

Co-Design of Signal, Power, and Thermal Distribution Networks for 3D ICs

Young-Joon Lee¹⁾, Yoon Jo Kim²⁾, Gang Huang³⁾,
Muhannad Bakir¹⁾, Yogendra Joshi²⁾, Andrei Fedorov²⁾, and Sung Kyu Lim¹⁾
¹⁾Electrical And Computer Engineering, Georgia Institute of Technology,
²⁾Mechanical Engineering, Georgia Institute of Technology, ³⁾Intel Corporation
email: yjlee@ece.gatech.edu

Abstract—Heat removal and power delivery are two major reliability concerns in the 3D stacked IC technology. Liquid cooling based on micro-fluidic channels is proposed as a viable solution to dramatically reduce the operating temperature of 3D ICs. In addition, designers use a highly complex hierarchical power distribution network in conjunction with decoupling capacitors to deliver currents to all parts of the 3D IC while suppressing the power supply noise to an acceptable level. These so called silicon ancillary technologies, however, pose major challenges to routing completion and congestion. These thermal and power/ground interconnects together with those used for signal delivery compete with one another for routing resources including various types of Through-Silicon-Vias (TSVs). This paper presents the work on routing with these interconnects in 3D: signal, power, and thermal networks. We demonstrate how to consider various physical, electrical, and thermo-mechanical requirements of these interconnects to successfully complete routing while addressing various reliability concerns.

I. INTRODUCTION

Historically, advances in the field of packaging and system integration have not progressed at the same rate as ICs. In fact, today's silicon ancillary technologies have truly become a limiter to the performance gains possible from advances in semiconductor manufacturing, especially due to cooling, power delivery, and signaling [1] [2]. Today, it is widely accepted that three-dimensional (3D) system integration is a key enabling technology and has recently gained significant momentum in the semiconductor industry. Three-dimensional integration may be used either to partition a single chip into multiple dies to reduce on-chip global interconnects length and/or used to stack chips that are homogenous or heterogeneous.

There are a number of interconnect challenges that need to be addressed to enable stacking of high-performance dies, especially in the area of cooling. When two $100W/cm^2$ microprocessors are stacked together, for example, the net power density becomes $200W/cm^2$, which is beyond the heat removal limits of today's air cooled heat sinks. This issue has recently been addressed with a novel 3D integration technology that features the use of a microchannel heat sink in each die of the 3D system and the use of wafer-level batch fabricated electrical and fluidic chip input/output (I/O) interconnects [3], [4].

Another major challenge is power delivery. As the fabrication technology advances, power consumption of the chip increases. According to ITRS projection, the power consumed by a single chip will reach $200W$ in a few years. Even in the packages of today's industry designs, more than half of the IO pins (or C4 bumps) are dedicated for power and ground connections. As multiples chips are stacked together into a smaller footprint, delivering current to all parts of the 3D stack while meeting the noise constraints becomes highly challenging. This is mainly because the number of through-silicon-vias (TSVs) available for signal nets and power/ground (P/G) nets is limited, causing routing congestion if many 3D connections are desired. This issue is further exacerbated if micro-fluidic channels

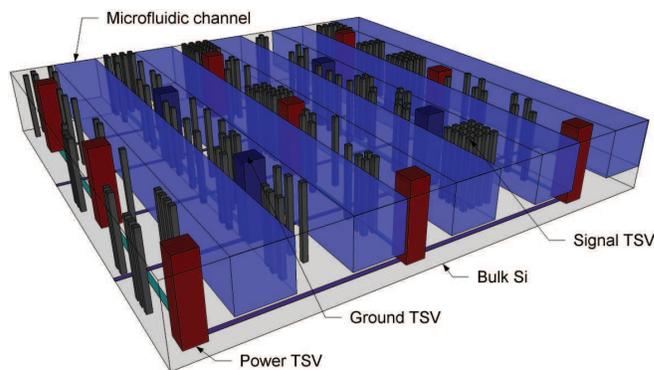


Fig. 1. A 3D die structure with MFCs, P/G TSVs, and signal TSVs. Transistors and signal wires are not shown for simplicity.

(MFCs) are used for liquid cooling.¹ Figure 1 shows a possible configuration of MFCs and signal and P/G TSVs, all competing for layout resource.

Our main contributions are as follows:

- We present the work on routing with three kinds of interconnects—signal, power, thermal—in 3D ICs. We demonstrate how to consider various physical, electrical, and thermo-mechanical requirements of these interconnects to successfully complete routing while addressing thermal and noise concerns.
- We present a compact physical model to analyze the thermal performance of MFC-based liquid cooling. We also discuss the routing challenges posed by this thermal interconnect and propose a way to optimize the geometries of the interconnects.
- We demonstrate the effectiveness of our approach using large-scale gate-level benchmarks that contain up to 1.2 million gates. We report routing congestion, thermal distribution, and power supply noise based on full-layouts of 4-die 3D IC.

