Multi-Objective Module Placement For 3-D System-On-Package

Eric Wong, Jacob Minz, and Sung Kyu Lim

***Abstract*—System-on-package (SOP) is a viable alternative to system-on-chip (SOC) for meeting the rigorous requirements of today's mixed-signal system integration. Thermal integrity is arguably the most crucial issue in three-dimensional (3-D) SOP due to the compact nature of the 3-D integration. In addition, the power supply noise issue becomes more serious as the supply voltage continues to decrease while the number of active devices consuming power increases. We propose a 3-D module and decap (decoupling capacitance) placement algorithm that evenly distributes the thermal profile and reduces the power supply noise. In addition, we allocate white spaces around the modules that require decaps to suppress the power supply noise while minimizing the area overhead. In our experimentation, we achieve improvements in both maximum temperature and decap amount with only small increase in area, wirelength, and runtime.**

***Index Terms*—Floorplanning, thermal via, three-dimensional (3-D) ICs.**

I. Introduction

The semiconductor industry is beginning to question the viability of the system-on-chip (SOC) approach due to its disadvantages of low yield and high cost. Recently, three-dimensional (3-D) packaging via system-on-package (SOP) [1] has been proposed as an alternative solution to meet the rigorous requirements of today's mixed-signal system integration. The SOP is about 3-D integration of multiple functions in a miniaturized package achieved by thin-film embedding. The 3-D SOP concept optimizes ICs for transistors and the package for integration of digital, RF, optical, sensor, and others. It accomplishes this by both build-up SOP, which is similar to IC fabrication, and by stacked SOP, which is similar to parallel board fabrication. The uniqueness of 3-D SOP is in the highly integrated or embedded RF, optical, or digital functional blocks and sensors, in contrast to stacked ICs and stacked package.

Thermal issues can no longer be ignored in high-performance 3-D packages due to higher power densities and other issues. High temperatures not only require more advanced heat sinks, but they also degrade circuit performance. Interconnect delay increases with temperature, which degrades circuit timing. If timing deteriorates enough, logic faults can occur. Hence, thermal issues must be considered early on in the design process. The continuing trend of reducing power supply voltage has resulted in a reduced noise margin, which effects reliability and may even cause functional failures due to spurious transitions. Active devices in 3-D packaging draw a large volume of instantaneous current during switching, which causes simultaneous switching noise (SSN). Existing approaches consider the thermal and SSN issues as an afterthought, which may require expensive cooling and an excessive amount of decoupling capacitance (decap) to suppress SSN. In addition, many iterations are required between full-length thermal and SSN simulation and manual layout repair until convergence to a satisfactory result. Thus, our goal in this study is to achieve a simultaneous thermal and SSN-aware physical design for 3-D SOP. Our design automation tool aims at reducing the hot spots and the amount of decap required to suppress SSN in a 3-D SOP design without compromising traditional design metrics such as area and wirelength.

Manuscript received August 26, 2005; revised December 31, 2005. This work was supported in part by MARCO GSRC/C2S2.

The authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA (e-mail: ewong,@ece.gatech.edu; jrminz@ece.gatech.edu; limsk@ece.gatech.edu).

Digital Object Identifier 10.1109/TVLSI.2006.876111

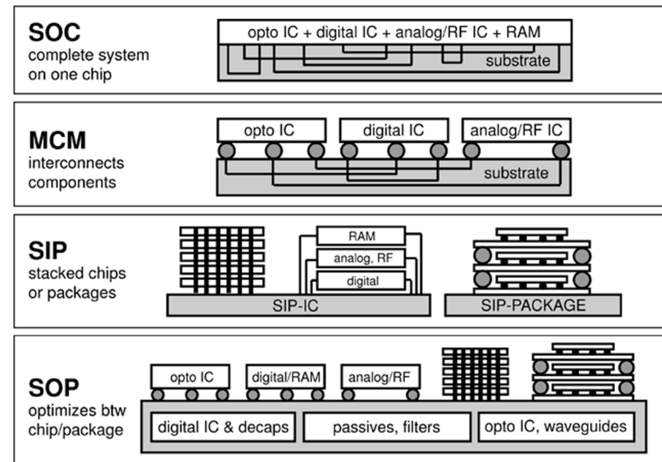

Fig. 1. Comparison among SOC, MCM, SIP, and SOP.

