Chapter 15
High-Level Synthesis Algorithms for Power and Temperature Minimization

Li Shang, Robert P. Dick, and Niraj K. Jha

Abstract Increasing digital system complexity and integration density motivate automation of the integrated circuit design process. High-level synthesis is a promising method of increasing designer productivity. Continued process scaling and increasing integration density result in increased power consumption, power density, and temperature. High-level synthesis for integrated circuit (IC) power and thermal optimization has been an active research area in the recent past. This chapter explains the challenges power and temperature optimization pose for high-level synthesis researchers and summarizes research progress to date.

Keywords: Behavioral synthesis, High-level synthesis, Power, Temperature, Thermal modeling, Reliability

15.1 Power and Temperature Optimization

In this section, we give an overview of the key motivations for, and challenges of, optimizing power consumption and temperature during high-level synthesis.

15.1.1 Brief Introduction to High-Level Synthesis

High-level synthesis [1–4] is the process of automatically converting a behavioral, algorithmic, specification to an optimized register-transfer level digital design. The specification indicates the behavior of an algorithm and available hardware resources such as multipliers and multiplexers, but does not indicate the manner in which the algorithm should be implemented. A high-level synthesis algorithm automatically selects the set of hardware resources to use, determines the connections between them, binds operations to functional units such as multipliers, determines a clock frequency, and produces a schedule of operations. High-level synthesis can

P. Coussy and A. Morawiec (eds.) *High-Level Synthesis*.

therefore be formulated as an optimization problem with functionality constraints. Performance, power consumption, temperature, IC area, reliability, or other metrics may be optimized or constrained [5–15].

15.1.2 Importance of Power Consumption and Temperature

Power is the source of the greatest problems facing IC designers. High-power ICs rapidly deplete battery energy. Rapid changes in power consumption result in on-chip voltage fluctuations that lead to transient errors. High spatial and temporal power densities lead to high temperatures, which result in decreased lifetime reliability. High temperatures also increase leakage power consumption, thereby closing a self-reinforcing power–temperature feedback loop. The effects of increasing power consumption, power variation, and power density are expensive to handle. The wages of power are bulky short-lived batteries, huge heatsinks, large on-die capacitors, high server electric bills, and unreliable ICs. The only alternative is optimizing IC power consumption, temperature, and reliability. Power optimization within high-level synthesis has a long history, which we will review in this chapter. In contrast, temperature optimization during high-level synthesis began to receive widespread attention fairly recently, although some researchers foresaw the coming importance of the problem a decade ago.

Temperature is increased by both IC dynamic and leakage power. In addition, IC on-die temperature profiles depend on the temporal and spatial distribution of IC power as well as the packaging and cooling solution. Increasing IC power consumption increases IC peak temperature as well as on-die spatial and temporal thermal variation, which have significant impact on IC power consumption, temperature, reliability, cooling cost, and performance. A high IC temperature increases charge carrier concentrations, resulting in increased subthreshold leakage power consumption. In addition, it decreases charge carrier mobility, decreasing transistor and interconnect performance, and decreases threshold voltage, increasing transistor performance. Moreover, temperature heavily influences the fault processes, i.e., electromigration, dielectric breakdown, and power–thermal cycling, that lead to a large number of IC permanent faults. Finally, increasing IC power density requires the use of more effective cooling and packaging solutions to ensure IC reliable runtime operation, resulting in a significant increase in cooling and packaging cost. In summary, thermal issues have become a major concern in IC design. Modeling and optimizing IC thermal properties is thus essential for reliability, power consumption, and performance.

15.1.3 Power Analysis and Optimization

IC power analysis and optimization have been an active research areas for decades. Researchers developed power modeling techniques at all levels of the IC design

hierarchy. High-level synthesis poses unique challenges for IC power modeling and analysis. During behavioral synthesis, the lack of low-level implementation details, such as interconnect length and timing information permitting estimation of transient glitches, makes accurate power analysis challenging. In addition, power optimization during high-level synthesis typically involves the evaluation of numerous optimization decisions, requiring highly-efficient power analysis techniques. Most existing power-aware high-level synthesis systems use microarchitectural or structural power modeling methods to permit fast power estimation. These modeling methods are capable of approximately estimating the relative power savings of behavioral optimization decisions, but unable to characterize the accurate IC power profile.