The remainder of the paper is organized as follows. We present micro-fluidic cooling and thermal analysis in Section II. Section III presents our P/G network design and noise analysis. Section IV presents our routing approach for the three kinds of interconnects. Experimental results are presented in Section V, and we conclude in Section VI.

¹These silicon ancillary technologies continue to advance, and the size of the related interconnects and TSVs continues to scale down. Our study will be helpful in evaluating these advances in the context of full layout and routing environment.

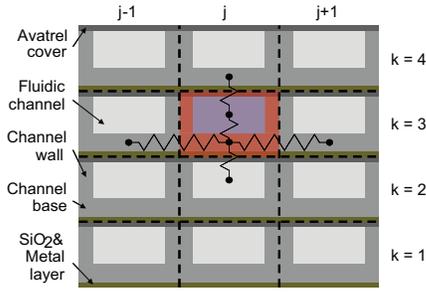


Fig. 2. Side view of the thermal grid structure used for 3D stacked IC with MFCs.

II. THERMO-FLUIDIC INTERCONNECT FOR 3D ICs

A. Fabrication of MFCs and TSVs

Unlike air-cooled heat sinks, liquid cooling using microchannels offers a larger heat transfer coefficient and chip-scale cooling solution. In [3], [4], both electrical and fluidic TSVs and I/Os were demonstrated. The electrical interconnects are used for power delivery and signaling between dies, and the fluidic interconnects are used to deliver a coolant to each microchannel heat sink in the 3D stack and thus enable the rejection of heat from each stratum in the 3D stack. The thermal resistance of the microchannel heat sink for single chip was previously measured [2]. When de-ionized water was used as coolant, the junction-to-ambient thermal resistance of the heat sink was $0.24^{\circ}\text{C}/\text{W}$ at a flow-rate of about $65\text{mL}/\text{min}$ without TSVs (impact of copper TSVs on thermal conductivity of the silicon microchannel wall is negligible), which is significantly better than current state-of-the-art air cooled heat sinks [2].

B. Thermal Analysis

Three-dimensionally stacked ICs bring several challenges in thermal management. By stacking layers, the heat dissipation per unit volume and per unit horizontal footprint area are significantly increased. Also, the interior layers of the 3D structure are thermally detached from the heat sink. Heat transfer is further restricted by the low thermal conductivity bonding interfaces and thermal obstacles in multiple IC layers.

Microchannels capped with the thin polymer (Avatrel 2000 P) coating ($\sim 30\mu\text{m}$) were tested up to 2.5atm pressure with no leakage observed during continuous operation [5]. In the following sections we analyze thermal performance of micro-channel cooling for 3D ICs using numerical simulations.

A three-dimensional thermal model of Koo et. al. [6] is modified to consider the lateral temperature and fluid flow rate distribution caused by non-uniform power/heat flux distribution. Figure 2 shows the side view of the 3D stacked IC with embedded MFCs. It is assumed that the temperatures of the fluid and the solid domains are different but uniform at each cross section within each control volume. Thermal and fluid flow in MFCs are described by the following energy and momentum conservation equations:

$$\dot{m} \frac{di}{dz} = \eta_0 h_c \tilde{P} (T_{w,k} - T_{f,k}) + h_c w (T_{w,k+1} - T_{f,k}) \quad (1)$$

$$-\frac{dP}{dz} = \frac{2fG^2}{d_h \rho} \quad (2)$$

$$\frac{\partial}{\partial x} (k \frac{\partial T_w}{\partial x}) + \frac{\partial}{\partial y} (k \frac{\partial T_w}{\partial y}) + \frac{\partial}{\partial z} (k \frac{\partial T_w}{\partial z}) + \dot{q}_g + \dot{q}_c = 0 \quad (3)$$

T_w and T_f represent the temperatures of solid and fluid, respectively, \dot{m} , i and h_c are mass flow rate, enthalpy, and convective heat

transfer coefficient, respectively. For each MFC, heat is directly supplied only to the channel base, and the channel wall is analyzed as a fin attached to the base (η_0 is the overall surface efficiency for heat transfer, including an array of fins and the base surface). Microchannel geometry is described by the channel perimeter \tilde{P} and the width w . Equation (1) represents the fluid enthalpy change due to the convective heat transfer owing to the temperature difference between the solid and fluid, as well as fluid convective motion. The pressure drop along the MFC is obtained by solving the fluid momentum balance Equation (2), wherein P , G and ρ are pressure, mass flux and density of the fluid, respectively, f is the fluid friction factor and d_h is the hydraulic diameter of a MFC. Equation (3) is the three-dimensional thermal transport equation for the solid. It has two source/sink terms owing to heat generated from the active and oxide-metal layers and convective heat transfer to the fluid (k denotes the thermal conductivity of solid).

