Exploiting Die-to-Die Thermal Coupling in 3D IC Placement*

Krit Athikulwongse, Mohit Pathak, and Sung Kyu Lim
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA, USA
{krit,mohitp,limsk}@ece.gatech.edu

ABSTRACT

In this paper, we propose two methods used in 3D IC placement that effectively exploit the die-to-die thermal coupling in the stack. First, TSVs are spread on each die to reduce the local power density and vertically aligned across dies simultaneously to increase thermal conductivity to the heatsink. Second, we move high-power logic cells to the location that has higher conductivity to the heatsink while moving TSVs in the upper dies so that high-power cells are vertically overlapping below the TSVs. These methods are employed in a force-directed 3D placement successfully and outperform several state-of-the-art placers published in recent literature.

Categories and Subject Descriptors

B.7.2 [**Integrated Circuits**]: Design Aids—*Placement and routing*

General Terms

Algorithms, Design, Reliability

Keywords

3D IC, TSV, Temperature

1. INTRODUCTION

Stacking thinned dies in 3D ICs results in increasing power density, thus rising temperature, which leads to other reliability problems, such as electromigration and negative-bias-temperature instability. Because of low thermal conductivity, polymer adhesive exacerbates the problem. Moreover, if the thinned dies are silicon on insulator, an extremely high temperature can be expected. Heat must be removed from the die quickly; otherwise, reliability problems may arise.

A few recent works on temperature-aware placement for 3D ICs have been published. In [2], a force-directed approach was proposed for 3D thermal placement; however, it did not include through-silicon vias (TSVs), which are commonly found in 3D ICs. In [3], a partitioning-based approach was proposed for 3D thermal placement. The work considered the impact of parasitic resistance and capacitance of signal TSVs on power, but failed to include thermal properties of TSVs. Failing again to acknowledge TSV area, it also reported unreasonably large numbers of TSVs even for small circuits. The work in [1] considered TSV thermal properties; however, it assumed that adhesive is an ideal insulator. In reality, heat can still flow through (silicon and) adhesive because of its thinness. Based on the assumption, the work balanced only the number of TSVs in a bin to heat dissipated from cells in the same bin and bins vertically below.

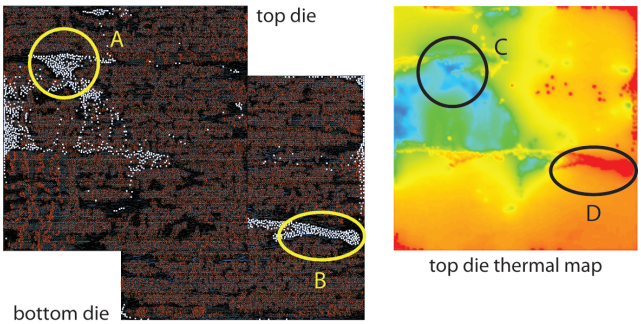

Figure 1: Die-to-die heat coupling from TSVs. TSVs are shown in white. The top die is closer to heatsink. The cold spot C is caused by the TSVs in spot A on the same die. The hot spot D is caused by the TSVs in spot B from the bottom die.

The contributions of this work are as follows: (i) We propose two effective heuristics, namely TSV spread and alignment method (TSA) and thermal coupling-aware placement (CA), that exploit the die-to-die thermal coupling in 3D ICs in force-directed temperature-aware placement. We present new forces, and discuss how to manage them to obtain high quality placements. (ii) We perform extensive experiments to show the trade-off among wirelength, delay, power, and temperature results obtained from GDSII layouts. (iii) Our placers outperform several state-of-the-art placers published in recent literature [2, 5, 3, 1, 4].

2. MOTIVATION

In a 3D IC layout, logic gates cannot overlap with TSVs. Area occupied by TSVs becomes "power whitespace" because no power is consumed and thus no heat is generated. In addition, TSVs conduct majority of heat through polymer adhesive between dies toward the heatsink as shown in Fig. 1. In the figure, the hotspot D on the top metal layer of the top die is caused by the TSVs in spot B from the bottom die. Heat flows through TSVs so intensely that its effect still remains on the top die. Thus, the temperature

*This research is funded by SRC ICSS Task 1836.075 and SRC CADTS Task 2239.001.

DAC 2012, June 3-7, 2012, San Francisco, California, USA.

distribution of the top die results from the combination of power profile of the top die and heat flowing from the bottom die through TSVs. Our TSV spread and alignment method presented in this paper exploits these thermal properties of TSVs by distributing TSVs evenly to reduce power density in local power hotspots and vertically aligning TSVs of adjacent dies to establish direct paths to the heatsink.

Using Ansys FLUENT, we simulate a part of bulk silicon with and without TSVs (and their related structures, e.g., landing pad and liner). We fix the temperature on the top side of the models, apply constant power density on the bottom side, and obtain the temperature distribution. The simulation results indicate that heat flowing through a TSV increases temperature far less than the same amount of heat flowing through bulk silicon and adhesive. We also observe that the temperature slowly increases in bulk silicon with TSVs. On the other hand, in bulk silicon without any TSV, low thermal conductivity of bonding adhesive results in steep temperature rise at first, but temperature does not rise as much inside the silicon. We compute the average thermal conductivity of bulk silicon with and without TSVs, and use them to guide our thermal coupling-aware placer presented in this paper.

3. GLOBAL PLACEMENT ALGORITHMS

We extend the force-directed placer [7] in two ways to perform thermal optimization in 3D ICs. In the first algorithm, we laterally spread TSVs in each die to form even thermal conductivity while perturbing TSV position to increase vertical overlap among TSVs across the dies in 3D stack. In the second algorithm, the logic cells on each die are positioned by using thermal conductivity-based force while TSVs are positioned by using power density-based force.

