

Full Length Article

3D thermal-aware floorplanner using a MILP approximation[☆]David Cuesta^{*}, José L. Risco-Martin, José L. Ayala

Facultad de Informática, C/ Prof. García Santesmases, S/N. 28040 Madrid, Spain

ARTICLE INFO

Article history:
Available online 22 February 2012

Keywords:
Design
Performance
Real-time and embedded systems
Problem Solving
Control Methods and Search-Heuristic methods

ABSTRACT

One of the most important concerns in 3D technology is heat removal. In this paper we propose a 3D thermal-aware floorplanner. Our contributions include: (1) a novel multi-objective formulation to consider the thermal and performance constraints in the optimization approach; (2) an efficient Mixed Integer Linear Programming (MILP) representation of the floorplanning model; and (3) a smooth integration of the MILP model with an accurate thermal modelling of the architecture. The experimental results for several realistic 3D stacks based on the Niagara system show promising improvements of the main thermal metrics, with a reduced overhead in the wire length of the system.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The exponentially increasing power densities that have been reached in current technologies, the values of leakage currents, the cooling costs and the recent reliability constraints in microprocessor-based systems have motivated the cooling down of the chip temperature to be one of the main concerns in system design.

The operating temperature has a significant impact on microprocessor design. At higher temperatures, transistors work slower due to the degradation of the carrier mobility. The resistivity of the metal interconnects also increases, causing longer delays and, therefore, performance degradation.

Reliability is also strongly related to temperature, and increasing the temperature will exponentially decrease the lifetime of the chip. The time to failure has been shown to be a function of $e^{-Ea/kT}$, where Ea is the activation energy of the failure mechanism being accelerated by the increased temperature, k is the Boltzmann's constant, and T is the absolute temperature.

When the temperature of the chip increases 10°, the life of the component is reduced by a 50%. In order to achieve a maximum reliability in the components of the system, it is important to keep the temperature as cool as possible. However, the absolute temperature of the chip is not the only factor that affects performance;

moreover, the thermal gradients that appear on the chip surface degrade the system reliability by creating dangerous electro-migrations.

In order to maintain the chip temperature under a certain limit, the power density of the hardware modules can be decremented by increasing the chip area. However, this is not admissible in terms of cost, and the problem of meeting all the geometric constraints should be solved.

Another important factor that affects the temperature distribution of the chip is the lateral spreading of heat in silicon. This depends on the placement of the units and their proximity to the chip border, as well as on other cool units that behave as thermal sinks, or hot units that are considered as thermal sources. Thermal-aware floorplanning algorithms are able to even out the temperature of the hardware modules through spreading of the heat dissipation. This aspect of floorplanning is particularly attractive in comparison with static external cooling, that reduces the temperature of the chip surface by a constant factor.

Three-dimensional (3D) multi-processor chips have been proposed as an effective mechanism to significantly improve system performance by reducing interconnect delays and increasing the density of the integrated logic. They also allow the integration of multiple and disparate technologies, such as radio frequency and mixed signal components.

A major concern in 3D architectures is the increased power densities reached as a result from placing a hard computational unit over another in the 3D stack. The thermal conductivity of the dielectric layers inserted between device layers for insulation is very low compared to silicon and metal. Since power densities are already a major concern in 2D architectures, the move to 3D architectures will accentuate the thermal problem. Consequently,

[☆] This work has been supported by Spanish Government grants TIN2008-00508 and MEC Consolider Ingenio CSD00C-07-20811 of the Spanish Council of Science and Technology, and partially supported by the Swiss National Science Foundation (SNF) grant 200021-127282.

^{*} Corresponding author.

E-mail addresses: dcuestag@pdi.ucm.es (D. Cuesta), jlrisko@dacya.ucm.es (J.L. Risco-Martin), jayala@fdi.ucm.es (J.L. Ayala).

it is mandatory to devise efficient 3D floorplanning mechanisms that optimize the thermal profile of these complex 3D multi-processor architectures.

This work continues the initiated in [1] that proposes a set of design rules for the generation of thermal-aware floorplans for the 3D Niagara architecture. This previous work obtained improvements of the thermal metrics with respect to the baseline architecture and compared to traditional thermal-aware floorplanner. The work presented in this paper outperforms this results with a MILP formulation and an efficient solver that manages multiple objectives in the minimization problem.