Deionized water is considered as a representative working fluid. The governing Equations (1), (2), and (3) are integrated over a control volume and then discretized using the upwind scheme [7]. The resulting system of linear algebraic equations is simultaneously iteratively solved using successive under-relaxation (SUR) method. More details on the experimental settings and results are provided in Section V-B.

C. Routing Requirements for Thermo-fluidic Network

The on-chip thermo-fluidic network is composed of fluidic TSVs and MFCs. In this work, we assume that all the fluidic TSVs are located outside the main region where all the gates and metal wirings are distributed. Thus, only MFCs are considered for routing requirement analysis. MFCs do not affect routing capacity on metal layers. However, MFCs obstruct TSV connections. Considering the significant size of MFCs, they decrease the routing capacity of TSVs quite considerably.

The following geometries related to MFCs have impacts on thermal and routability objectives:

- MFC depth - By increasing the MFC depth, mass flow rate and thus cooling capability can be improved. However, it also increases die thickness, and with the fixed aspect ratio of TSVs, the diameter of signal TSVs gets increased proportionally. Thus, it decreases routing capacity between dies, aggravating signal net routability.
- MFC width - Increasing the MFC width also improves mass flow rate and cooling capability, but the space available for signal TSVs will be decreased.
- MFC pitch - By increasing the MFC pitch, total thermal contact area between the channel wall and the working fluid will be reduced and thus the cooling performance is degraded, but the routability of TSVs gets improved.

III. POWER DELIVERY NETWORK FOR 3D ICs

A. Noise Analysis

In a 3D stack, global power distribution networks on each die are distributed using grids made of orthogonal interconnects on the top wiring levels. Power is fed from the package through P/G I/O bumps distributed over the bottom-most die and travels to the upper dice using TSVs and solders [8]. High performance systems require dense P/G grids for both power distribution and signal current return purposes.

For a 3D chip stack structure with footprint size of 1cm^2 , we may have thousands of P/G I/Os for each die and millions of wire segments on the P/G grids in each die. Our power noise analyzer

is based on modified nodal analysis (MNA) [9]. We use domain decomposition (DD) [10] to increase the maximum circuit size the analyzer can handle. The DD technique basically decomposes the circuit into several parts and uses a mathematical technique to reduce the time needed for matrix inversion.

B. Routing Requirements for Power Delivery Network

In our work, P/G TSVs are placed regularly in a mesh structure with a predefined pitch ($p_{pg.tsv} = 400\mu m$). The width of a P/G tile is a half of the power TSV pitch ($w_{pg.tile} = p_{pg.tsv}/2 = 200\mu m$) and contains one quarter of power TSV and one quarter of ground TSV. The total number of P/G TSVs in a die is:

$$N_{pg.tsv} = \left\lfloor \frac{w_{chip}}{w_{pg.tile}} \right\rfloor \times \left\lfloor \frac{l_{chip}}{w_{pg.tile}} \right\rfloor \times (1/4 + 1/4)$$

where w_{chip} and l_{chip} are the width and the length of the chip.

P/G nets are routed on metal layer 7 and 8. The following three-levels of wiring hierarchy is used:

- Thick wires have $10\mu m$ width and runs between P/G TSVs.
- Between the thick wires, thin wires of $1\mu m$ width and $4.64\mu m$ pitch are placed.
- Between two thin wires, up to six signal wires can be routed.

Thus, the area ratio of signal wire on metal layer 7 and 8 is:

$$R_{sig}(M78) = \frac{6}{4.64/0.56} = 0.724$$

In our 3D technology, two TSVs from two adjacent dies are connected using the metal layers and vias. That is, two stacked signal TSVs are not directly connected. On the other hand, P/G TSVs pierce through all 4 dies for efficient power delivery (see Figure 3). Thus, no cell can be placed at P/G TSV locations. Considering the size of P/G TSVs, this area is not negligible. The total area required for P/G TSVs can be calculated by multiplying the number of P/G TSVs by the P/G TSV area:

$$S_{tot.pg.tsv} = N_{pg.tsv} \times S_{pg.tsv}$$

where $S_{pg.tsv}$ is the area of a P/G TSV. Based on our structural assumptions, P/G TSVs occupy around 2% of the chip area. In the routing tile with some part of a P/G TSV, routing capacity is decreased by a large amount.

IV. SIGNAL INTERCONNECT FOR 3D ICs

A. Overview of 3D Physical Design

Our routing package is divided into several steps. After reading in the input circuit, the partitioning stage starts. In this stage, the input circuit is divided into four dies. Since the total number of signal TSVs is determined at this stage, we minimize the cutsize by utilizing a mincut algorithm [11]. In the placement stage, we perform global placement onto $n_p \times n_p$ grid. Cell occupancy ratio (COR) of a placement tile at (x, y, z) is defined as:

$$COR(x, y, z) = \frac{\sum_{\forall cell \in r_{tile}(x,y,z)} S_{cell}}{S_{r_{tile}}}$$

where S_{cell} is the area of a cell in the routing tile $r_{tile}(x, y, z)$, and $S_{r_{tile}}$ is the area of a routing tile. We perform congestion-driven placement based on simulated annealing technique with a predefined target $COR (= tCOR)$ to distribute gates evenly.