3.1 Design Flow

Fig. 2 shows the overall flow of our placement, where the position of cells and TSVs is determined simultaneously. Given a netlist, we partition cells into dies if the partition is not also given. Then, we insert the minimum number of TSVs required to connect cells on different dies. Once this die partitioning is fixed, we do not move cells across dies during placement. The reason is that changing cell partition results in change in the number of TSVs, and this change causes the complexity of problem to become unmanageable. Next, we minimize wirelength to obtain initial placement, which may contain high overlap among cells and TSVs. In the main loop to resolve the overlap, we use TSV density and TSV position to compute target point for TSVs in the first algorithm. In the second algorithm, we periodically perform 3D power analysis based on current cell and TSV position. Then, we use the cell power, TSV density, and average thermal conductivity of bulk silicon obtained from the simulation results in Section 2 to compute target points for cells and TSVs to move towards. After updating force equations and solving them, we update the position of cells and TSVs. This loop continues until the overlap is sufficiently reduced.

3.2 Force-directed 3D Placement

In a quadratic placement [7], quadratic wirelength Γ_{x} and Γ_{y} along x- and y-axis are separately minimized to obtain the placement result. Treated Γ_{x} as spring energy, its derivative can be regarded as net force $\mathbf{f}_{\mathrm{x}}^{\mathrm{net}}$. By setting $\mathbf{f}_{\mathrm{x}}^{\mathrm{net}}$ to zero, the minimum Γ_{x} and the corresponding placement are found; however, cells may overlap in few small areas. Hold force $\mathbf{f}_{\mathrm{x}}^{\mathrm{hold}}$ prevents $\mathbf{f}_{\mathrm{x}}^{\mathrm{net}}$ from pulling cells back to the initial placement. In addition, density-based force $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$ reduces the overlap by spreading cells in high density region.

To extend [7] for 3D ICs, cells are not moved across dies during placement in [4] because they are already assigned into dies by the partitioner. In addition, $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$ is computed die-by-die based on the placement density D_d of each die d, which is defined as

$$D_d(x,y) = D_d^{\mathrm{cell}}(x,y) - D_d^{\mathrm{die}}(x,y), \tag{1}$$

where D_d^{cell} is the cell density on die d, and D_d^{die} is the die capacity scaled to match the total cell area on the die. Then, the placement potential Φ_d is computed by solving Poisson's equation

$$\Delta\Phi_d(x,y) = -D_d(x,y). \tag{2}$$

The target point $\mathring{x}_i^{\mathrm{d}}$ to connect density-based spring of cell i is computed by

$$\mathring{x}_i^{\mathrm{d}} = x_i' - \frac{\partial}{\partial x}\Phi_d(x,y)\Big|_{(x_i',y_i')}, \tag{3}$$

where x_i' is the x-position of cell i on die d from the last iteration. Lastly, for each placement iteration, the placement result can be obtained by setting total force $\mathbf{f}_{\mathrm{x}}$ to zero, and solve

$$\mathbf{f}_{\mathrm{x}} = \mathbf{f}_{\mathrm{x}}^{\mathrm{net}} + \mathbf{f}_{\mathrm{x}}^{\mathrm{hold}} + \mathbf{f}_{\mathrm{x}}^{\mathrm{den}} = \mathbf{0}. \tag{4}$$

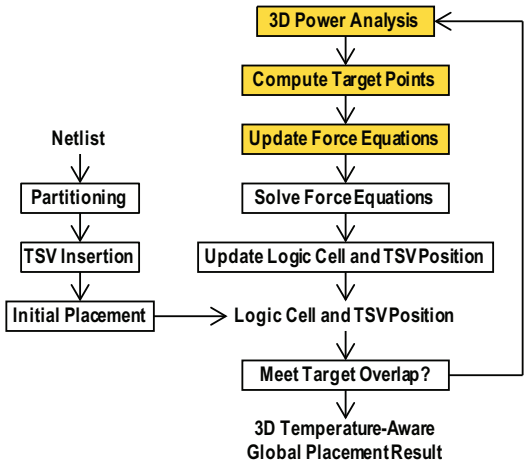

Figure 2: Design flow for our 3D IC global placement.

3.3 TSV Spread and Alignment

In this algorithm, we exploit one of thermal properties of TSVs to help alleviate thermal problems as shown in Fig. 3(a). TSVs occupy placement area, but do not dissipate power. The existence of TSVs among cells with high power dissipation reduces local dissipated power density, which in-turn helps reduce local temperature. Therefore, spreading TSVs evenly on each die should help reduce intra-die thermal variation in 3D ICs. We propose this algorithm because it is simple yet effective. It can be viewed as a method to mimic uniform TSV position. Instead of moving TSVs based on the placement density computed from both TSV and cell area, we move TSVs based on TSV density only. In other words, we compute D_d^{cell} in Equation (1) from TSV area only, and scale D_d^{die} to match the total TSV area on the die.