This paper specifically makes the following contributions:

1. It provides a novel multi-objective formulation of the floorplanning problem in 3D multi-processor architectures with thermal constraints.
2. It performs an efficient resolution of the optimization problem by the use of a Mixed Integer Linear Programming (MILP) framework.
3. It shows good response in terms of the main thermal metrics (mean temperature, peak temperature and thermal gradient) for a real architecture based on Niagara multi-processor, outperforming previous approaches.

The paper is organized as follows, first the design flow of the optimization is explained. Secondly the setup of the system is exposed and finally the results obtained are shown.

2. Related work

The impact of the floorplanning on the thermal distribution of real microprocessor based systems is analyzed in [2], where the placement of components for Alpha and Pentium Pro is evaluated.

Some initial works on thermal aware floorplanning [3] propose a combinatorial optimization problem to model our problem. However, the simplification of the considered floorplan and the lack of a real experimental framework motivated the further research on the area. Thermal placement for standard cell ASICs is a well researched area in the VLSI CAD community, where we can find works as [4].

In the area of floorplanning for microprocessor-based systems, some authors consider the problem at the microarchitectural level [5,6], where it is shown that significant peak temperature reduction can be achieved by managing lateral heat spreading through floorplanning. Other works [7] use genetic algorithms to demonstrate how to decrease the peak temperature while generating floorplans with area comparable to that achieved by traditional techniques. Ref. [8] uses a simulated annealing algorithm and an interconnect model to achieve thermal optimization. These works have a major restriction since they do not consider multiple objective factors in the optimization problem, as opposed to our work. Our floorplanner will optimize jointly both thermal metrics (mean temperature, peak temperature and gradient) with a strong impact on the reliability of the system, and the performance of the system (through the minimization of the wire length delay). Moreover, the thermal models used in these studies do not reflect the complex diffusion processes that exist in current technologies. More recent works [9] have tackled the problem of thermal-aware floorplanning with geometric programming but, in this case, the area of the chip is not considered constant.

Thermal-aware floorplanning for 3D stacked systems has also been investigated. Cong et al. [10] proposed a thermal-driven floorplanning algorithm for 3D ICs, which is a natural extension of his previous work on 2D. In [11], Healy et al. implemented a multi-

objective floorplanning algorithm for 2D and 3D ICs, combining linear programming and simulated annealing.

Our work presents more similarities with Ref. [12] by Hung et al., where they propose a thermal-aware floorplanner for 3D architectures. However, this preliminary work does not consider the multi-objective approach proposed in our work, and does not consider the minimization of those thermal variables with a strong impact on the reliability of the system.

Thus, an efficient model of the optimization problem and an effective solver are required to achieve good tradeoff between thermal optimization and performance constraints. In the case of 3D IC design, incremental optimization is a promising way to handle multi-objective optimization with complicated constraints and facilitate the design reuse technology. Several works concerned with incremental floorplanning for 2D IC design [13–16] have been proposed, but none has been proposed, or has taken thermal-aware 3D IC design into consideration. [17] has recently proposed an incremental MILP algorithm. However the design process could take several iterations, whereas our methodology perform the thermal-aware and total wire length optimization in two steps.

In this paper, we propose a novel algorithm to optimize the 3D layout in order to eliminate the hotspots, reduce the peak temperature and decrease the reliability risks. Given a 3D packing and a chip area, we formulate the thermal-aware and total wire length optimization into two MILP problems. The former is defined in terms of power density, moving hottest blocks until they are as far away as possible from each other. Then, with hottest blocks fixed in space, we perform another optimization trying to move the remaining blocks to reduce total wire length. Experiments results show that we can reduce the maximum on-chip temperature in 80° on average, outperforming previous thermal-aware floorplan designs.

3. Design flow

In order to reduce maximal on-chip temperature as much as possible, we propose a novel thermal-aware incremental optimization flow. To this end, we have developed three algorithms. The first one performs an accurate analysis of the thermal behavior in the 3D IC. The second moves all the blocks until the hottest ones are as far away as possible from each other. The last one, having fixed the hottest, tries to move the remaining blocks while total wire length is minimized.

Fig. 1 shows the design flow of our thermal-aware 3D micro architectural floorplanner. Such flow can be divided in two phases. The first one is the thermal analysis of an initial configuration of the 3D IC. Since we are studying the Niagara system, such initial configuration is available in [18]. If a baseline layout is not available, we can obtain an initial configuration by running the second MILP 1 algorithm proposed in this work. This initial scenario will be the seed for the optimizer. The second phase is the optimization loop (rest of the diagram). Next, we describe these two phases in detail.