II. Why 3-D SOP?

A. SOC Versus SOP

The SOC paradigm is a new system integration approach, where not only more and more transistors but various mixed-signal active and passive components are also integrated into a single chip. However, the systems community is beginning to realize that SOC presents fundamental, engineering, and investment limits [1]. A complete integration of RF, digital, and optical technologies on a single chip poses numerous challenges. RF circuit performance is a tradeoff between the quality factor (Q) of passive components and power. Low-power circuit implementations for mobile applications require high-Q passive components. In standard silicon technologies, the Q-factor is limited to 5–25 due to the inherent losses of silicon [2]. Antennas are another example that cannot be integrated on silicon due to size restrictions [2]. On standard silicon, a major concern is substrate coupling caused by the finite resistivity of the silicon substrate. For multiple voltage levels, distributing power to the digital and RF circuits while simultaneously maintaining isolation and low electromagnetic interference (EMI) can be a major challenge [3].

This led to the 3-D system-in-package (SIP) approach, where one could stack multiple ICs or multiple package-stacked ICs at a much lower cost and size. The SIP, while providing major opportunities in both miniaturization and integration for advanced and portable electronic products, is still a subsystem, limited by the CMOS process just like the SOC. SIP could go one step further in embedding both active and passive components, but passive component embedding is bulky and thick film discrete components.

This distinguishes SIP from SOP [1] (see Fig. 1), which is a new emerging 3-D system integration concept that involves embedding of both active and passives, but the passives are by incorporation of ultra-thin films at micro-scale. With SOP, the package, not the board, is the system. Therefore, SOP is capable of addressing the shortcomings of both SOC and SIP as well as traditional packaging, which is bulky, costly, and lower in performance and reliability. This improvement comes in two ways: 1) SOP uses CMOS-based silicon for what it is good for, namely, transistor integration and 2) SOP uses the package for what it is good for, namely, RF, optical, and digital integration by means of IC-package-system co-design. Therefore, SOP overcomes both the

1063-8210/$20.00 © 2006 IEEE

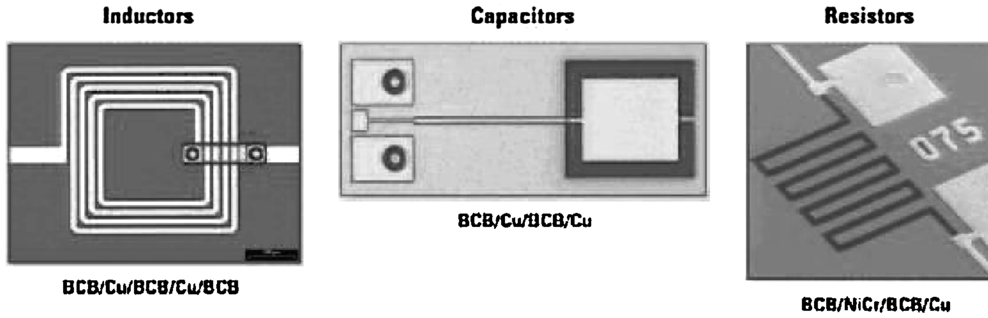

Fig. 2. Embedded passive components on a high-resistivity silicon substrate (courtesy of Semiconductor International).

computing limitations and integration limitations of SOC, SIP, multichip module (MCM), and traditional system packaging. It does this by having global wiring as well as RF, digital, and optical component integration in the package, and not in the chip. Moreover, 3-D SOP addresses the wire delay problem by enabling the replacement of long and slow global interconnects with short and fast vertical routes.

While 3-D SOP manufacturing technology continues to advance, the research in how to actually make use of the technology lags behind. Due to the high complexity in designing large-scale 3-D SOP under multiple objectives and constraints, computer-aided design (CAD) tools have become indispensable. Unlike the active CAD research effort in mixed-signal SOC [4], [5] and 3-D stacked ICs [6], [7], however, the CAD research for 3-D SOP considerably lags behind due to its short history.

III. Problem Formulation

The location of individual modules available from 3-D placement has a huge impact on many important metrics. First, different placements can impact the performance of a given 3-D SOP design as the delay of global interconnects between modules dominate the overall timing. Second, placement impacts the thermal and leakage profile. This is because the temperature of SOP modules is not only dependent on the heat generation rate of each individual module but also the heat coupling between neighboring modules. Moreover, the leakage power of each module is exponentially dependent on the temperature. Lastly, placement has a significant impact on SSN profile since the distance to the power pins determines the amount of IR drop experienced by each module.