In addition to TSV spread, we exploit another thermal property of TSVs to help alleviate thermal problems as shown in Fig. 3(b). TSVs conduct majority of heat through polymer adhesive between dies, causing local hot spots on the adjacent die between the TSVs and heatsink. Therefore, aligning TSVs on each die to TSVs on the adjacent die should help prevent this kind of hot spots, and direct the heat toward the heatsink quickly, resulting in overall temperature decrease. To align TSVs during global placement, we introduce an additional force for TSVs, alignment force denoted

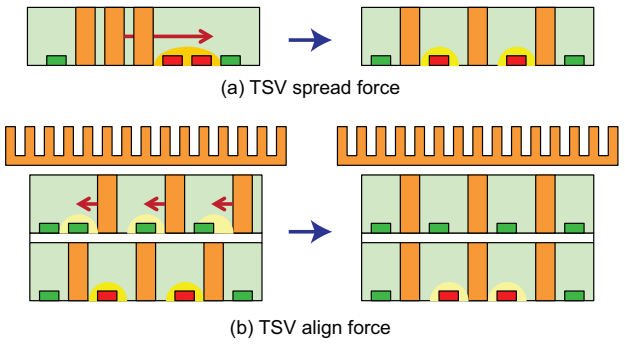

Figure 3: TSV spread and TSV align forces.

$\mathbf{f}_{\mathrm{x}}^{\mathrm{align}}$, into Equation (4). This force can be represented by alignment springs connected to TSVs, and defined as

$$\mathbf{f}_{\mathrm{x}}^{\mathrm{align}} = \mathring{\mathbf{C}}_{\mathrm{x}}^{\mathrm{a}}(\mathbf{x} - \mathring{\mathbf{x}}^{\mathrm{a}}), \tag{5}$$

where vector $\mathring{\mathbf{x}}^{\mathrm{a}}$ represents the x-position of target points to connect alignment springs to TSVs, and diagonal matrix $\mathring{\mathbf{C}}_{\mathrm{x}}^{\mathrm{a}}$ collects spring constants $\mathring{w}_{\mathrm{x},i}^{\mathrm{a}}$ of the alignment spring connected to TSV i.

We apply alignment force to TSV i only when its closest TSV j on the adjacent die farther from the heatsink is within a certain range so that we do not excessively increase wirelength. The range is set to the size of TSV because of the high probability of aligning the TSVs in few iterations. We balance $\mathbf{f}_{\mathrm{x}}^{\mathrm{align}}$ against other forces by setting $\mathring{w}_{\mathrm{x},i}^{\mathrm{a}}$ to density-based spring constant $\mathring{w}_{\mathrm{x},i}^{\mathrm{d}}$ of $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$ and setting alignment target point $\mathring{x}_i^{\mathrm{a}}$ to x_j', the x-position from last iteration of TSV j (on the adjacent die farther from heatsink) closest to TSV i. This method naturally balances $\mathbf{f}_{\mathrm{x}}^{\mathrm{align}}$ against $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$.

The intuition is that because of the high cell overlap in the early placement iterations, the target point $\mathring{x}_i^{\mathrm{d}}$ is farther away from TSV i than the alignment target point $\mathring{x}_i^{\mathrm{a}}$. Thus, $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$ dominates. When cells are evenly distributed in the late iterations of placement, $\mathring{x}_i^{\mathrm{d}}$ is closer to TSV i. Then, $\mathbf{f}_{\mathrm{x}}^{\mathrm{den}}$ becomes weaker, and $\mathbf{f}_{\mathrm{x}}^{\mathrm{align}}$ affects the TSV position more.

3.4 Thermal Coupling-aware Placement

In this algorithm, we consider the die-to-die thermal coupling during placement. The basic approach is to introduce two new forces, the first that moves cells and the second that moves TSVs, both in an attempt to place high-power cells closer to the TSV-to-heatsink path. Since the heat dissipated by a cell must flow toward heatsink, we place cells based on their power density and the effective thermal conductivity computed using the same die *and the dies above*. In addition, since TSV conducts heat without raising temperature too much, we place TSVs based on the total power density of the same die *and the dies below*.

Our basic approach is that the area with high power density and low thermal conductivity leads to high temperature. Thus, the temperature at a certain position depends on the difference (or imbalance) between power density and thermal conductivity. The force that moves cells (TSVs) on a die also changes the power density (thermal conductivity) distribution of the die. Our goal is to use these forces to balance the power density and the thermal conductivity at each position on the die. The force in an area with high difference should be stronger than the force in an area with low difference. The strength of a spring force depends on the distance to the connection point, so we set the strength based on this difference. Based on this concept, we first build a map of the difference, and smooth the map in an iterative fashion.

3.4.1 For Cell Movement

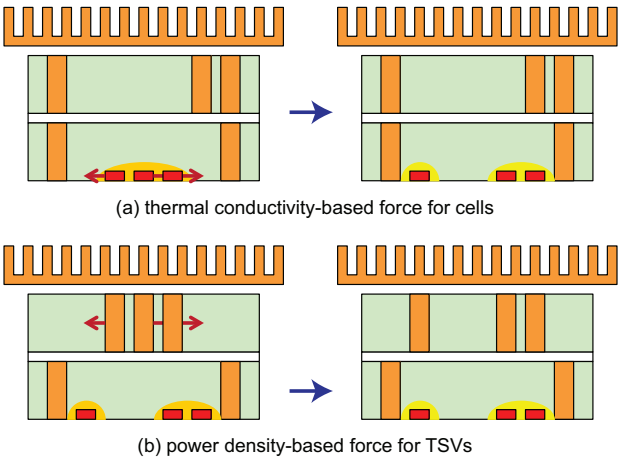

Figure 4: Thermal conductivity-based vs power density-based forces.