Power optimization has been a primary focus of high-level synthesis for more than a decade. A variety of power optimization techniques have been proposed to tackle IC dynamic and leakage power consumption during high-level synthesis. IC dynamic power consumption can be reduced by attacking supply voltage, capacitance, switching activity, and frequency. Among these, voltage scaling is the most promising technique for reducing IC dynamic power consumption, due to the fact that IC dynamic power is quadratically proportional to supply voltage. Techniques, such as voltage and frequency scaling, multi-V_{dd}, and voltage islands, have been widely adopted by recently-developed low-power high-level synthesis systems. However, voltage reduction has a negative impact on circuit performance. Moreover, the effectiveness of voltage scaling diminishes as the supply voltage of nanometer-scale ICs approaches the sub-volt range. IC leakage power consumption was once a second-order consideration. However, it is becoming increasingly significant as a result of continued IC process scaling. Leakage accounts for 40% of the power consumption of today's high-performance microprocessors [16]. Leakage power can be the primary limitation on the lifetime of battery-powered systems. Leakage power optimization techniques, such as body biasing and transistor sizing, have been used in several high-level synthesis systems [17–20]. IC subthreshold leakage increases superlinearly with temperature. Due to the increase of IC power density and thermal effects, thermal-aware leakage analysis has gained prominence in high-level synthesis [21, 22].

15.1.4 Thermal Analysis and Optimization

An IC's thermal profile is a complex, time-varying function of its power consumption profile. The chip average temperature is determined by IC average power density and cooling package efficiency. The run-time chip thermal profile, on the other hand, depends on IC spatial and temporal power variation. The occurrence of on-die hotspots is often the result of transient activation of functional units with a high power density.

Behavioral design changes alone cannot effectively solve the IC temperature optimization problem. IC thermal analysis requires detailed physical information,

i.e., IC floorplan, interconnect, and chip-package configuration. IC thermal optimization requires the use of behavioral power optimization techniques to minimize IC average power density and temperature-aware physical design to balance and optimize the chip thermal profile. A unified high-level and physical analysis and optimization flow is critical for IC thermal optimization.

One primary challenge of IC thermal optimization comes from the high computational complexity of IC thermal analysis. IC thermal analysis is the process of characterizing the three-dimensional temperature profile of IC chip and cooling package. It requires a detailed simulation of heat conduction from an IC's power sources, i.e., transistors and interconnects, through cooling package layers, to the ambient environment, which can be described using the following equation:

$$\rho c \frac{\partial T(\mathbf{r},t)}{\partial t} = \bigtriangledown \cdot (k(\mathbf{r}) \bigtriangledown T(\mathbf{r},t)) + p(\mathbf{r},t), \tag{15.1}$$

where ρ is the material density, c is the mass heat capacity, $T(\mathbf{r},t)$ and $k(\mathbf{r})$ are the temperature and thermal conductivity of the material at position $\mathbf{r}$ and time t, and $p(\mathbf{r},t)$ is the power density of the heat source. Steady-state thermal analysis characterizes the chip temperature distribution when the IC power consumption does not vary with time, i.e., when the heat capacity, c, is neglected. Dynamic thermal analysis is used to characterize the temporal variations of the IC thermal profile. This problem is analogous to transient analysis of an electrical circuit [23], with electrical resistance and capacitance replaced with thermal resistance and heat capacity. The rate of temperature change in response to a change in power density is related to the thermal RC time constant of the IC region of interest. The major challenges of numerical IC thermal analysis are high computational complexity and memory usage. For steady-state thermal analysis, high modeling accuracy requires fine-grain modeling of IC chip and cooling package, resulting in high memory usage and long analysis time. For dynamic thermal analysis using time-domain methods, such as the fourth-order Runge-Kutta method, higher modeling accuracy requires fine spatial and temporal discretization granularity, increasing computational overhead and memory usage. Recent IC thermal analysis techniques use spatially and temporally adaptive numerical modeling methods to control the computational complexity and memory usage of IC thermal analysis while maintaining high accuracy [24].

15.2 High-Level Synthesis Algorithms for Power Optimization

Research on power-aware high-level synthesis can be traced back to the early 1990s. This section reviews existing low-power high-level design methodologies and synthesis tools.