3.1. Thermal analysis

As Fig. 1 shows, we first perform a thermal analysis. To this end, we have developed an accurate thermal model, which is briefly described in the following.

3D integration consists on placing different active layers using silicon dioxide and joining them with a glue material. If inter layer communication is required, Trough Silicon Vias (TSVs) allow it. Some of the goals on the design of 3D stacks are to achieve a reduction in area and also to decrease the length of the interconnections, that would be translated into a decrease in the data transfer time




Fig. 1. Iterative flow of our approach.

and the power consumption. The 3D stack is built over an adiabatic PCB surface and then, traditional technological dies composed by silicon dioxide and silicon, are placed one over the others.

The 3D stacks that will be studied in the experimental work are composed of 2, 4 and 5 layers, as will be shown in Section 4.

For the simulation we use a 3D RC thermal model as the one presented in [19]. The stack can be characterized by this model because heat flow through the 3D stack is mainly diffusive.

The thermal modeling of the stack is performed after splitting the chip into small cubic unitary cells as the one in Fig. 2. These cells are characterized by six thermal resistances connecting the cell with its neighbors, representing heat transfer, and a capacitance which represents the heat storage inside the cell. The values of the conductances and the capacitance are calculated using these expressions:

$$G_{top/bottom} = k_{th}(l \cdot w)/(h/2) \quad (1)$$

$$G_{north/south} = k_{th}(l \cdot w)/(h/2) \quad (2)$$

$$G_{east/west} = k_{th}(l \cdot w)/(h/2) \quad (3)$$

$$C_{top} = sc_{th}(l \cdot w \cdot h) \quad (4)$$

where north, south, east and west indicate the direction in which heat is diffused; k_{th} and sc_{th} are the thermal conductivity and specific heat capacity per volume unit of the material, respectively.

The model also considers the heat diffusion to the surrounding environment. It is possible to simulate different chip packages by tuning the resistance and conductance parameters. As the PCB base is considered to be adiabatic, no heat transfer occurs in the bottom direction of the first layer. Also, the vertical diffusion that occurs inside the chip depends on the thermal characteristics of the interface material. The interface material that exists in between two silicon layers, used as a glue, is modeled as an epoxy layer, a pure resistant material. The existence of TSVs is considered in the model, also as a resistance element. The most important thermal properties of the material used in the model are listed in Table 1.

Once the technological values have been calculated, a set of equations for the RC grid is created. After that, an iterative method (Forward Euler) is used to solve it.




Fig. 2. Equivalent RC circuit of a single cell.

Table 1
Thermal properties of materials.

Silicon linear thermal conductivity	295 W/(m K)
Silicon quadratic thermal conductivity	−0.491 W/(m K ²)
Silicon dioxide thermal conductivity	1.38 W/(μm K)
Silicon specific heat	1.628 × 10 ⁶ J/m ³ K
Silicon dioxide specific heat	4.180 × 10 ⁶ J/m ³ K

The active elements in the 3D stack can be considered as heat sources or heat sinks. Processors are considered as strong heat sources because they dissipate power in the die, and this heat is then spread throughout the chip. On the other hand, memories have a lower power activity and they can be considered almost as heat sinks. The floorplanner will try to place both heat sinks and heat sources as close as possible (provided the routing and performance constraints) to balance the thermal profile.

Once the previous model has been applied to the 3D IC, we obtain mean and peak temperatures, as well as the thermal gradient and power density, which are used later in the optimization phase.

3.2. Optimization phase

In the optimization phase, MILP approach is used because of two main reasons:

1. MILP solvers check immediately if a scenario is feasible or not.
2. If the problem is well formulated, this approach offers feasible solutions in a short time.

In order to use this approach, several linear approximations must be performed to the non-linear models, starting with the thermal model, which includes non-linear and differential equations. The temperature of a single cell depends not only on the power dissipated by the cell itself, but also on the power dissipated by its neighbors. The first and main factor, refers to the activity of the cell, while the second one is related to the diffusion process of heat [20].

For the linear approximation we use the power density that every cell dissipates in the steady state. This is a valid approximation because the main term of the temperature of a cell is given by the power dissipated in the cell. The contribution of the neighbours have a lower impact on the cell temperature.