The following are given as the input to our 3-D SOP placement problem: 1) a set of blocks that represent the various active and passive components in the given SOP design (see Fig. 2); 2) width, height, and maximum switching currents for each block; 3) a netlist that specifies how the blocks are connected via electrical wires; 4) the number of placement layers in the 3-D packaging structure; 5) the number of power/ground signal layers along with the location of the power/ground pins; and 6) tolerance on simultaneous switching noise. Our 3-D placement formulation is applicable to general SOP designs, where each SOP component is modeled as a rectangle with noise-related attributes.[1]

Let A^{tot} and W^{tot} denote the final footprint area and total wirelength of the 3-D placement. Let D^{tot} denote the total amount of decoupling capacitance required to suppress the SSN under the given tolerance value. Let T^{tot} denote the maximum temperature of the substrate. The goal of the *3D SOP Placement Problem* is to find the location of each block in the placement layers such that the following cost function is minimized while SSN constraint is satisfied: $w_1 \cdot A^{\text{tot}} + w_2 \cdot W^{\text{tot}} + w_3 \cdot D^{\text{tot}} + w_4 \cdot T^{\text{tot}}$. In addition, decaps are required to be placed adjacent to the blocks that require them.[2]

[1]In this paper, the SOP modules are two-dimensional (2-D) in a sense that each module only spans a single layer. However, it is possible to imagine 3-D modules in SOP that span multiple device layers.

IV. 3-D SOP Placement Algorithm

A. Overview of the Algorithm

Simulated annealing (SA) is a popular approach for module placement due to its high-quality solutions and flexibility in handling various constraints. We choose the sequence pair representation [8] and extend it to handle our 3-D floorplanning solutions. Specifically, we use k sequence pairs to represent the block placement of k device layers. Our SA procedure starts with an initial multilayer placement along with its cost in terms of area, wirelength, thermal, and decap. We perform a one-time setup of the thermal resistance matrix R, which is used during incremental temperature calculations to evaluate the thermal cost. In our perturbation scheme, we swap a random pair of blocks from the same or two different layers. The new area and wirelength are first computed from a given candidate solution. Next, our incremental thermal analyzer computes the temperature of the active and passive blocks from a given 3-D SOP placement. Lastly, the following steps are performed to measure the decap cost for a given 3-D floorplanning solution.

1) SSN noise analysis: The amount of SSN for each block is computed based on the location of the blocks and power pins.
2) Decap budget calculation: The amount of decap needed for each block based on its SSN is computed so that the overall SSN constraint is satisfied.
3) White space detection and insertion: We first detect and allocate existing white space (decap) to the blocks. In case the existing white space is not enough to suppress the SSN, more white space is added by expanding the area in the X and Y directions for decap implementation.
4) Decap allocation: White spaces (decaps) are allocated to the blocks that need them so that the utilization of white space is maximized.

Due to the runtime overhead involved, we use the first two steps to estimate the decap budget during the high-annealing-temperature region. During the low-annealing-temperature region, however, we perform all four steps to accurately compute the decap cost as well as area overhead. In this case, we compute the area, wirelength, and thermal cost *after* the decap insertion to reflect the impact of area expansion.

B. 3-D Power Supply Noise Modeling

We model the P/G network for 3-D SOP as a 3-D mesh. The edges in the mesh have inductive and resistive impedances. The mesh contains

[2]The density of interlayer via is very high in 3-D SOP technology, which explains why our formulation does not include via cost. Our simulated annealing-based placer, however, can be easily extended to consider via as objective or constraint.

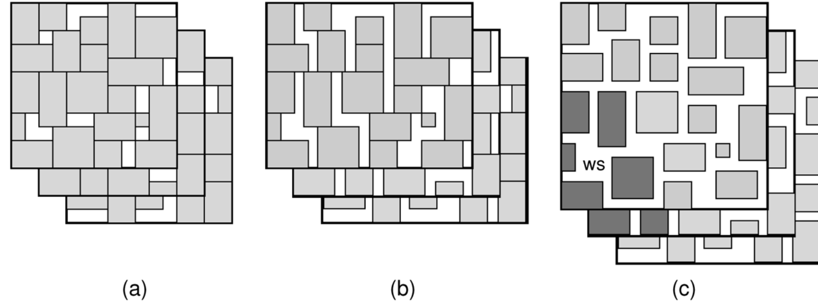

Fig. 3. Illustration of 3-D decap allocation. (a) 3-D placement. (b) X-expansion. (c) XY-expansion, where the darker blocks denote the neighboring blocks of the decap (= white space) inserted. Note that blocks from other layers can utilize the white space for decap insertion.