Next, we perform global routing onto $n_r \times n_r$ grid. In our routing strategy, the signal, thermal, and power nets are routed sequentially. We first route the signal nets, followed by MFCs, P/G TSVs and P/G wires. Lastly, we rip up and reroute signal nets with routing capacity

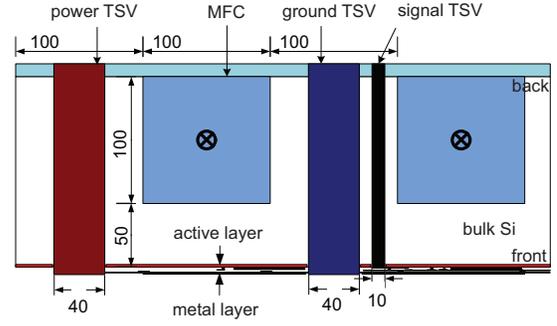


Fig. 3. Side view of a die layer in a stacked chip. The die is flipped over and the active layer is facing down. Shapes are drawn to scale. Unit is μm .

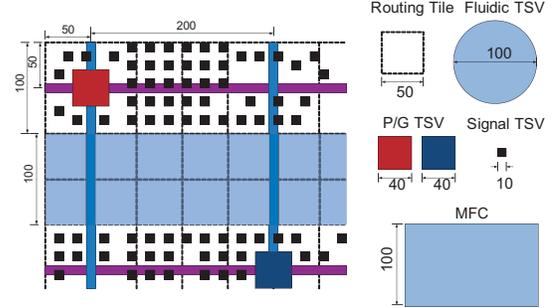


Fig. 4. Top view of the routing tile objects on routing grid. Objects are drawn to scale.

violations. We used the thermal-aware 3D maze router [12]. After the routing is finished, we run power noise and thermal analysis to see if assumed constraints are satisfied. If needed, we may repeat the entire or some parts of the physical design steps with changed configurations.

B. Geometries of Wires and Vias

As for the signal wires, we use the metal interconnect dimensions similar to the ones in Intel's 45nm technology [13]. The TSV formation approach was assumed to be via-first. TSV aspect ratio was assumed to be 15:1.

Figure 3 shows the side view of a die. The diameter of signal TSVs is set to the minimum size to accommodate as many connections as possible. In contrast, the diameter of P/G TSVs is $40\mu m$, which is comparable to an existing work [8]. Table II provides more details on related geometries. We fix the width of our routing tile $w_{r.tile} = 50\mu m$. Figure 4 shows the routing tile objects.

C. Routing Capacity Calculation

For each routing tile, there are x-, y-, and z-direction routing capacity values. x- and y-direction capacity represents available routing space on metal layers, while z-direction capacity is for signal TSVs. Basically, x- and y-direction capacity values of a metal layer are calculated from dividing the routing tile size by the pitch of the metal layer. Since the benchmark circuits [14] are based on gate-level netlist, most of metal layer 1 and 2 are used for local wiring. Thus, we assume that only 20% of the routing capacity is available in metal 1-2. Metal 3-6 are dedicated to signal routing. In metal 7 and 8, we decrease number of routing capacity values due to the P/G nets. The capacity values based on multiple metal layer stack are added together for each tile. If the tile is pre-occupied with P/G

TABLE I

ISPD 2006 BENCHMARK CIRCUITS. WE REPORT THE TOTAL NUMBER OF CELLS AND NETS, WIDTH (μm) AND FOOTPRINT AREA (mm^2) OF THE 3D STACK. WE ALSO REPORT THE DIMENSIONS OF THE ROUTING, P/G, AND THERMAL GRIDS BASED ON THE CHIP AREA.

	<i>adaptec1</i>	<i>newblue1</i>	<i>newblue3</i>	<i>adaptec5</i>	<i>newblue5</i>
# cells	211,447	330,474	494,011	843,128	1,233,058
# nets	221,142	338,901	552,199	867,798	1,284,251
Width	6,000	7,200	10,000	12,000	15,000
Area	36	51.84	100	144	225
R-grid	120x120x4	144x144x4	200x200x4	240x240x4	300x300x4
P-grid	30x30x4	36x36x4	50x50x4	60x60x4	75x75x4
T-grid	30x80x4	36x80x4	50x80x4	60x80x4	75x80x4

TSVs, we decrease the capacity accordingly:

$$Cap_x(x, y, z) = \left(\frac{w_{r.tile}}{w_{M1}} \times 0.2 + \frac{w_{r.tile}}{w_{M3}} + \frac{w_{r.tile}}{w_{M5}} + \frac{w_{r.tile}}{w_{M7}} \times R_{sig}(M78) \right) \times (1 - R_{pg.tsv})$$

$$Cap_y(x, y, z) = \left(\frac{w_{r.tile}}{w_{M2}} \times 0.2 + \frac{w_{r.tile}}{w_{M4}} + \frac{w_{r.tile}}{w_{M6}} + \frac{w_{r.tile}}{w_{M8}} \times R_{sig}(M78) \right) \times (1 - R_{pg.tsv})$$

where $R_{pg.tsv}$ is the ratio of the length by which P/G TSVs occupy in the tile.