15.2.1 Dynamic Power Optimization in High-Level Synthesis

In the past, IC power consumption was dominated by dynamic power. Therefore, early research on low-power synthesis focused on dynamic power optimization. IC dynamic power consumption is a quadratic function of supply voltage. Voltage scaling is therefore the most effective dynamic power optimization technique. However, voltage scaling may have a negative impact on circuit performance. Therefore, the tradeoff between power and performance has been a central theme in power-aware high-level synthesis. Johnson and Roy developed MESVS, a behavioral scheduling algorithm, that minimizes IC power consumption by using multiple supply voltages [25]. This work uses integer linear programming to produce an optimal schedule with discrete voltage-level assignment under timing constraints. Unfortunately, optimal integer linear programming formulations generally cannot be used for large problem instances due to high computational complexity. Raje and Sarrafzadeh proposed a heuristic to solve the voltage assignment problem [26]. The computational complexity of this method is $\mathcal{O}(N^2)$. Chang and Pedram developed a dynamic programming technique to solve the multi-voltage scheduling problem [27]. This technique reduces supply voltages along non-critical paths to optimize IC power consumption and minimize performance impact. Hong et al. designed a multi-voltage scheduling algorithm to minimize the power consumption of core-based systems-on-a-chip [28]. Helms et al. propose a behavioral synthesis system which uses multi-voltage assignment and adaptive body biasing to minimize IC power consumption [29]. These studies demonstrate that voltage scaling can reduce IC power consumption. However, the extra power saving decreases with the number of voltage levels. Recently, Liu et al. propose an approximation algorithm for IC power optimization using multiple supply voltages [30]. The computational complexity of the proposed approximation algorithm is $\mathcal{O}(dkN)$, where d and k are small constants. This work shows significant runtime advantage over the past work.

IC dynamic power consumption can be reduced by minimizing circuit capacitance and run-time switching activity. Chatterjee and Roy designed a behavioral synthesis system, which uses architectural transformation to minimize circuit switching activity [31]. Raghunathan and Jha developed the first optimal, ILP-based formulation of high-level synthesis for switching power minimization [32]. Chandrakasan et al. developed HYPER-LP, a high-level synthesis system using algorithmic transformation to reduce circuit capacitance, thereby reducing IC power consumption [9]. Chang and Pedram developed an low-power allocation and resource binding technique to minimize the switching activity in registers [11] and datapath functional components [33]. In this work, the power-optimal register and functional component assignment problem is formulated as a max-cost flow problem. Dasgupta and Karri developed binding and scheduling techniques to minimize the switching activity of buses [6]. Musoll and Cortadella developed a high-level synthesis system, which uses loop interchange, operand reordering, operand sharing, idle units, and operand correlation, for reducing the activities of IC functional units [34]. Raghunathan and Jha designed SCALP, an iterative-improvement-based high-level synthesis system [13], which integrates a variety

of power optimization techniques, including architectural transformation, scheduling, clock selection, module selection, and hardware allocation and assignment. Lakshminarayana et al. proposed a power-aware register binding technique for high-level synthesis, which provides the first formulation of a perfect power management philosophy, i.e., no functional unit that does not need to be active in a given cycle should consume any switching power in that cycle [35]. Dasgupta and Karri developed a high-level synthesis system for IC energy and reliability optimization [36]. They proposed a resource binding and scheduling algorithm to minimize circuit switching activity, thereby optimizing IC power consumption and minimizing electromigration-induced failure effects in on-chip buses. Ercegovac et al. proposed a behavioral synthesis system [37] that uses multi-gradient search for system resource allocation using multiple-precision arithmetic units. Karmarkar-Karp's number partitioning heuristic is used to determine task assignment. Lakshminarayana et al. proposed a high-level power optimization technique which extracts common-case behavior from the given behavioral description and then synthesizes an RTL implementation of the common-case circuit, which is a much smaller than the circuit that implements the complete behavior and runs most of the time [38]. Wang et al. proposed a high-level design methodology for IC energy and performance optimization [39] called input space adaptive design. This technique identifies the behavioral equivalence among sub-circuits and eliminates redundant logical operations, thereby optimizing IC energy and performance.