Another approximation that must be done is related to the distance between active elements. This distance is approximated as the Manhattan distance.

Next, step after analyzing the temperature behavior of different blocks, we firstly sort the functional units, creating a list in descending temperature order, according to their power density.

We select the first N_h , $h = 1$ blocks from the previous list, that are supposed to be the hottest ones, according to the assumption explained before. The number of those N_h blocks is user-defined,

and it mainly depends on the architecture to optimize. In our work, we included in the list all the cores of the system because heat sources mostly decide the final temperature behavior of the 3D integrated circuit.

Every functional unit in the model $i (i = 1, 2, \dots, n)$ is characterized by a width w_i , a height h_i , and a length l_i , as can be seen in Fig. 2. It is also characterized by its normalized power density $dp_i (i = 1, 2, \dots, n)$. The whole chip is considered as a design volume that has a maximum width W , maximum height H , and maximum length L . In the first MILP search algorithm (MILP1) we aim to find a feasible floorplan by maximizing distance between hottest elements (N_h). We define the vector (x_i, y_i, z_i) as the geometrical location of block i , where $0 \leq x_i \leq L - l_i$, $0 \leq y_i \leq W - w_i$, $0 \leq z_i \leq H - h_i$. We use (x_i, y_i, z_i) to denote the left-bottom-back coordinate of block i while we assume that the coordinate of left-bottom-back corner of the resultant IC is $(0, 0, 0)$. The first proposed MILP for the 3D IC, noted by MILP1, is formulated in Fig. 3.

In Fig. 3, binary variables l_{ij} , r_{ij} , u_{ij} , a_{ij} , b_{ij} and f_{ij} are equal to 1 if block i is in the left of block j , and respectively, right, under, above,

behind and in front. Similarly, kx_{ij} is equal to 1 if block i is in the left of block j and in the same layer. With this argument we can easily place the blocks in the design volume. The condition of no overlap between two blocks is guaranteed by constraints (8)–(14). The distance between two blocks in the x axis is computed through constraints (15)–(20). The same constraints are contained in the complete MILP1 model for both the y and z axis (constraints 21–32). Finally, constraint (33) computes the sum of the Manhattan distances among hottest blocks in N_h divided by the product of normalized power densities.

The optimization can be repeated several times allocating the following set of $N_h, h = 2$ blocks of the remaining sorted list (having the previous N_1 blocks fixed in the final design through constraints (5)–(7)). This procedure can be repeated until the sorted list is empty.

Finally, we move the remaining blocks (those blocks that have less power density, and are considered as heat sinks), using a second search algorithm, called MILP2 in Fig. 1. The algorithm works as MILP1 but, in this case, we do not try to maximize the distance