For z-direction capacity, we calculate the remaining surface area of each routing tile. Starting from the routing tile area, we extract the placed cell area and the P/G TSV area. Since we place P/G TSVs at the center of four routing tiles, only one quarter of P/G TSV is included in a routing tile. Then, we divide the resulting area by the area calculated with the signal TSV pitch.

$$Cap_z(x, y, z) = \frac{S_{r.tile} - S_{placed.cells} - R_{pg.tsv} \cdot S_{pg.tsv}}{p_{sig.tsv}^2}$$

In case of x-direction interconnects, P/G TSVs occupy 2.5% of the available routing area. Thus, 97.5% of x-direction routing resource is available for signal net routing. The same is true for y-direction. In case of z-direction, the power TSVs and the thermal interconnects (=MFCs) occupy 2% and 50% of the surface area, respectively. Considering the placed cell area with $tCOR = 0.25$, about 36% of the total surface area is available for signal TSVs.

V. EXPERIMENTAL RESULTS

We implemented our design package in C++/STL and MATLAB. The simulations were done on a 64-bit Linux server with two quadcore Intel Xeon 2.5GHz CPUs and 16GB main memory. The circuits used for experiments are from the ISPD 2006 Placement Contest benchmark [14] that range from 200K to 1.2M gates as shown in Table I. The technology and setting parameters are shown in Table II. The target cell occupancy ratio is lower than that of 2D cases [14], because of MFCs and TSVs.

A. Routability and Congestion Analysis

Table III shows the signal, P/G, and thermal interconnect routing results, where we report the average and maximum utilization of the routing tiles in x-, y-, and z-directions on die 1. Only die 1 result is shown because it is usually the most congested layer. We also report the number of signal TSVs for all dies. The run time of signal routing stage was about 50 hours with *newblue5*.

Looking at the average routing tile usage, we note that the z-direction usage is higher than x- and y-direction ones. This is expected because the size of signal TSVs is significantly larger than

TABLE II

VARIOUS TECHNOLOGY AND SETTING PARAMETERS.

Item	Value
Number of dies	4
Bonding type	face-to-back
Die thickness (μm)	150
Bonding layer thickness (μm)	10
TSV aspect ratio	15:1
Routing grid size (μm)	50
Signal TSV diameter (μm)	10
Signal TSV minimum pitch (μm)	20
P/G TSV diameter (μm)	40
P/G TSV pitch (μm)	400
P/G grid size (μm)	200
MFC depth (μm)	100
MFC width (μm)	100
MFC pitch (μm)	200
MFC occupancy ratio	0.5
Target cell occupancy ratio	0.25

TABLE III

ROUTING RESULTS, WHERE WE REPORT THE AVERAGE AND MAXIMUM UTILIZATION OF THE ROUTING TILES IN X-, Y-, AND Z-DIRECTIONS ON DIE 1. WE ALSO REPORT THE NUMBER OF SIGNAL TSVs FOR ALL DIES.

	<i>adaptec1</i>	<i>newblue1</i>	<i>newblue3</i>	<i>adaptec5</i>	<i>newblue5</i>
Average routing tile usage					
x, die 1	0.153	0.121	0.202	0.273	0.290
y, die 1	0.125	0.106	0.159	0.219	0.235
z, die 1:2	0.343	0.417	0.273	0.249	0.351
Maximum routing tile usage					
x, die 1	1.000	1.000	1.000	1.000	1.000
y, die 1	0.992	0.959	1.000	1.000	1.000
z, die 1:2	1.000	1.000	1.000	1.000	1.000
TSV 0:1	14,057	28,707	36,052	102,771	154,544
TSV 1:2	19,356	38,686	40,692	59,337	139,515
TSV 2:3	6,557	33,659	27,181	13,184	127,743

that of signal wires, and some tiles have MFCs. We also note that the difference of average usage between x-/y- and z-direction gets smaller with bigger circuits. This is because the ratio of x-/y- and z-direction connections is dependent on the chip size and the number of dies. Maximum routing usage is almost 1 for all cases. And, the variation of number of TSVs among layers is significant for *adaptec5*. Considering the z-direction congestion, even distribution is preferred, but this is not optimized in this study. Figure 5 shows x-, y-, and z-direction routing tile utilization on die 1 for *newblue5*. We note that the congestion of z-direction is more severe than that of x- and y-directions.