We introduce the thermal conductivity-based force $\mathbf{f}_{\mathrm{x}}^{\mathrm{cond}}$ as illustrated in Fig. 4(a). It moves high-power cells toward the position with high thermal conductivity to heatsink, and is defined as

$$\mathbf{f}_{\mathrm{x}}^{\mathrm{cond}} = \mathring{\mathbf{C}}_{\mathrm{x}}^{\mathrm{c}}(\mathbf{x} - \mathring{\mathbf{x}}^{\mathrm{c}}), \tag{6}$$

where the vector $\mathring{\mathbf{x}}^{\mathrm{c}}$ represents the x-position of target points to connect thermal conductivity-based springs to cells, and the diagonal matrix $\mathring{\mathbf{C}}_{\mathrm{x}}^{\mathrm{c}}$ contains spring constants $\mathring{w}_{\mathrm{x},i}^{\mathrm{c}}$ of the spring connected to cell i.

We compute $\mathbf{f}_{\mathrm{x}}^{\mathrm{cond}}$ die-by-die by balancing the cell power density P_d^{cell} of each die d against its effective thermal conductivity to heatsink, denoted K_d^{sink}. Under the demand-supply system of the force-directed framework in [7], P_d^{cell} and K_d^{sink} represent the demand and supply to remove the heat from die d in the 3D stack. We define the thermal conductivity-based balance factor B_d^{cond} for die d as (see Fig. 5)

$$B_d^{\mathrm{cond}}(x,y) = P_d^{\mathrm{cell}}(x,y) - s_d^{\mathrm{cond}} \cdot K_d^{\mathrm{sink}}(x,y), \tag{7}$$

where s_d^{cond} is a scaling factor to match K_d^{sink} to P_d^{cell} across the die. We use s_d^{cond} to balance the total supply (K_d^{sink}) and the total demand (P_d^{cell}), and compute it by

$$s_d^{\mathrm{cond}} = \frac{\int\int P_d^{\mathrm{cell}}(x,y)\,dx\,dy}{\int\int K_d^{\mathrm{sink}}(x,y)\,dx\,dy}. \tag{8}$$

Here, K_d^{sink} is computed as

$$K_d^{\mathrm{sink}}(x,y) = \frac{1}{\sum_{j=d}^{N_{\mathrm{die}}} \frac{1}{K_j^{\mathrm{die}}(x,y)}}, \tag{9}$$

where K_j^{die} is the thermal conductivity of die j, and die N_{die} is the die closest to the heatsink (see Fig. 6). Here, $K_{N_{\mathrm{die}}}^{\mathrm{die}}$ includes the thermal conductivity of thick substrate and heatsink, and K_j^{die} is computed based on the TSV density at each position on the die and the average thermal conductivity of bulk silicon with and without TSVs, obtained from the simulation results in Section 2.

The potential Φ_d^{cond} for B_d^{cond} is computed by solving Poisson's equation

$$\Delta\Phi_d^{\mathrm{cond}}(x,y) = -B_d^{\mathrm{cond}}(x,y). \tag{10}$$

The target point $\mathring{x}_i^{\mathrm{c}}$ of cell i is computed by

$$\mathring{x}_i^{\mathrm{c}} = x_i' - \frac{\partial}{\partial x}\Phi_d^{\mathrm{cond}}(x,y)\Big|_{(x_i',y_i')}, \tag{11}$$

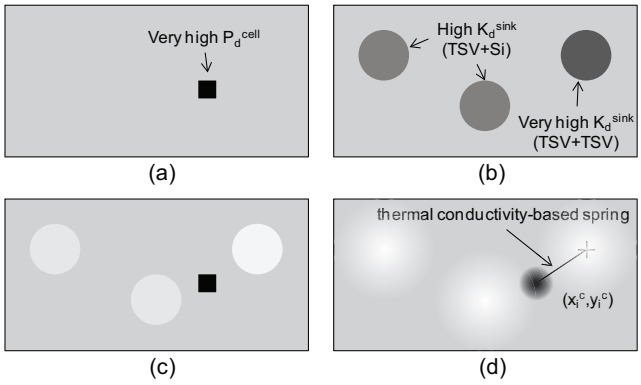

Figure 5: Illustration of B_d^{cond}. (a) P_d^{cell}, (b) $s_d^{\text{cond}} \cdot K_d^{\text{sink}}$, (c) B_d^{cond}, (d) potential for B_d^{cond} after solving Poisson's equation.

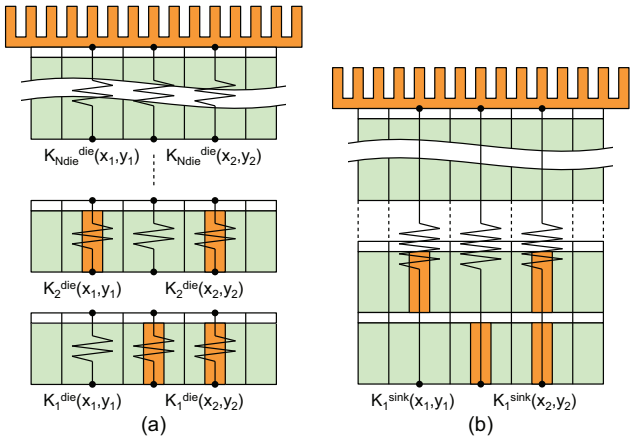

Figure 6: Computation of K_d^{sink}. (a) K_j^{die}, (b) K_1^{sink}.

where x_i' is the x-position of cell i on die d from the last iteration. We set spring constant $\mathring{w}_{\text{x},i}^{\text{c}}$ for cell i based on cell power and the total cell power by

$$\mathring{w}_{\text{x},i}^{\text{c}} = p_i / \sum_{\forall j} p_j, \tag{12}$$

where p_i is the power of cell i, and j is a cell on die d. Therefore, a high-power cell is connected to a strong thermal conductivity-based spring.