$$\begin{aligned}
 & (MILP1) \max J_1 \\
 & s.t. \\
 & xmin_i \leq x_i \leq xmax_i \quad i = 1, \dots, n, \quad (5) \\
 & ymin_i \leq y_i \leq ymax_i \quad i = 1, \dots, n, \quad (6) \\
 & zmin_i \leq z_i \leq zmax_i \quad i = 1, \dots, n, \quad (7) \\
 & l_{ij} + r_{ij} + u_{ij} + a_{ij} + b_{ij} + f_{ij} \geq 1 \quad i < j = 1, \dots, n, \quad (8) \\
 & x_i - x_j + L * l_{ij} \leq L - l_i \quad i < j = 1, \dots, n, \quad (9) \\
 & x_j - x_i + L * r_{ij} \leq L - l_j \quad i < j = 1, \dots, n, \quad (10) \\
 & y_i - y_j + W * b_{ij} \leq W - w_i \quad i < j = 1, \dots, n, \quad (11) \\
 & y_j - y_i + W * f_{ij} \leq W - w_j \quad i < j = 1, \dots, n, \quad (12) \\
 & z_i - z_j + H * u_{ij} \leq H - h_i \quad i < j = 1, \dots, n, \quad (13) \\
 & z_j - z_i + H * a_{ij} \leq H - h_j \quad i < j = 1, \dots, n, \quad (14) \\
 & dx_{ij} \geq x_i + l_i/2 - x_j - l_j/2 \quad i < j \in N_h, \quad (15) \\
 & dx_{ij} \geq x_j + l_j/2 - x_i - l_i/2 \quad i < j \in N_h, \quad (16) \\
 & x_i + l_i/2 + L * kx_{ij} \geq x_j + l_j/2 \quad i < j \in N_h, \quad (17) \\
 & x_j + l_j/2 + L * (1 - kx_{ij}) \geq x_i + l_i/2 \quad i < j \in N_h, \quad (18) \\
 & dx_{ij} \leq 2 * L * kx_{ij} + x_i + l_i/2 - x_j - l_j/2 \quad i < j \in N_h, \quad (19) \\
 & dx_{ij} \leq 2 * L * (1 - kx_{ij}) + x_j + l_j/2 - x_i - l_i/2 \quad i < j \in N_h, \quad (20) \\
 & dy_{ij} \geq y_i + w_i/2 - y_j - w_j/2 \quad i < j \in N_h, \quad (21) \\
 & dy_{ij} \geq y_j + w_j/2 - y_i - w_i/2 \quad i < j \in N_h, \quad (22) \\
 & y_i + w_i/2 + W * ky_{ij} \geq y_j + w_j/2 \quad i < j \in N_h, \quad (23) \\
 & y_j + w_j/2 + W * (1 - ky_{ij}) \geq y_i + w_i/2 \quad i < j \in N_h, \quad (24) \\
 & dy_{ij} \leq 2 * W * ky_{ij} + y_i + w_i/2 - y_j - w_j/2 \quad i < j \in N_h, \quad (25) \\
 & dy_{ij} \leq 2 * W * (1 - ky_{ij}) + y_j + w_j/2 - y_i - w_i/2 \quad i < j \in N_h, \quad (26) \\
 & dz_{ij} \geq z_i + h_i/2 - z_j - h_j/2 \quad i < j \in N_h, \quad (27) \\
 & dz_{ij} \geq z_j + h_j/2 - z_i - h_i/2 \quad i < j \in N_h, \quad (28) \\
 & z_i + h_i/2 + H * kz_{ij} \geq z_j + h_j/2 \quad i < j \in N_h, \quad (29) \\
 & z_j + h_j/2 + H * (1 - kz_{ij}) \geq z_i + h_i/2 \quad i < j \in N_h, \quad (30) \\
 & dz_{ij} \leq 2 * H * kz_{ij} + z_i + h_i/2 - z_j - h_j/2 \quad i < j \in N_h, \quad (31) \\
 & dz_{ij} \leq 2 * H * (1 - kz_{ij}) + z_j + h_j/2 - z_i - h_i/2 \quad i < j \in N_h, \quad (32) \\
 & J_1 = \sum_{i < j \in N_h} (dx_{ij} + dy_{ij} + dz_{ij}) / (dp_i * dp_j) \quad (33)
 \end{aligned}$$

Fig. 3. MILP1 3D placement.

$$\begin{aligned}
& (MILP3) \min J_3 \\
& s.t. \\
& d1_{ij} = dx_{ij} + dy_{ij} & i < j \in C_0, \quad (35) \\
& d2_{ij} = 0 & i < j \in C_0, \quad (36) \\
& d1_{ij} = 0 & i < j \in C_1, \quad (37) \\
& d2_{ij} \geq 1 & i < j \in C_1, \quad (38) \\
& d2_{ij} = \sum_{t \in T: l_t \leq z_i, l_t \leq z_j} (dx_{it} + dy_{it} + dx_{tj} + dy_{tj} + dz_{ij}) \cdot kt_t & i < j \in C_1, \quad (39) \\
& J_3 = \sum_{i < j \in C} (d1_{ij} + d2_{ij}) & (40)
\end{aligned}$$

Fig. 4. MILP3 TSV placement.




Fig. 5. Niagara original floorplan.

among hottest blocks; instead, we minimize total wire length, approximated as the Manhattan distance between connected blocks (C). This way, the thermal profile of the stack will be slightly changed (no more than 0.5° in average) but wire length can be minimized. Note that MILP2 is quite similar to MILP1, with the difference that dx , dy and dz are computed for all the interconnected blocks, and J_1 is replaced by the following objective J_2 .

$$\min J_2 = \sum_{i < j \in C} (dx_{ij} + dy_{ij} + dz_{ij}) \quad (34)$$

The last step for the optimization phase is to calculate the optimum number of TSVs to minimize wire length.

Since we have already placed the functional units in the previous phases, we examine the remaining free cells in the resultant stack and build an array of x - y coordinates of allowed TSVs. Thus,

a TSV $t \in 1, 2, \dots, T$ is characterized by a bottom level l_t (all the TSVs start at the top level) and the corresponding x - y coordinate (x_t, y_t) .