15.2.2 Leakage Power Optimization in High-Level Synthesis

IC leakage power consumption is becoming increasingly significant as a result of technology scaling. Therefore, leakage power optimization during high-level synthesis has drawn significant attention. Khouri and Jha [17] developed a behavioral, iterative algorithm to minimize IC leakage power consumption using dual-V_{th} technology. The proposed algorithm is a greedy approach that iteratively identifies the operation with the maximum leakage power reduction potential and binds it with a high-V_{th} implementation. Gopalakrishnan and Katkoori developed a leakage-aware resource allocation and binding algorithm using multi-V_{th} technology [18]. This algorithm seeks to maximize the idle time slots of datapath components. Idle functional modules are scheduled to enter the sleep mode at runtime to minimize the IC leakage power consumption. Tang et al. formulated the leakage optimization problem as the maximum weight independent set problem [19]. A heuristic was proposed to identify the datapath components with maximum or near-maximum leakage reduction potentials, which are then replaced with low-leakage alternatives. Dal et al. developed a low-power high-level synthesis algorithm using power islands [20]. The supply voltage of each power island can be controlled independently. The proposed algorithm conducts circuit partitioning and assigns circuit components with overlapping idle times to the same power island. Idle power islands are then scheduled to be power-gated to minimize leakage power consumption. IC sub-threshold leakage power is a strong function of chip temperature. Therefore,

thermal effects must be considered during leakage power optimization. We will later survey thermal-aware leakage optimization techniques.

15.2.3 Importance of Incorporating Physical Design Within High-Level Synthesis

It is becoming increasingly important to consider physical design decisions within high-level synthesis. Interconnect power consumption and delay are increasing relative to logic delay. Increasing power densities are making it necessary to determine and optimize the IC thermal profile at design time; computing a thermal profile requires a power profile. Determining the interconnect structure and power profile depends on the knowledge of the IC floorplan. As a result, a number of researchers have considered the impact of physical details, e.g., floorplanning information, on high-level synthesis [40–46].

Taking interconnect power consumption and delay into consideration during high-level synthesis has attracted significant attention. In previous work [47–51], the number of interconnects or multiplexers was used to estimate the interconnect cost. The performance and power impact of the interconnect and interconnect buffers are now first-order considerations [52]. It is no longer possible to accurately predict the power consumption and performance of a design without first knowing enough about its floorplan to predict the structure of its interconnect. This change has complicated both design and synthesis. For this reason, a number of researchers have worked on interconnect-aware high-level synthesis algorithms [53–55]. These approaches typically use a loosely coupled independent floorplanner for physical estimation. This technique has the advantage of allowing estimation of physical properties but has a drawback. Creating a floorplan from scratch for each high-level synthesis move is inefficient, given the fact that the new floorplan frequently has only small differences with the previous one. The constructive approach works for small problem instances but is unlikely to scale to large designs. New techniques for tightly coupling behavioral and physical synthesis that dramatically improve their combined performance and quality are now necessary.

Incremental automated design promises to build tighter relationship between high-level synthesis and physical design, improving the quality of each [56, 57]. A number of high-level synthesis algorithms are based on incremental optimization and are therefore amenable to integration with incremental physical design algorithms. This has the potential of improving both quality and performance. Incremental methods improve quality of results by maintaining important properties across consecutive physical estimations during synthesis. Moreover, they shorten CPU time by reusing and building upon previous high-quality physical design solutions that required a huge amount of effort to produce. Recent work has proposed unified incremental behavioral synthesis and floorplanning to permit more accurate communication delay, communication power consumption, and power profile estimation [58].

15.3 Modeling and Optimizing Temperature in High-Level Synthesis

This section introduces the main challenges of temperature-aware high-level synthesis and describes a number of recent techniques to overcome them.

15.3.1 Thermal Model Selection for Use in High-Level Synthesis

It is important to select appropriate thermal modeling and analysis techniques for use in temperature-aware high-level synthesis. In reality, ICs experience temporal and spatial temperature variation. However, accurately modeling spatial and temporal variation during thermal analysis can be the most time consuming part of high-level synthesis. Given a fixed amount of time for synthesis, there is a trade-off between the amount of time spent on thermal analysis and the number of tentative behavioral synthesis solutions that can be considered. Therefore, it is important to model temporal and spatial temperature variation with as much detail as necessary for accuracy, but no more.