3.4.2 For TSV Movement

We introduce power density-based force $\mathbf{f}_{\text{x}}^{\text{pow}}$ as illustrated in Fig. 4(b). It moves TSVs toward the position with high cell power density on the same die and the dies below. We define $\mathbf{f}_{\text{x}}^{\text{pow}}$ as

$$\mathbf{f}_{\text{x}}^{\text{pow}} = \mathring{\mathbf{C}}_{\text{x}}^{\text{p}}(\mathbf{x} - \mathring{\mathbf{x}}^{\text{p}}), \tag{13}$$

where the vector $\mathring{\mathbf{x}}^{\text{p}}$ represents the x-position of target points to connect power density-based springs to TSVs, and the diagonal matrix $\mathring{\mathbf{C}}_{\text{x}}^{\text{p}}$ contains spring constants $\mathring{w}_{\text{x},i}^{\text{p}}$ of the spring connected to TSV i.

We compute $\mathbf{f}_{\text{x}}^{\text{pow}}$ die-by-die by balancing the thermal conductivity K_d^{die} of each die d against the total power density $\sum P_j^{\text{cell}}$ that flows through the die toward heatsink. Under the demand-supply system of the force-directed framework in [7], K_d^{die} and $\sum P_j^{\text{cell}}$ represent the demand and supply to conduct heat from the same die and dies below to heatsink. We define the power density-based balance factor B_d^{pow} for die d as

$$B_d^{\text{pow}}(x,y) = K_d^{\text{die}}(x,y) - s_d^{\text{pow}} \cdot \sum_{j=1}^{d} P_j^{\text{cell}}(x,y), \tag{14}$$

where s_d^{pow} is a scaling factor to match $\sum P_j^{\text{cell}}$ to K_d^{die} across the die. We use s_d^{pow} to balance the total supply ($\sum P_j^{\text{cell}}$) and the total demand (K_d^{die}), and compute it by

$$s_d^{\text{pow}} = \frac{\int\int K_d^{\text{die}}(x,y)\,dx\,dy}{\int\int \sum_{j=1}^{d} P_j^{\text{cell}}(x,y)\,dx\,dy}. \tag{15}$$

The potential Φ_d^{pow} for B_d^{pow} is computed by solving Poisson's equation

$$\Delta\Phi_d^{\text{pow}}(x,y) = -B_d^{\text{pow}}(x,y). \tag{16}$$

The target point $\mathring{x}_i^{\text{p}}$ of TSV i is computed by

$$\mathring{x}_i^{\text{p}} = x_i' - \frac{\partial}{\partial x}\Phi_d^{\text{pow}}(x,y)\Big|_{(x_i',y_i')}, \tag{17}$$

where x_i' is the x-position of TSV i on die d from the last iteration. We set spring constant $\mathring{w}_{\text{x},i}^{\text{p}}$ to $1/N_d^{\text{TSV}}$, where N_d^{TSV} is the total number of TSVs on die d. Therefore, the power density-based spring for each TSVs has the same strength.

3.4.3 Balancing the Forces

We balance the new forces against $\mathbf{f}_{\text{x}}^{\text{den}}$ because $\mathbf{f}_{\text{x}}^{\text{den}}$ is the main force that moves cells and TSVs. First, we scale the new forces so that they have the same magnitude as $\mathbf{f}_{\text{x}}^{\text{den}}$. Then, we apply weighting constants to $\mathbf{f}_{\text{x}}^{\text{den}}$, $\mathbf{f}_{\text{x}}^{\text{cond}}$, and $\mathbf{f}_{\text{x}}^{\text{pow}}$ so that we can control their contribution to the total force.

First, to scale $\mathbf{f}_{\text{x}}^{\text{cond}}$ to $\mathbf{f}_{\text{x}}^{\text{den}}$, we normalize P_d^{cell}, the demand for B_d^{cond} in Equation (7), to D_d^{cell} by a scaling factor s_d^{PD} defined as

$$s_d^{\text{PD}} = \frac{\int\int D_d^{\text{cell}}(x,y)\,dx\,dy}{\int\int P_d^{\text{cell}}(x,y)\,dx\,dy}. \tag{18}$$

Then, we replace P_d^{cell} in Equation (7) and Equation (8) by $s_d^{\text{PD}} \cdot P_d^{\text{cell}}$. Second, to scale $\mathbf{f}_{\text{x}}^{\text{pow}}$ to $\mathbf{f}_{\text{x}}^{\text{den}}$, we normalize K_d^{die}, the demand for B_d^{pow} in Equation (14), to D_d^{cell} by a scaling factor s_d^{KD} defined as

$$s_d^{\text{KD}} = \frac{\int\int D_d^{\text{cell}}(x,y)\,dx\,dy}{\int\int K_d^{\text{die}}(x,y)\,dx\,dy}. \tag{19}$$

Then, we replace K_d^{die} in Equation (14) and Equation (15) by $s_d^{\text{KD}} \cdot K_d^{\text{die}}$.