power-supply points and connection points. The connection points consume currents. The current is drawn from all the sources by the consumers, and the amount of current drawn along a path is inversely proportional to the impedance of the path in the power supply mesh.[3] The *dominant current source* for a block is defined as the voltage source supplying significantly more power to the block than any other neighboring sources. The *dominant path* for a block is the path from the dominant supply to the block causing the most drop in voltage. It has been shown experimentally in [9] that the shortest path between the dominant current source (nearest Vdd pins) and the block offers highly accurate SSN estimation within reasonable runtime. In our 3-D SSN analysis engine, we compute dominant paths (shortest paths to the nearest Vdd pins) for all blocks. This information is then dynamically updated whenever a new placement solution is evaluated in terms of SSN. Let P_k be a dominant current path for block k. Then, $T^k = \{P_j : P_j \cap P_k \neq \emptyset\}$ denotes the set of dominating paths overlapping with P_k (T^k includes P_k itself). Let P_{jk} be the overlapping segments between path P_j and P_k. Let $R_{P_{jk}}$ and $L_{P_{jk}}$ denote the resistance and inductance of P_{jk}. After the current paths and their values have been determined for all blocks, the SSN for B_k is given by

$$V_{\text{noise}}^k = \sum_{P_j \in T^k} (i_j \cdot R_{P_{jk}} + L_{P_{jk}} \frac{di_j}{dt}$$

where i_j is the current in the path P_j, which is the sum of all currents through this path to various consumers. The weight of i_j and its rate of change are the resistive and inductive components of the path.

Let Q^k denote the maximum charge drawn from the power supply by block B_k. If $\theta = \max(1, V_{\text{noise}}^k / V_{\text{noise}}^{\text{lim}})$, where $V_{\text{noise}}^{\text{lim}}$ is the noise tolerance, the decap allocated to block B_k is given by

$$D^k = \frac{(1 - 1/\theta) Q^k}{V_{\text{noise}}^{\text{lim}}}, \quad 1 \leq k \leq M$$

where M denotes the total number of blocks to be placed. If $\theta = 1$, this means the noise level for block k is below the constraint and thus $D^k = 0$. If $\theta > 1$, D^k is proportional to the noise level. The maximum value for D^k is $Q^k / V_{\text{noise}}^{\text{lim}}$, where the decap fulfills the entire current demand of block k. Finally, the decap cost is given by $D^{\text{tot}} = \sum_{k=1}^{M} D^k$.

C. 3-D Decoupling Capacitor Placement

The 3-D design offers a new opportunity to improve the decap efficiency via the concept of *footprint-aware decap insertion*. For a given block that requires decap to suppress its switching noise, we allow its decap to be inserted in other layers as long as the distance to the decap is close enough and the overall footprint area is minimally affected. In general, the distance to the power pins in a 3-D design is reduced compared to its 2-D counterpart, so we expect that the overall decap cost (= the area overhead) will be less in 3-D designs. Moreover, our footprint-aware decap insertion further reduces the area cost involved with decap. In our LP-based 3-D decap allocation is formulated as follows:

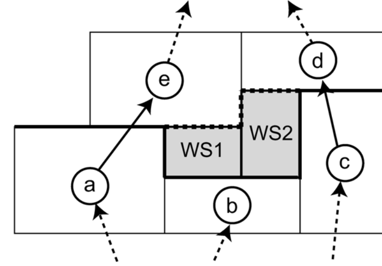

Fig. 4. White-space detection. Blocks a, b, and c are in the lower level. Blocks d and e are in the next level. The bold line is the lower boundary, while the dotted line is the upper boundary. $ws1$ and $ws2$ are the detected white spaces.

$$\text{Maximize } S = \sum_{k=1}^{H} \sum_{j \in N_k} \beta_{kj} x_k^{(j)} \tag{1}$$

subject to

$$\sum_{j \in N_k} x_k^{(j)} \leq A_k, \qquad k = 1, 2, \cdots, H \tag{2}$$

$$\sum_{k=1}^{k=H} x_k^{(j)} \leq S^{(j)}, \qquad j = 1, 2, \cdots, M \tag{3}$$

$$x_k^{(j)} \geq 0 \; \forall k, \forall j \tag{4}$$

where $x_k^{(j)}$ denotes the amount of decap allocated from white space k to block j. The objective is to maximize the utilization of white spaces for decap allocation. The constraint (2) limits the total allocation of a white space to its total area, where A_k denotes the area of white space k, and N_k denotes the neighboring blocks of k. The constraint (3) ensures that the total amount of decap allocated to each block does not exceed its demand, where $S^{(j)}$ denotes the demand.