B. Thermal Analysis

The benchmark circuits do not have gate-level switching activities nor current demand information. Thus, we generated a realistic power map. Each hot spot is rectangular with a randomly chosen dimensions ranging 150 – 300 μm and a power density of 400 – 800W/cm².

Pressure drop between inlet and outlet was constrained to 140kPa for all MFCs. Inlet fluid temperature was assumed to be 20°C. The thermal analyzer was written in MATLAB, and the run time was about 3 minutes with *newblue5* circuit.

Figure 6 shows that the die temperature is maintained well below 85°C. On each die, the temperature on the right side is higher than the one on the left side. This is because the fluid enters on the left side and as it goes to the right side the fluid temperature gets higher.

Table IV shows the summary of thermal analysis results. For all the cases, the coolant temperature never exceeded 30°C. Maximum

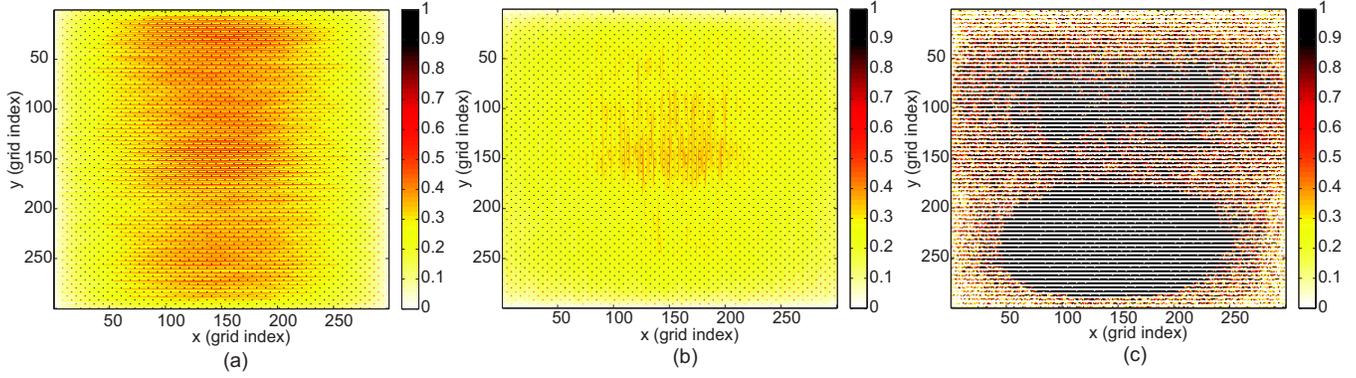


Fig. 5. Routing usage of *newblue5*: (a) x-direction, (b) y-direction, (c) z-direction (=TSV between die 1 and die 2), where the horizontal white lines denote MFCs.

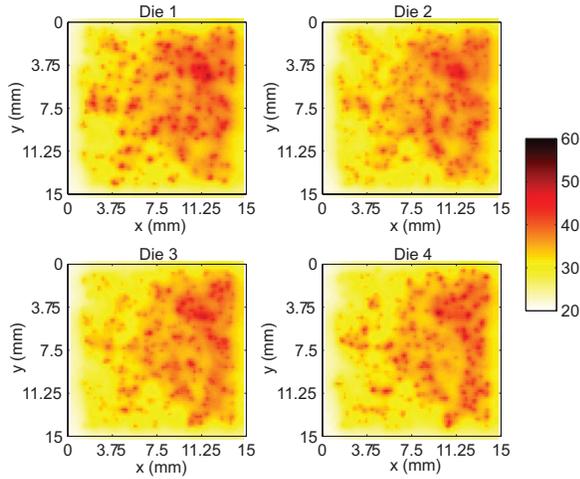


Fig. 6. Silicon wall temperature of *newblue5*. Unit is $^{\circ}C$.

TABLE IV
SUMMARY OF THERMAL ANALYSIS RESULTS. POWER AND POWER DENSITY VALUES ARE REPORTED IN W AND W/cm^2 RESPECTIVELY, AND TEMPERATURES ARE IN $^{\circ}C$.

	<i>adaptec1</i>	<i>newblue1</i>	<i>newblue3</i>	<i>adaptec5</i>	<i>newblue5</i>
Ave die power	14.35	21.40	33.92	54.39	80.79
Ave pwr density	39.87	41.28	33.92	37.77	35.91
Ave fluid temp.	20.65	20.97	21.55	22.52	23.52
Max fluid temp.	21.82	22.71	24.04	26.47	29.68
Ave wall temp.	30.62	31.15	29.94	31.82	32.34
Max wall temp.	43.93	49.48	45.86	49.10	50.91
σ wall temp.	2.963	3.753	3.247	3.953	4.826
Pump power	1.500	1.496	1.494	1.518	1.541

silicon wall temperature was kept under $51^{\circ}C$ for all dies in all cases. This means that microfluidic cooling scheme had enough cooling capability for thermal management of the 3D circuits. For *newblue5*, the maximum temperature difference between the die and the working fluid is observed to be around $20^{\circ}C$, which means that the convective heat transfer coefficients was higher than estimated and/or thermal diffusion due to conduction played a significant role. The temperature difference was supposed to be around $36^{\circ}C$ considering the die power density (about $36W/cm^2$) and the expected heat transfer coefficient (around $10,000W/m^2K$). Standard deviations of wall temperatures were around $3 - 5^{\circ}C$.