We scale both $\mathbf{f}_{\text{x}}^{\text{cond}}$ and $\mathbf{f}_{\text{x}}^{\text{pow}}$ to $\mathbf{f}_{\text{x}}^{\text{den}}$ based on D_d^{cell}, not on the gradient of Φ_d because of the stability issue. After normalizing P_d^{cell} and K_d^{die} to D_d^{cell} as shown in Equation (18) and Equation (19), the magnitude of B_d^{cond} and B_d^{pow} and gradient of their potential are properly normalized. At an equilibrium, a small magnitude of the gradients results in a small magnitude of $\mathbf{f}_{\text{x}}^{\text{cond}}$ and $\mathbf{f}_{\text{x}}^{\text{pow}}$. If we scale $\mathbf{f}_{\text{x}}^{\text{cond}}$ and $\mathbf{f}_{\text{x}}^{\text{pow}}$ to $\mathbf{f}_{\text{x}}^{\text{den}}$ based on the gradient of Φ_d instead, the magnitude of the gradient of potential of B_d^{cond} and B_d^{pow} would be exaggerated after the normalization, which in turn causes instability.

In summary, $\mathbf{f}_{\text{x}}^{\text{cond}}$ moves cells in such a way that high power density flows through the paths with high thermal conductivity to heatsink. In addition, $\mathbf{f}_{\text{x}}^{\text{pow}}$ moves TSVs in such a way that each TSV establishes a heat path for the high-power cells in the same die and the dies below. Our overall force equation is as follows:

$$\mathbf{f}_{\text{x}} = \mathbf{f}_{\text{x}}^{\text{net}} + \mathbf{f}_{\text{x}}^{\text{hold}} + (1-\alpha)\mathbf{f}_{\text{x}}^{\text{den}} + \alpha(\mathbf{f}_{\text{x}}^{\text{cond}} + \mathbf{f}_{\text{x}}^{\text{pow}}) = \mathbf{0}. \tag{20}$$

By increasing α, the forces $\mathbf{f}_{\text{x}}^{\text{cond}}$ and $\mathbf{f}_{\text{x}}^{\text{pow}}$ dominate the movement of cells and TSVs for more thermal optimization.

Table 1: Benchmark circuits.

Ckt.	#Gates	#TSVs	Util.	Footprt (mm^2)	Profile
ckt1	119,040	5,725	0.66	0.50×0.50	Data encryption
ckt2	191,420	24,540	0.63	0.90×0.90	Graphic accelerator
ckt3	280,933	17,362	0.49	0.98×0.98	Video compression
ckt4	383,329	17,436	0.53	1.04×1.04	Signal processing
ckt5	644,357	15,024	0.53	1.16×1.16	Image encoder

4. EXPERIMENTAL RESULTS

We use 45-nm technology from FreePDK45 for our experiments. TSV diameter is 5 μm, and the landing pad width is 7 μm. TSV liner thickness is 250 nm. We use copper TSVs with SiO_2 liner and 2.6-μm-thick benzocyclobutene bonding adhesive for our experiments. Each die in the 3D chip stack is thinned to 30 μm except that the topmost die, which is attached to heatsink, retains its thickness at 530 μm. The ambient temperature on top of the heatsink is 300 K. The TSV parasitic resistance and capacitance are $0.1\,\Omega$ and $125\,fF$, respectively. We base all our experiments on 4-die chip stacks.

We use IWLS 2005 benchmarks and several industrial circuits from OpenCores. We synthesize the circuits using Synopsys Design Compiler to obtain gate-level netlist, and use the target clock period of each circuit when performing all analyses. The benchmark characteristics are listed in Table 1. The numbers of TSVs are based on partitioning results from our own implementation of [3]. We use the same die partitioning results for all algorithms for fair comparison in Section 4.2. Because [3] does not consider TSV area, it inserts high number of TSVs, resulting in low placement utilization.

We do not optimize the circuits after placement because buffers and sized gates can change power profile, thus affecting temperature. The results reported in this paper are from commercial tools. We use Cadence Encounter to route the layouts, Synopsys PrimeTime to analyze timing and power, and Ansys FLUENT to analyze temperature. We report all our temperature results in terms of the increase from the ambient temperature measured at the top of the heatsink.

4.1 Impact of TSV Density Uniformity

In this experiment, we show how TSV density uniformity impacts thermal profile. Our two baseline 3D placements are wirelength-driven placement with uniform TSV position [4] and wirelength-driven placement with non-uniform TSV position [4]. First, we obtain both baseline placements using our own implementation of [4]. Then, we perform power and thermal analyses on both placement results. The routed wirelength, longest path delay, and power are shown in Table 2, and temperatures are shown in Table 3. Although the placement with non-uniform TSV position has shorter wirelength, better timing, and lower power than the placement with uniform TSV position, its temperature, especially the thermal variation, is worse. Both the non-uniform power density and the non-uniform thermal conductivity, caused by the non-uniform distribution of TSVs in the 3D chip stack, contribute to the problem. In the placement with non-uniform TSV position, we observe that the area with high TSV density has low power density and low temperature, vice versa. These two opposite trends are responsible for high thermal variation.

4.2 Comparison with State-of-the-Art

We compare our temperature-aware global placement algorithms with the following recent state-of-the-art temperature-aware placers:[1]

[1]This task is challenging due to the discrepancy among the settings and assumptions made in each work. However, we made our best effort to provide fair and meaningful comparison.

Table 2: Routed wirelength, longest path delay, and power of placements with uniform [4] and non-uniform [4] TSV position.