Unlike the 2-D decap assignment as done in [9], the neighboring blocks in the 3-D case are the adjacent blocks either from the same placement layer or from neighboring layers. The assumption here is that the white space from different layers can be used to allocate decap. In order to facilitate this, we introduce parameters β_{kj} to control decap

[3]We assume that the dimension of our 3-D power mesh as well as 3-D thermal mesh discussed in Section IV-E remains fixed during the module placement. We determine the dimension so that there always exists at least one node in the mesh that covers each module regardless of the module location change during the optimization.

TABLE I
FOUR-LAYER FLOORPLANNING WITH VARIOUS OBJECTIVES. AREA IS IN mm^2, WIRELENGTH IN METER, DECAP IN nF, AND TEMPERATURE IN $^\circ C$

ckts		area/wire-driven				decap-driven				thermal-driven				multi-objective			
name	size	area	wire	decap	temp	area	wire	decap	temp	area	wire	decap	temp	area	wire	decap	temp
n50	50	221	26.6	18.0	87.2	232	30.5	5.2	85.2	377	35.9	29.7	68.9	294	35.5	9.3	76.2
n50b	50	255	25.1	13.1	75.5	262	27.6	4.1	74.4	381	34.2	15.0	63.2	279	27.8	6.5	67.7
n50c	50	212	23.5	15.9	80.6	229	29.2	2.2	70.9	314	30.1	19.5	65.1	233	28.2	5.4	72.2
n100	100	315	66.6	78.2	86.5	343	73.1	69.2	81.7	493	84.1	93.6	69.8	410	77.0	77.9	78.5
n100b	100	280	48.3	73.0	104.8	303	56.1	59.1	98.4	501	69.5	98.7	78.2	311	54.2	64.9	100.6
n100c	100	328	62.1	80.3	90.9	342	67.6	67.5	90.4	529	84.3	94.9	73.8	421	77.9	76.6	81.7
n200	200	560	17.1	226.3	96.4	693	20.5	223.2	96.2	1077	24.5	243.6	76.2	824	21.3	229.1	85.4
n200b	200	556	18.2	233.2	97.7	653	21.6	221.5	90.1	1031	25.6	251.2	76.2	555	18.2	233.2	97.7
n200c	200	542	16.9	237.4	100.9	613	19.6	228.6	103.1	1135	24.3	261.3	76.1	541	16.9	237.4	100.9
n300	300	846	28.6	393.8	100.1	843	28.6	393.8	100.1	1310	38.8	405.3	86.6	844	28.6	393.8	100.1
gt100	100	191	13.2	60.8	71.0	207	16.8	42.5	70.9	474	20.4	92.7	52.3	264	18.6	55.2	59.2
gt300	300	238	19.6	342.5	93.2	248	22.3	334.9	99.5	528	28.0	392.1	72.1	256	22.3	343.9	85.3
gt400	400	270	28.1	493.1	114.0	268	32.5	482.0	111.6	362	37.0	512.0	89.2	282	34.6	492.6	91.1
gt500	500	316	30.3	645.3	99.7	321	35.4	632.4	98.0	541	38.5	684.4	80.3	321	34.8	635.8	95.8
gt600	600	475	60.8	806.5	115.6	509	71.1	794.8	106.8	777	76.5	835.0	84.6	527	68.5	794.2	96.5
RATIO		1.00	1.00	1.00	1.00	1.07	1.15	0.79	0.97	1.76	1.28	1.19	0.79	1.15	1.16	0.87	0.91
TIME		563				304				512				367			

allocation to block j from white space module k. β_{kj} evaluates the usefulness of whitespace k to be used as a decap for block j. Given a 3-D placement, β_{kj} is computed as follows: $\beta_{kj} = \gamma/dist(k,j)$, where $dist(k,j)$ denotes the 2-D distance between k and j, and γ is a user-specified parameter. This equation prefers a shorter distance between k and j. In addition, it captures the case when a neighboring white space located in different layer is easier to access. γ determines how big should be the impact of this distance-based term has on the decap utilization.