TABLE V
VARIOUS MFC GEOMETRY SETTINGS.

	<i>base</i>	<i>d50</i>	<i>w50</i>	<i>p400</i>
MFC depth (μm)	100	50	100	100
MFC width (μm)	100	100	50	100
MFC pitch (μm)	200	200	200	400
MFC occupancy ratio	0.5	0.5	0.25	0.25
Die thickness (μm)	150	90	150	150
Signal TSV diameter (μm)	10	6	10	10

TABLE VI
IMPACT OF MFC GEOMETRIES ON ROUTING CONGESTION.

	<i>base</i>	<i>d50</i>	<i>w50</i>	<i>p400</i>
Average routing file usage				
x, die 1	0.153	0.152	0.151	0.151
y, die 1	0.125	0.124	0.124	0.124
z, die 1:2	0.343	0.166	0.377	0.375
Maximum routing file usage				
x, die 1	1.000	1.000	1.000	1.000
y, die 1	0.992	0.984	0.927	1.000
z, die 1:2	1.000	1.000	1.000	1.000
TSV 0:1	14,057	14,922	15,012	14,971
TSV 1:2	19,356	21,052	21,152	21,033
TSV 2:3	6,557	6,898	6,937	6,936

In our system, MFCs are the biggest objects that block P/G and signal nets and thus attract more attention for routing optimization. To investigate the impacts of MFC configurations on design quality, we conducted the experiments with three varied geometric parameters of MFCs shown in Table V. In setting *d50*, the depth of MFC was halved and the signal TSV dimension was decreased accordingly, in *w50*, the width of MFC was halved, and in *p400*, the pitch of MFC was doubled. Circuit *adaptec1* was used for the demonstration.

Table VI shows the routing results with the MFC variations. Compared to the baseline case (*base*), x- and y-direction utilizations of variants are almost the same. Since z-direction capacity has been increased for all variants, the router uses more TSVs to decrease wirelength. In *d50* case, the average usage of z-direction is lower than that of other cases, because the reduced signal TSV dimension increased the z-direction capacity.

Table VII shows the thermal analysis results with MFC variations. Compared to *base*, in *d50* case, the fluid temperature becomes higher because the mass flow rate decreases with smaller MFC cross-sectional area. In *w50* case, the maximum wall temperature is lower than that of *d50* case, because of increased cooling efficiency due to

TABLE VII
IMPACT OF MFC GEOMETRIES ON THERMAL METRICS.

	<i>base</i>	<i>d50</i>	<i>w50</i>	<i>p400</i>
Ave fluid temp. ($^{\circ}C$)	20.65	23.05	23.08	20.51
Max fluid temp. ($^{\circ}C$)	21.82	28.33	27.81	22.56
Ave wall temp. ($^{\circ}C$)	30.62	31.08	30.71	36.69
Max wall temp. ($^{\circ}C$)	43.93	46.44	43.32	49.07
σ wall temp. ($^{\circ}C$)	2.963	3.749	3.272	3.486
Pump power (W)	1.500	0.322	0.322	0.755

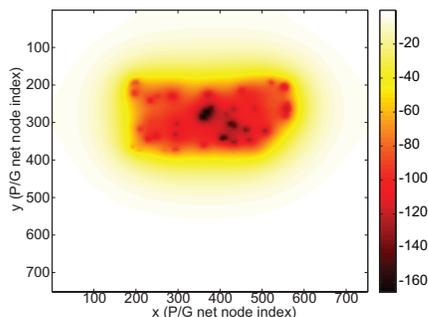


Fig. 7. Peak power noise on the top die of *newblue5*. Unit is *mV*.

the thicker wall between two adjacent MFCs. And in *p400* case, although fluid temperature is not much higher than that of *base*, the maximum silicon wall temperature increases due to the reduced contact area for heat transfer between the MFC wall and the fluid.

C. Power Noise Analysis

For the purpose of power noise analysis, P/G grid was formed by superposing *RLC* mesh structures of a half P/G TSV pitch ($p_{pg,tsv}/2$) on each layer. The power consumption for each grid location was modeled as a current source. We used the power consumption values of routing tiles in thermal analysis stage to determine the values of current sources. The gate oxide thickness was set to $1nm$ for decoupling capacitor (decap) size calculation. And the inductance and the resistance of package pins was assumed to be $0.3nH$ and $3m\Omega$, respectively.

In order to determine decap area ratio (R_{decap}) at each grid point, we calculated the unused silicon area, and assumed that 80% of the unused area is used for decap. For *newblue5*, the average of R_{decap} for the entire stack was found to be around 50.6%.