	Uniform			Non-uniform		
Ckt.	rWL (m)	D_{max} (ns)	P (W)	rWL (m)	D_{max} (ns)	P (W)
ckt1	3.897	5.320	0.752	3.014	4.836	0.728
ckt2	11.718	16.510	2.661	7.744	13.694	2.463
ckt3	13.532	8.814	2.353	9.326	6.535	2.288
ckt4	19.355	20.788	2.710	12.457	12.515	2.640
ckt5	22.708	19.772	3.209	18.711	13.798	3.122
ratio	1.405	1.350	1.039	1.000	1.000	1.000

Table 3: Temperature (°C) of placements with uniform [4] and non-uniform [4] TSV position. ($\Delta T_{ja} = T_{ja,max} - T_{ja,min}$)

	Uniform			Non-uniform		
Ckt.	$T_{ja,max}$	ΔT_{ja}	$T_{ja,ave}$	$T_{ja,max}$	ΔT_{ja}	$T_{ja,ave}$
ckt1	71.55	17.60	64.50	74.13	18.33	63.98
ckt2	101.14	47.14	69.41	94.41	50.19	64.78
ckt3	70.38	31.01	55.06	80.09	42.81	55.48
ckt4	64.91	18.76	54.32	75.98	38.01	55.16
ckt5	66.77	35.40	53.13	75.24	39.32	54.50
ratio	1.000	1.000	1.000	1.081	1.325	0.995

[2] (force-directed placer): In this work, thermal analysis is performed at the beginning of every global placement iteration. The thermal gradient obtained from the analysis is used to compute repulsive force, which moves logic cells from high-temperature area toward low-temperature area. We implement our own version of this work by calling Ansys FLUENT from inside our placer, and combining scaled thermal gradient into density-based force $\mathbf{f}_x^{den}$.

[5] (force-directed placer): Instead of moving logic cells based on placement area density, it moves logic cells based on placement power density. Therefore, logic cells are spread according to their power dissipation, and logic cells with high power dissipation occupy more space than logic cells with low power dissipation, leading to uniform power density and thermal profile across the die. We implement our own version of this work.

[3] (partitioning-based placer): In this work, logic cells are partitioned into placement area and different dies based on the switching activity and parasitic capacitance of connecting wires and TSVs. We perform global routing to determine the position of TSVs as proposed in [6] after performing global placement using our own implementation of [3].

[1] (analytical placer): We implement this method by balancing the power density combined across dies in vertical direction against the TSV density and solving the density for potential function. The gradient of potential is used to compute a force to move cells and TSVs to maintain the balance. The force is added to $\mathbf{f}_x^{den}$ with a user-defined parameter β to provide temperature-wirelength tradeoff similar to the work.

Table 4 shows the routed wirelength, delay, power, and temperature comparison based on the GDSII layouts we build using these placers. The wirelength, delay, and power values are normalized to the wirelength-driven non-uniform TSV placement [4] shown in Table 2. The temperature values are normalized to the wirelength-driven uniform TSV placement [4] shown in Table 3. Recall that non-uniform placer achieves high-quality wirelength, delay, and power results while uniform placer leads to high-quality temperature values.

First, we observe that [2] produces comparable wirelength, delay, and power results to non-uniform TSV placer [4]. In case of temperature, [2] obtains worse result compared with uniform TSV placer [4]. We tried increasing the magnitude of thermal-gradient-

Table 4: Comparison with state-of-the-art temperature-aware placers [2, 5, 3, 1, 4]. Our placers are TSA (TSV spread and alignment) and CA (Coupling-aware placement). The routed wirelength, delay, and power values are normalized to the non-uniform TSV placement [4] shown in Table 2. The temperature values are normalized to the uniform TSV placement [4] shown in Table 3.

Ckt.	routed wirelength (m)						longest path delay (ns)						power consumption (W)					
	[2]	[5]	[3]	[1]	TSA	CA	[2]	[5]	[3]	[1]	TSA	CA	[2]	[5]	[3]	[1]	TSA	CA
ckt1	3.046	3.109	3.784	3.240	3.250	3.133	4.935	4.796	5.128	5.067	4.786	4.871	0.729	0.734	0.776	0.736	0.736	0.732
ckt2	7.740	8.780	14.924	8.349	7.892	8.314	13.679	15.004	15.231	14.416	13.588	14.785	2.463	2.548	2.564	2.521	2.487	2.523
ckt3	9.347	10.544	16.028	10.706	10.355	10.261	6.567	6.797	7.865	7.276	6.530	6.906	2.290	2.331	2.351	2.318	2.306	2.321
ckt4	12.480	13.902	19.871	15.234	14.901	14.545	12.518	12.695	16.158	13.609	13.695	13.113	2.640	2.671	2.737	2.682	2.672	2.675
ckt5	18.869	21.482	27.649	20.125	19.845	19.994	13.931	16.427	15.649	13.674	13.799	14.664	3.127	3.194	3.255	3.166	3.130	3.156
ratio	**1.005**	1.112	1.595	1.120	1.093	1.090	**1.007**	1.066	1.160	1.058	1.015	1.051	**1.001**	1.019	1.043	1.015	1.009	1.014

Ckt.	max junc.-to-amb. temp, $T_{ja,max}$ (°C)						temp difference, $T_{ja,max} - T_{ja,min}$ (°C)						average temp, $T_{ja,ave}$ (°C)					
	[2]	[5]	[3]	[1]	TSA	CA	[2]	[5]	[3]	[1]	TSA	CA	[2]	[5]	[3]	[1]	TSA	CA
ckt1	72.48	73.12	82.86	70.69	70.85	70.41	16.29	14.94	28.12	14.69	15.55	14.16	63.80	63.70	69.52	63.32	63.27	63.35
ckt2	91.70	74.21	101.00	76.89	100.19	73.05	46.96	15.15	51.16	22.39	53.87	17.15	64.81	66.84	69.36	66.07	65.14	66.14
ckt3	77.74	64.39	69.80	66.34	72.41	65.60	39.89	19.68	28.69	23.82	33.65	22.97	55.41	55.49	55.97	54.53	54.14	55.08
ckt4	73.79	62.43	80.11	60.14	65.50	59.31	35.46	16.69	39.76	15.87	21.83	14.27	55.07	54.35	60.42	53.91	53.63	53.85
ckt5	74.86	79.22	76.25	61.95	64.45	61.60	38.08	36.39	38.02	23.77	33.07	24.53	54.51	55.08	57.97	53.22	51.91	52.90
ratio	1.056	0.964	1.105	0.909	0.997	**0.895**	1.235	0.744	1.360	0.719	1.042	**0.673**	0.994	0.999	1.059	0.984	**0.973**	0.984