The decap allocation involves several iteration between white space insertion and LP solving to reach a good solution both in terms of completeness of decap allocated and the additional increase in area. The white space is generated by expanding the area in the X and Y directions, as illustrated in Fig. 3. In a multilayer placement, however, we have to take additional care to balance the expansion in each layer, so as to minimize the expansion of the total footprint area. The expansion is proportional to the decap demand of each module. In our sequence pair-based 3-D placement, we modify the horizontal and vertical constraint graphs to expand the placement into X and Y directions, respectively.

D. Whitespace Detection

The white space present in a placement can be used to fabricate decap. If the existing white space is insufficient or unreachable by modules needing decap, then white space insertion through area expansion may be necessary. Hence, detection of all existing white spaces in a placement is highly desirable. This is done by using the longest path-tree calculation based on the vertical constraint graph. All nodes at the ith level in the tree are at an edge distance of i from the source node. Each level is ordered by the horizonal constraint graph. The white spaces at level i are detected by comparing the *upper* boundary of blocks at level i and the *lower* boundary of the blocks at level $i+1$. If the boundaries are not incident on each other, then there is whitespace. In Fig. 4, blocks a, b, and c are in the same level and blocks d and e are in the next level. The algorithm compares the upper boundary of a, b, and c, to the lower boundary of d and e. The mismatched boundaries allows the algorithm to find white spaces $ws1, ws2$. This algorithm is capable of detecting all white spaces, and runs in $O(n)$ time, given the ordered longest path tree, where n is the total number of blocks.

E. 3-D Thermal Analysis

We use a 3-D mesh to apply the well-known finite-difference approximation for thermal analysis. Each node models a small volume of the 3-D circuit, and each edge denotes the connectivity between two adjacent regions. Our thermal model is based on the steady-state thermal equation $-k\nabla^2 T = P$, where k is thermal conductivity, T is temperature, and P is power. Taking the finite-difference approximation to form a thermal resistance grid gives the matrix equation $R \cdot P = T$, where R is a thermal resistance matrix, T is a temperature vector, and P is a power vector. Assuming that the thermal conductivity of device blocks are similar (they are mostly silicon), swapping the location device blocks would not change the thermal resistance matrix R. This means that matrix R only needs to be computed once in the beginning. To calculate the temperature profile of a new block configuration, the power profile P needs to be updated and then multiplied by R. Alternatively, a change in power profile ΔP can be defined. Multiplying R and ΔP will give a change in temperature profile ΔT. Adding ΔT to the old temperature profile will give the new temperature profile.

Our thermal analyzer takes the interconnect capacitance into consideration during the ΔP computation. For each active node in the thermal mesh, ΔP includes both module and interconnect power. Since $\Delta T = R \cdot \Delta P$, our thermal model considers the interconnect for the temperature calculation. A more accurate way is to update R, but this is too costly and is not necessary for placement optimization.

V. EXPERIMENTAL RESULTS

We implemented the proposed algorithms and analysis tools using C++/STL. Our program was evaluated using the GSRC and a set of synthesized benchmarks, referred to as GT. The GT benchmarks were generated from ISPD98 benchmarks [10] by partitioning the circuits to desired sizes. The number of layers were fixed to four for all benchmarks. In our experiments, we used the same simulated annealing parameters across the benchmarks. The technology parameters used were obtained from [11]. These are our observations from Table I.

- We achieved a 21% improvement over baseline in decap cost using our decap-driven algorithm, with a 7% increase in final area and 15% increase in wirelength. The temperature decreases by 3%.
- Thermal-driven floorplanning achieved a 21% improvement over baseline with a dramatic increase in total area of 76%. Wirelength increased by 28% and decap requirement *increases* by 19%. We note that it may be possible to reduce area and wirelength costs by fine-tuning parameters.
- Simultaneous thermal and decap optimization during floorplanning leads to an improvement in *both* decap and temperature over baseline by 13% and 9%, respectively. There is a reasonable increase in final area and wirelength by 19% and 15%, respectively.