In Fig. 4, C_0 is the set of connected blocks in the same level, whereas C_1 is $C - C_0$. Parameter $d1$ represents the distance between connected blocks in the same level (no TSV is needed) and $d2$ is the distance between connected blocks placed at different levels. Finally, binary variable kt_t is equal to 1 if the TSV t must be inserted in the chip. Parameters dx , dy and dz are computed from the previous MILP2 solution. As can be seen in Fig. 4, the optimization model tries to minimize wire length, and thus, maximize the number of TSVs. Obviously, the number of TSVs can be limited with an additional constraint:

$$\sum_{t \in T} kt_t \leq K \quad (41)$$

where K may define the bandwidth constraints between connections.

In the following, we will define the experimental set-up, showing the floorplans that will be thermally analyzed and compared with the results obtained by our floorplanner.

4. Experimental set-up

The 3D multiprocessor scenario studied in our experimental work is based on the Niagara architecture, fabricated in 90 nm technology (these cores are much more powerful than the Power cores found in SCC and will also exhibit higher thermal issues). This architecture has been extended and an increased number of cores can be placed in several layers of the 3D stack. For this purpose, the original architecture has been replicated several times and stacked vertically to build the 3D system.




Fig. 6. Original floorplan for scenarios 2 and 3.




Fig. 7. Original Niagara (A) and Floorplanner Niagara (B) thermal behavior.




Fig. 8. 48 Homogeneous core original thermal behavior.

The communication subsystem of the cores has been also modified from the original. In particular, Niagara crossbar considers a maximum bandwidth of 130 GB/s for the eight cores that integrate the chip. Since our floorplanner can place a variable number of cores in every layer, the power consumption of the crossbar is scaled accordingly to the number of cores found in every layer and their required bandwidth. The inter-layer communication is resolved with a set of TSVs that route the communication signals.

The floorplanner will place the functional units that compose the 3D multi-processor architecture targeting both temperature and wire length optimization.

The experimental work will analyze the thermal optimization achieved by the floorplanner in three different scenarios. The first

scenario presents the original Niagara architecture (Fig. 5 with 8 cores in one layer).

The second and third scenarios are shown in Fig. 6. The second scenario resembles the SCC architecture with a system where 48 and 64 SPARC cores are integrated in the 3D stack, while the third one models an heterogeneous system where the 48 and 64 cores are composed of SPARC and Power6 cores. The ratio of Power6 cores will be 1/4. This setup will show the optimized thermal profile that can be expected when multiple core architectures are considered, as well as the extra optimization opportunities that the floorplanner will find.

In these patterns, the empty space is dedicated to routing and communication units, not considered in the description for their negligible thermal impact.




Fig. 9. 48 Homogeneous core optimized thermal behavior.




Fig. 10. 48 Heterogeneous core optimized thermal behavior.




Fig. 11. 48 Cores average temperature.




Fig. 12. 48 Cores thermal gradient.




Fig. 13. 48 Cores maximum temperature.

5. Results

In this section we present the thermal profiles, estimated by the thermal model, of the floorplanning configurations described in Section 4.

The metrics considered for the analysis of the experimental results are the mean and maximum temperature of the layer and the maximum thermal gradient. These metrics are usually found in all the thermal-related analysis. In order to evaluate the impact of decreasing the temperature over the performance of the system we also include in our analysis the wirelength overhead.

Table 2

48-Core wire length.

Floorplan	Wire length (μm)	Mean temperature (K)
Original homogeneous	3742	347
Optimized homogeneous	4370	340
Original heterogeneous	3696	341
Optimized heterogeneous	4519	333

The worst case of power consumption in the Niagara2 (84 W at 1.1 V and 1.4 GHz [18]) is considered to extract the power densities of every functional unit. Also, the area of the layers has been kept constant while the number of integrated cores and layers is increased. Power consumption of the Power6 core has been set to 2.6 W. This value can be found in [21].

TSVs will be shown in the images as black spots. TSVs connect the top layer with the other ones to guarantee communication among layers.

5.1. Scenario 1.8 cores Niagara architecture

In this first scenario a simple optimization of a 2D system is considered. Taking as initial point the original Niagara distribution, we optimize the placement of the units with our floorplanner. As is explained in Section 3, spreading heat sources as far as possible will contribute to a decrease in the thermal parameters.