based force, and found large increase in wirelength without much additional temperature improvement. Moving cells out of a high-temperature area on a die may not reduce temperature if the high temperature is a result from thermal coupling with other dies. Also, without considering TSV thermal properties during thermal analysis, the thermal gradient does not capture the impact of TSVs on temperature accurately, thereby misguiding the placement. Second, we see that [5] obtains wirelength and delay results that are significantly worse than non-uniform TSV placer. This is mainly because it moves logic cells based only on power density. However, this move helps reduce maximum temperature and thermal variation inside the 3D chip stack significantly. Although it attempts to spread power over placement area, we observe that TSVs obstruct this effort frequently.

Third, the routed wirelength and delay of results from [3] are worse than all other placers. The main reason is that [3] does not consider TSV area during placement. Thus, the TSVs inserted during routing affects the placement quality significantly. The maximum temperature, thermal variation, and average temperature are also worse than uniform TSV placer. The router tends to insert TSVs in the middle of the die to minimize wirelength, leaving low thermal conductivity at chip corners, thus high temperature. Fourth, although the wirelength of result from [1] is worse than other placers, temperature improvement is among the best. Because the algorithm considers the impact of TSV on chip area and temperature, it utilizes TSVs more effectively to help improve temperature results.

Fifth, we observe that our TSV spread and alignment method (TSA) achieves comparable delay and power results at the cost of wirelength degradation compared with non-uniform placer. In case of temperature, TSA obtains better average temperature than uniform TSV and comparable maximum temperature and temperature difference. But, the wirelength of TSA method is significantly better than that of uniform TSV placer. These results show that our TSA method is better in reducing wirelength while optimizing temperature compared with uniform TSV placer.

Lastly, our thermal coupling-aware placement (CA) achieves the best temperature results among all placers [2, 5, 3, 1], including uniform TSV placer [4]. In particular, our CA method outperforms uniform TSV placer by 10% and 33% in terms of maximum temperature and temperature difference. CA obtains 9% worse wirelength and 5% worse delay results compared with non-uniform TSV placer, but CA is among the best in terms of wirelength and delay among other placers [2, 5, 3, 1]. The power overhead is negligible. The TSVs in the placement by our CA method are not spread as evenly as our TSA placer and uniform TSV placer, but they are spread only sufficiently to help remove heat from the dies in the stack while maintaining high-quality wirelength. In addition, we observe that high-power logic cells are also placed effectively to dissipate heat using the nearby TSVs that are vertically aligned all the way to the heatsink.

The runtime of all placement algorithms is roughly in the same magnitude. Except for our TSA method, all other placement algorithms require power simulation (and thermal simulation in the case of [2]), resulting in larger runtime than [4].

5. CONCLUSIONS

In this paper we showed that temperature-aware placers must consider TSV thermal properties and die-to-die thermal coupling during placement. We presented two temperature-aware placement algorithms for 3D ICs. TSVs are spread and aligned in the first algorithm. In the second algorithm, logic cells are moved based on the thermal conductivity to the heatsink, and TSVs are moved based on the power density of the neighboring dies. Experimental results show that our placers achieve the best temperature results among all placers used in our comparison.

6. REFERENCES

[1] J. Cong, G. Luo, and Y. Shi. Thermal-aware cell and through-silicon-via co-placement for 3D ICs. In *Proc. ACM Design Automation Conf.*, pages 670–675, San Diego, CA, Jun. 5–9 2011.

[2] B. Goplen and S. Sapatnekar. Efficient thermal placement of standard cells in 3D ICs using a force directed approach. In *Proc. IEEE Int. Conf. on Computer-Aided Design*, pages 86–89, San Jose, CA, Nov. 9–13 2003.

[3] B. Goplen and S. Sapatnekar. Placement of 3D ICs with thermal and interlayer via considerations. In *Proc. ACM Design Automation Conf.*, pages 626–631, San Diego, CA, June 4–8 2007.

[4] D. H. Kim, K. Athikulwongse, and S. K. Lim. A study of through-silicon-via impact on the 3D stacked IC layout. In *Proc. IEEE Int. Conf. on Computer-Aided Design*, pages 674–680, San Jose, CA, Nov. 2–5 2009.

[5] B. Obermeier and F. M. Johannes. Temperature-aware global placement. In *Proc. Asia and South Pacific Design Automation Conf.*, pages 143–148, Yokohama, Japan, Jan. 27–30 2004.

[6] M. Pathak, Y.-J. Lee, T. Moon, and S. K. Lim. Through-silicon-via management during 3D physical design: When to add and how many? In *Proc. IEEE Int. Conf. on Computer-Aided Design*, pages 387–394, San Jose, CA, Nov. 7–11 2010.

[7] P. Spindler, U. Schlichtmann, and F. M. Johannes. Kraftwerk2–A fast force-directed quadratic placement approach using an accurate net model. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 27(8):1398–1411, Aug. 2008.