A number of high-level synthesis formulations consider energy consumption or average power consumption. This is equivalent to optimizing temperature while neglecting temporal and spatial variation in temperature. In some applications, this is legitimate. In others, it can result in extremely large errors. Let us now consider the circumstances in which it is necessary to model spatial and temporal variation in temperature.

IC packaging has a strong influence on heat flow, and therefore on the importance of modeling spatial temperature variation. Packaging and cooling solutions that more efficiently remove heat tend to be more expensive. In order to minimize cost, it is reasonable to select a cooling solution that permits the temperature to approach its constraint under worst-case or average-case conditions. As a result, in low power density designs the package will have poor thermal conductance, e.g., a plastic package without heatsink. Is this case, the conductance between different points on the silicon die is high relative to the conductance between a point on the die, through the package, to the ambient. As a result, the temperature of the active layer will generally be fairly uniform despite spatial variation in power density. For this reason, a simple thermal model is sufficient for low power density ICs using low thermal conductance packages and cooling solutions [59,60]. High power density designs require more efficient packaging and cooling solutions to maintain safe temperatures. As a result, the thermal conductance between different points on the silicon die can decrease relative to the thermal conductance to the ambient. In this case, spatial variation in the power profile will result in spatial variation in temperature.

The properties of temporal variation in IC power consumption have a strong influence on the thermal modeling requirements. Most existing work on

temperature-aware high-level synthesis assumes that power density does not vary with time and uses steady-state thermal analysis based on the temporal averages of power density. This is legitimate when the temporal variation of power densities occurs in a much shorter timescale than the IC thermal RC time constants, e.g., a high-frequency periodic system in which power density does not change on long time scales due to changing input data. However, it is not legitimate when there are long time scale changes in power density. If the interval of change in power density is long relative to the thermal RC time constants, it may be possible to accurately approximate the temperature by conducting steady-state analysis for each power density phase. However, in general, accurately modeling the thermal impact of time-varying power profiles requires dynamic thermal analysis, which is generally much more time-consuming than steady-state analysis.

Thus far, we have considered the conditions in which spatial and temporal thermal variation can be entirely neglected. However, once the decision is made to model spatial and/or temporal variation, it is still necessary to determine the required modeling resolution. Increasing the number of thermal elements or temperature evaluation time instants can dramatically increase the run-time of thermal analysis.

The required thermal model spatial resolution depends on material properties, cooling environment, and power density variation. During thermal analysis, it is common for an IC to be partitioned into multiple elements, each of which is assumed to be isothermal, i.e., to have internally-uniform temperature. To minimize analysis time, thermal elements should generally be as large as possible while still honoring the isothermal assumption. Note that an element with uniform power density does not necessarily honor the isothermal assumption because its neighboring thermal elements may have different temperatures, resulting in a substantial temperature gradient. The architectural thermal analysis tools commonly used in high-level synthesis thermal analysis support manual [61] or automatic [60] adaptation of spatial modeling granularity.

Dynamic thermal analysis is frequently formulated as a time-domain initial value problem in which the thermal profile is iteratively updated at increasing time instants. There is a tradeoff between the number of time instants, at which the temperature is explicitly evaluated, and accuracy. Assuming a constant error bound, the duration between explicit temperature evaluations depends on the rate and complexity of changes in the power profile. Therefore, dynamic adaptation is required to minimize analysis time under a constraint on maximum error. The thermal analysis tools commonly used in high-level synthesis support dynamic temporal adaptation to varying degrees [60, 61].

15.3.2 High-Level Synthesis Algorithms for Temperature Optimization

Temperature-aware high-level synthesis is currently a thriving research area, with new work appearing monthly in top conferences and journals. Ten years ago, Weng

and Parker were the first to address the problem by moving high power density functional units away from high-temperature areas to reduce the spatial power density and introducing redundant operators to reduce the temporal power density [62]. It is interesting to note that Prakash and Parker were also the first to formulate the system-level heterogeneous distributed system synthesis problem, also 10 years before it became a highly-active research area [63]. Mukherjee et al. proposed to incrementally improve binding decisions to reduce the temperature of the hottest functional unit, thereby reducing both dynamic and leakage power consumption [21]. Gu et al. designed TAPHS, a temperature-aware unified physical and behavioral synthesis system [64]. TAPHS integrates behavioral and physical thermal optimization techniques, including voltage assignment, voltage island generation, and floorplanning, to optimize chip temperature, power, performance and area. Lim and Kim propose a network flow based method for temperature-aware binding that minimizes both peak and average switched capacitance [65]. Ni and Öğrenci Memik proposed a technique to reduce leakage power consumption using selective resource redundancy [22].