To model simultaneous switching noise, it was assumed that 1/8 area of each die is turned on together with the current profile of $5ns$ rise time. The supply voltage was assumed to be $1V$. After simulation, we gathered the peak power noise voltage for each grid point. The runtime of a power noise simulation was about 30 minutes.

Figure 7 shows the power noise level at the grid with maximum peak noise on the top die of *newblue5*. Only the top die result is shown because it is the farthest one from power supply and thus tends to show the highest power noise level. Some parts have higher power noise level due to higher power demand.

Table VIII shows the summary of power noise simulation results. Peak power noise values are quite high for all circuits. And it was found that the maximum peak power noise location has correlation with the maximum power consumption location. To decrease the maximum noise level, we may need to put more decaps, widen the P/G net wires, or carefully place cells to avoid too high power density.

TABLE VIII
SUMMARY OF POWER NOISE SIMULATION RESULTS. WE REPORT AVERAGE AND MAXIMUM PEAK POWER NOISE IN *mV*.

	<i>adaptec1</i>	<i>newblue1</i>	<i>newblue3</i>	<i>adaptec5</i>	<i>newblue5</i>
Average	16.66	18.68	17.70	18.75	17.12
Maximum	95.02	108.64	131.31	166.11	165.87

VI. CONCLUSIONS

In this paper, we presented the work on routing with the signal, power, and thermal interconnects in 3D ICs. We discussed how to consider various physical, electrical, and thermo-mechanical requirements of these interconnects to successfully complete routing while addressing various reliability concerns. Our studies revealed that the liquid cooling based on MFCs is highly effective in removing the hotspots in 3D designs. We also learned that the P/G distribution network for 3D IC requires a high demand on TSVs and interfere with signal net routing. The major signal net routing bottleneck was related to TSVs.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under CAREER Grant No. CCF-0546382 and the Interconnect Focus Center (IFC).

REFERENCES

- [1] G. G. Shahidi, "Evolution of CMOS technology at 32 nm and beyond," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2007, pp. 413–416.
- [2] M. Bakir, B. Dang, and J. Meindl, "Revolutionary nanosilicon ancillary technologies for ultimate-performance gigascale systems," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2007, pp. 421–428.
- [3] D. Sekar, C. King, B. Dang, T. Spencer, H. Thacker, P. Joseph, M. S. Bakir, and J. D. Meindl, "A 3D-IC Technology with Integrated Microchannel Cooling," in *Proc. Int. Interconnect Technol. Conf.*, 2008.
- [4] M. S. Bakir, C. King, D. Sekar, H. Thacker, B. Dang, G. Huang, A. Naeemi, and J. D. Meindl, "3D heterogeneous integrated systems: liquid cooling, power delivery, and implementation," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2008.
- [5] B. Dang, M. S. Bakir, and J. D. Meindl, "Integrated thermal-fluidic I/O interconnect for an on-chip microchannel heat sink," *IEEE Electron Device Letter*, vol. 27(2), pp. 117–119, 2006.
- [6] J.-M. Koo, S. Im, L. Jiang, and K. E. Goodson, "Integrated microchannel cooling for three-dimensional electronic architecture," *J. Heat Transfer*, vol. 127, pp. 49–58, 2005.
- [7] S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*. Washington, DC, Hemisphere Publishing Corp., 1980.
- [8] G. Huang, M. Bakir, A. Naeemi, H. Chen, and J. Meindl, "Power Delivery for 3D Chip Stacks: Physical Modeling and Design Implication," in *Proc. IEEE Electrical Performance of Electronic Packaging*, 2007, pp. 205–208.
- [9] C.-W. Ho, A. E. Ruehli, and P. A. Brennan, "The Modified Nodal Approach to Network Analysis," *IEEE Transactions on Circuits and Systems*, vol. 22, no. 6, pp. 504–509, June 1975.
- [10] Q. Zhou, K. Sun, K. Mohanram, and D. C. Sorensen, "Large power grid analysis using domain decomposition," in *Proc. Design, Automation and Test in Europe*, vol. 1, 2006, pp. 1–6.
- [11] C. Fiduccia and R. Mattheyses, "A Linear Time Heuristic for Improving Network Partitions," in *Proc. ACM Design Automation Conf.*, 1982, pp. 175–181.
- [12] M. Pathak and S. K. Lim, "Thermal-aware Steiner Routing for 3D Stacked ICs," in *Proc. IEEE Int. Conf. on Computer-Aided Design*, 2007, pp. 205–211.
- [13] K. M. et.al., "A 45nm Logic Technology with High-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free Packaging," in *IEEE International Electron Devices Meeting*, 2007, pp. 247–250.
- [14] G.-J. Nam, "ISPD 2006 Placement Contest." [Online]. Available: <http://www.sigda.org/ispd2006/contest.html>