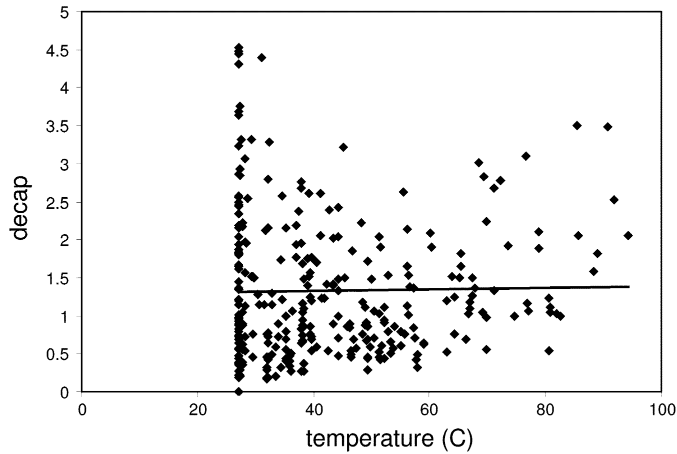

Fig. 5. Decap versus temperature correlation plot for n300.

Smaller decap cost implies the following advantages. First, smaller decap budget translates to a cheaper manufacturing cost for decap. Second, the leakage power associated with decap is rapidly increasing, so our decap saving translates to leakage power saving. Note that it is always possible to control the overall runtime of the placement by changing the cooling schedule of annealing process. This is done to facilitate fair comparison.

Fig. 5 shows the temperature and decap requirement of each block of the final floorplan of the n300 benchmark. The figure clearly shows that there is little correlation between temperature and decap. This shows that minimizing one objective does not necessarily minimize the other objective as a positive correlation would indicate. It also means that minimizing both objectives is not mutually exclusive, as a negative correlation would indicate. This matches with the block placement results.

VI. Conclusion

In this paper, we developed algorithms to handle thermal and power-supply noise issues during module placement for 3-D SOP. We extended the power supply network and thermal models to 3-D and used them to guide our 3-D module placement. In addition to estimating the amount of decap needed to keep the SSN at the circuits tolerance level, we also efficiently used nearby whitespace to allocate decap and adding more area to the placement if necessary. Our experimental results demonstrated the effectiveness of our approach.

References

[1] R. Tummala, "SOP: What is it and why? A new microsystem-integration technology paradigm-Moore's law for system integration of miniaturized convergent systems of the next decade," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 241–249, 2004.

[2] R. Li, G. DeJean, M. Tentzeris, and J. Laskar, "Integration of miniaturized patch antennas with high dielectric constant multilayer packages and soft-and-hard surfaces (SHS)," in *Proc. IEEE Electron. Compon. Technol. Conf.*, 2003, pp. 474–477.

[3] T. Sudo, H. Sasaki, N. Masuda, and J. Drewniak, "Electromagnetic interference (EMI) of system-on-package (SOP)," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 304–314, 2004.

[4] K. Kundert, H. Chang, D. Jefferies, G. Lamant, E. Malavasi, and F. Sendig, "Design of mixed-signal systems-on-a-chip," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 19, no. 12, pp. 1561–1571, Dec. 2000.

[5] G. Gielen and R. Rutenbar, "Computer-aided design of analog and mixed-signal integrated circuits," *Proc. IEEE*, vol. 88, no. 12, pp. 1825–1854, Dec. 2000.

[6] Y. Deng and W. Maly, "Interconnect characteristics of 2.5-D system integration scheme," in *Proc. Int. Symp. Phys. Design*, 2001, pp. 171–175.

[7] S. Das, A. Chandrakasan, and R. Reif, "Design tools for 3-D integrated circuits," in *Proc. Asia South Pacific Design Automat. Conf.*, 2003, pp. 53–56.

[8] H. Murata, K. Fujiyoshi, S. Nakatake, and Y. Kajitani, "Rectangle packing based module placement," in *Proc. IEEE Int. Conf. Comput.-Aided Design*, 1995, pp. 472–479.

[9] S. Zhao, C. Koh, and K. Roy, "Decoupling capacitance allocation and its application to power supply noise aware floorplanning," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 27, no. 1, pp. 81–92, 2002.

[10] C. J. Alpert, "The ISPD98 circuit benchmark suite," *Proc. Int. Symp. Phys. Design*, pp. 80–85, 1998.

[11] M. Arik, J. Garg, and A. Bar-Cohen, "Thermal modeling and performance of high heat flux SOP packages," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 398–412, May 2004.