15.4 Conclusions

This chapter has described the current state-of-the-art in high-level synthesis algorithms that optimize power consumption and temperature. International Technology Roadmap for Semiconductors imply that power consumption will continue to be a primary concern for IC designers. Emerging power and power-induced problems, such as process variation influenced IC leakage power consumption, IC leakage-thermal coupling, and power–thermal dependent IC lifetime reliability problems, further exacerbate the challenges for high-level synthesis algorithms. On the other hand, power optimization techniques that were widely used in the past, such as voltage scaling and body biasing, will soon start running out of steam as a result of continued process scaling. Moreover, as power-aware unified architectural–physical design flows cease to be luxuries and become necessities, it will become necessary to cooperatively solve many problems that were once orthogonal to high-level synthesis. These challenges may require fundamental changes to existing high-level synthesis flows.

References

1. R. Camposano and W. Wolf, *High Level VLSI Synthesis*. Kluwer, MA, 1991
2. D. C. Ku and G. D. Micheli, *High Level Synthesis of ASICs Under Timing and Synchronization Constraints*. Kluwer, MA, 1992
3. D. Gajski, N. Dutt, A. Wu and S. Lin, *High-Level Synthesis: Introduction to Chip and System Design*. Kluwer, MA, 1992
4. A. Raghunathan, N. K. Jha and S. Dey, *High-Level Power Analysis and Optimization*. Kluwer, MA, 1998

5. R. Mehra and J. Rabaey, "Behavioral level power estimation and exploration," in *Proc. Int. Wkshp Low Power Design*, Apr. 1994, pp. 197–202
6. A. Dasgupta and R. Karri, "Simultaneous scheduling and binding for power minimization during microarchitecture synthesis," in *Proc. Int. Symp. Low-Power Design*, Apr. 1994
7. L. Goodby, A. Orailoglu and P. M. Chau, "Microarchitecture synthesis of performance-constrained, low-power VLSI designs," in *Proc. Int. Conf. Computer Design*, Oct. 1994
8. A. Raghunathan and N. K. Jha, "Behavioral synthesis for low power," in *Proc. Int. Conf. Computer Design*, Oct. 1994, pp. 318–322
9. A. P. Chandrakasan, M. Potkonjak, R. Mehra, J. Rabaey and R. Brodersen, "Optimizing power using transformations," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 14, no. 1, pp. 12–31, 1995
10. R. S. Martin and J. P. Knight, "Power profiler: Optimizing ASICs power consumption at the behavioral level," in *Proc. Design Automation Conf.*, June 1995
11. J. M. Chang and M. Pedram, "Register allocation and binding for low power," in *Proc. Design Automation Conf.*, June 1995
12. N. Kumar, S. Katkoori, L. Rader and R. Vemuri, "Profile-driven behavioral synthesis for low-power VLSI systems," *IEEE Design Test*, vol. 12, no. 3, pp. 70–84, 1995
13. A. Raghunathan and N. K. Jha, "SCALP: An iterative-improvement-based low-power data path synthesis system," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 16, no. 11, pp. 1260–1277, 1997
14. K. S. Khouri, G. Lakshminarayana and N. K. Jha, "High-level synthesis of low power control-flow intensive circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 18, no. 12, pp. 1715–1729, 1999
15. H. P. Peixoto and M. F. Jacome, "A new technique for estimating lower bounds on latency for high level synthesis," in *Proc. Great Lakes Symp. VLSI*, Mar. 2000, pp. 129–132
16. S. Naffziger et al., "The implementation of a 2-core, multi-threaded Itanium family processor," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 1, pp. 197–209, 2006
17. K. S. Khouri and N. K. Jha, "Leakage power analysis and reduction during behavioral synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 10, no. 6, pp. 876–885, 2002
18. C. Gopalakrishnan and S. Katkoori, "KnapBind: An area-efficient binding algorithm for low-leakage datapaths," in *Proc. Int. Conf. Computer Design*, Oct. 2003, pp. 430–435
19. X. Tang, H. Zhou and P. Banerjee, "Leakage power optimization with dual-Vth library in high-level synthesis," in *Proc. Design Automation Conf.*, June 2005, pp. 202–207
20. D. Dal, A. Nunez and N. Mansouri, "Power islands: A high-level technique for counteracting leakage in deep sub-micron," in *Proc. Int. Symp. Quality of Electronic Design*, Mar. 2006, pp. 165–170
21. R. Mukherjee, S. Öğrenci Memik and G. Memik, "Temperature-aware resource allocation and binding in high-level synthesis," in *Proc. Design Automation Conf.*, June 2005
22. M. Ni and S. Öğrenci Memik, "Thermal-induced leakage power optimization by redundant resource allocation," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2006, pp. 297–302
23. G. S. Ohm, "The Galvanic circuit investigated mathematically," in *Die galvanische Kette: mathematisch bearbeitet*, 1827
24. Y. Yang, C. Zhu, Z. P. Gu, L. Shang and R. P. Dick, "Adaptive multi-domain thermal modeling and analysis for integrated circuit synthesis and design," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2006, pp. 575–582
25. M. Johnson and R. K. Roy, "Optimal selection of supply voltages and level conversion during datapath scheduling under resource constraints," in *Proc. Int. Conf. Computer Design*, Oct. 1996, pp. 72–77
26. S. Raje and M. Sarrafzadeh, "Variable voltage scheduling," in *Proc. Int. Symp. Low Power Electronics & Design*, Aug. 1995, pp. 9–14
27. J. Chang and M. Pedram, "Energy minimization using multiple supply voltages," in *Proc. Int. Symp. Low Power Electronics & Design*, Aug. 1996, pp. 157–162

