Estimation for Capacitively Coupled Interconnects," IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems, Vol.22, No.3, March 2003.

[18] E. Enomoto, "Low Power Design Technology for Digital LSIs," *IEICE Transactions on Electronics,* Vol. E79-C, No.12, pp.1639-1649, December 1996.

[19] S. Borkar, "Design Challenges of Technology Scaling," IEEE Micro, Vol. 19, No. 4, pp.23-29, July-August 1999.

[20] A.O. Adan and K. Higashi, "OFF-State Leakage Current Mechanisms in BulkSi and SOI MOSFETs and Their Impact on CMOS ULSIs Standby Current," IEEE Transactions on Electron Devices, Vol. 48, No. 9, pp.2050-2057, Sept. 2001.

[21] E.J. Nowack, "Maintaining the Benefits of CMOS Scaling when Scaling Bogs Down," *IBM Journal of Research and Development*, No. 2/3 March/May 2002.

[22] A.Dean, G. Garrett, M. Stan, S. Ventrone, "Low Power Design for ASIC Cores," *VLSI Design*, 2000, pp.1-15.

[23] J.G. Xi, W. W-M. Dai, "Low Power Clock Distribution," *Low Power Design Methodologies*, Copyright 1996 by Kluwer Academic Publishers, Chapter 5.

[24] T.Stoehr *et al*. "Analysis, Reduction and Avoidance of Crosstalk on VLSI Chips," *Proceedings of the ISPD,* 1998.

[25] J.A Darringer, *et al.*, "Early analysis tools for System-on-a-chip design," *IBM Journal of Research and Development*, Vol.46, No.6, pp. 691-707, Nov. 2002.

[26] D.E. Lackey, P.S. Zuchowski, T.R. Bednar, D.W. Stout, S.W. Gould, J.W. Cohn, "Managing Power and Performance for System-on-Chip Designs using Voltage Islands," *Proc. IEEE Conference on Computer-Aided Design*, Nov. 2002.

[27] P. S. Zuchowski, C. B. Reynolds, R. J. Grupp, S. G. Davis, B. Cremen, B. Troxel, "A Hybid ASIC and FPGA Architecture", *Proc. IEEE Conference on Computer-Aided Design*, Nov. 2002.