28. I. Hong, D. Kirovski, G. Qu, M. Potkonjak and M. B. Srivastava, "Power optimization of variable voltage core-based systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 18, no. 12, pp. 1702–1714, 1999
29. D. Helms, O. Meyer, M. Hoyer and W. Nebel, "Voltage- and ABB-island optimization in high level synthesis," in *Proc. Int. Symp. Low Power Electronics & Design*, Aug. 2007, pp. 153–158
30. H. Liu, W. Lee and Y. Chang, "A provably good approximation algorithm for power optimization using multiple supply voltages," in *Proc. Design Automation Conf.*, June 2007, pp. 887–890
31. A. Chatterjee and R. K. Roy, "Synthesis of low power linear DSP circuits using activity metrics," in *Proc. Int. Conf. VLSI Design*, Jan. 1994, pp. 261–264
32. A. Raghunathan and N. K. Jha, "An ILP formulation for low power based on minimizing switched capacitance during datapath allocation," in *Proc. Int. Symp. Circuits & Systems*, May 1995, pp. 1069–1073
33. J. Chang and M. Pedram, "Module assignment for low power," in *Proc. European Design Automation Conf.*, Sept. 1996, pp. 376–381
34. E. Musoll and J. Cortadella, "High-level synthesis techniques for reducing the activity of functional units," in *Proc. Int. Symp. Low Power Electronics & Design*, Aug. 1995, pp. 99–104
35. G. Lakshminarayana, A. Raghunathan, N. K. Jha and S. Dey, "A power management methodology for high-level synthesis," in *Proc. Int. Conf. VLSI Design*, Jan. 1998
36. A. Dasgupta and R. Karri, "High-reliability, low-energy microrchitecture synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 17, pp. 1273–1280, 1998
37. M. Ercegovac, D. Kirovski and M. Potkonjak, "Low-power behavioral synthesis optimization using multiple precision arithmetic," in *Proc. Design Automation Conf.*, June 1999, pp. 568–573
38. G. Lakshminarayana, A. Raghunathan, K. S. Khouri, N. K. Jha and S. Dey, "Common case computation: A high-level power-optimizing technique," in *Proc. Design Automation Conf.*, June 1999
39. W. Wang, A. Raghunathan, G. Lakshminarayana and N. K. Jha, "Input space adaptive design: A high-level methodology for energy and performance optimization," in *Proc. Design Automation Conf.*, June 2001, pp. 738–743
40. M. C. McFarland and T. J. Kowalski, "Incorporating bottom-up design into hardware synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 9, no. 9, pp. 938–950, 1990
41. J.-P. Weng and A. C. Parker, "3D scheduling: High-level synthesis with floorplanning," in *Proc. Design Automation Conf.*, June 1991, pp. 668–673
42. D. W. Knapp, "Fasolt: A program for feedback-driven data-path optimization," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 11, no. 6, pp. 677–695, 1992
43. J. P. Weng and A. C. Parker, "3D scheduling: High-level synthesis with floorplanning," in *Proc. Design Automation Conf.*, June 1992
44. Y. M. Fang and D. F. Wong, "Simultaneous functional-unit binding and floorplanning," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 1994
45. M. Xu and F. J. Kurdahi, "Layout-driven RTL binding techniques for high-level synthesis using accurate estimators," *ACM Transactions on Design Automation Electronic Systems*, vol. 2, no. 4, pp. 312–343, 1997
46. W. E. Dougherty and D. E. Thomas, "Unifying behavioral synthesis and physical design," in *Proc. Design Automation Conf.*, June 2000
47. P. G. Paulin and J. P. Knight, "Scheduling and binding algorithms for high-level synthesis," in *Proc. Design Automation Conf.*, June 1989, pp. 1–6
48. C. A. Papachristou and H. Konuk, "A linear program driven scheduling and allocation method followed by an interconnect optimization algorithm," in *Proc. Design Automation Conf.*, June 1990
49. T. A. Ly, W. L. Elwood and E. F. Girczyc, "A generalized interconnect model for data path synthesis," in *Proc. Design Automation Conf.*, June 1990

50. S. Tarafdar and M. Leeser, "The DT-model: High-level synthesis using data transfer," in *Proc. Design Automation Conf.*, June 1998
51. C. Jego, E. Casseau and E. Martin, "Interconnect cost control during high-level synthesis," in *Proc. Design Circuits & Integration Systems Conf.*, Nov. 2000
52. R. Ho, K. Mai and M. Horowitz, "The future of wires," *Proceedings of the IEEE*, vol. 89, no. 4, pp. 490–504, 2001
53. P. Prabhakaran and P. Banerjee, "Simultaneous scheduling, binding and floorplanning high-level synthesis," in *Proc. Int. Conf. VLSI Design*, Jan. 1998
54. L. Zhong and N. K. Jha, "Interconnect-aware low power high-level synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 3, pp. 336–351, 2005
55. A. Stammermann, D. Helms, M. Schulte, A. Schulz and W. Nebel, "Binding, allocation and floorplanning in low power high-level synthesis," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2003
56. O. Coudert, J. Cong, S. Malik and M. Sarrafzadeh, "Incremental CAD," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2000, pp. 236–244
57. W. Choi and K. Bazargan, "Hierarchical global floorplacement using simulated annealing and network flow migration," in *Proc. Design, Automation & Test in Europe Conf.*, Mar. 2003
58. Z. P. Gu, J. Wang, R. P. Dick and H. Zhou, "Unified incremental physical-level and high-level synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2007
59. G. Paci, P. Marchal, F. Poletti and L. Benini, "Exploring "temperature-aware design" in low-power MPSoCs," in *Proc. Design, Automation & Test in Europe Conf.*, Mar. 2006
60. Y. Yang, Z. P. Gu, C. Zhu, R. P. Dick and L. Shang, "ISAC: Integrated Space and Time Adaptive Chip-Package Thermal Analysis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2007
61. W. Huang et al., "HotSpot: A compact thermal modeling methodology for early-stage VLSI design," *IEEE Transactions on VLSI Systems*, vol. 14, no. 5, pp. 501–524, 2006
62. J.-P. Weng and A. C. Parker, "Taking thermal considerations into account during high-level synthesis," *VLSI Design*, vol. 5, no. 2, pp. 183–193, 1997
63. S. Prakash and A. Parker, "Synthesis of application-specific multiprocessor architectures," in *Proc. Design Automation Conf.*, June 1991
64. Z. P. Gu, Y. Yang, J. Wang, R. P. Dick and L. Shang, "TAPHS: Thermal-aware unified physical-level and high-level synthesis," in *Proc. Asia & South Pacific Design Automation Conf.*, Jan. 2006, pp. 879–885
65. P. Lim and T. Kim, "Thermal-aware high-level synthesis based on network flow method," in *Proc. Int. Conf. Hardware/Software Codesign and System Synthesis*, Oct. 2006