This fact is considered by our floorplanner and, as can be seen in Fig. 7, it succeeds on spreading the placement of the cores. Although our floorplanner places some cores together, most of them are located close to the border of the chip, making easier the heat dissipation. This placement obtains a decrease of 3° for the mean temperature, seven for the maximum temperature and nine in the gradient.

Once that the capability of our floorplanner for decreasing the temperature of a 2D design has been checked, we will study more restrictive 3D designs with an increased number of cores distributed in 4 and 5 layers, where TSVs will come into play.

5.2. Scenarios 2 and 3

A more exhaustive analysis has been conducted for the 48 and 64 homogeneous and heterogeneous 3D configurations. These results analyze the maximum and mean temperature, thermal gradient and wirelength of the optimized distributions compared to the original one described in Section 4.

The results will be compared according to the number of cores. Firstly the thermal behavior of the 48-core optimized floorplan for the homogeneous and heterogeneous configurations can be seen in Figs. 9 and 10, respectively. If we compare it with the thermal distribution of the original floorplan in Fig. 8 we can easily find out how our floorplanner has succeeded in optimizing the heat spread across the chip, achieving a reduction of the maximum temperature.

The results for the mean temperature, thermal gradient and maximum temperature for the 48 core system are shown in Figs. 11–13, where original and optimized systems for both, homogeneous and heterogeneous configurations are drawn. The comparison with the baseline homogeneous and heterogeneous systems, shows that the floorplanner is capable of optimizing the maximum temperature in 80° in average, the mean temperature in 8 and the thermal gradient is decreased in 90°.

These optimal results can be explained because our floorplanner spreads heat sources (cores) as much as possible, trying to place them close the border of the chip (helping, in this way, to the cooling down of the cores). The floorplanner also takes into account vertical heat spread, and each layer will show a significantly different layout, avoiding placing heat sources one over the others.




Fig. 14. 64 Homogeneous core optimized thermal behavior.




Fig. 15. 64 Homogeneous core optimized thermal behavior.




Fig. 16. 64 Cores average temperature.




Fig. 17. 64 Cores maximum temperature.




Fig. 18. 64 Cores thermal gradient.

Table 3
Thermal deviation (K).

Scenario	Layer 1	Layer 2	Layer 3	Layer 4
Original homogeneous	30.04	27.33	24.57	22.61
Optimized homogeneous	9.14	9.00	8.64	8.07
Original heterogeneous	26.69	24.19	21.70	19.98
Optimized heterogeneous	7.26	7.19	6.89	6.36

Table 4
Wirelength improvement.

Floorplan	Wire length without TSVs (μm)	Wire length with TSVs (μm)	improvement (%)
Original homogeneous	9402	4800	49
Optimized homogeneous	9561	4848	49
Original heterogeneous	13969	7923	43
Optimized heterogeneous	13794	6284	54

Table 2 contains the wirelength associated with the 3D stack and the mean temperature of the chip for the original and optimized floorplans.

Table 2 shows the wirelength associated with the 3D system. As can be seen, the overhead incurred by the floorplanner has been a 17% when compared to the original distribution for the homogeneous design and 22% for the heterogeneous one. This overhead in the wiring is not directly translated into an increase of the communication delay because core-to-core communication is regulated by the crossbar. As the crossbar is the module that limits the bandwidth and speed of the link, this overhead is seen minimized. On the other hand, the big savings reached in the mean and maximum temperature justify the overhead in wiring.

A similar study was conducted for the 64 core design. Figs. 14 and 15, show the thermal behavior of the homogeneous and heterogeneous optimized floorplans.

Similarly to the previous setup, thermal metrics for every layer of the stack have been calculated. Comparing each optimized system with the baseline case, it can be seen a reduction of 73° and 78° in maximum temperature for the homogeneous and heterogeneous systems respectively, 7° in the mean temperature and a reduction of 83° in the gradient (see Figs. 16, 18 and 17). Also, the heterogeneous architecture outperforms the results of the homogeneous system in 20° for the maximum temperature, 11 for the mean temperature and 23 for the thermal gradient.

This results can be explained because Power cores will not be considered as hotspots by the optimizer, since their power density consumption is much lower than SPARC's. Therefore, the floorplanner will place SPARC cores considered as strong heat sources near the Power cores, achieving a better thermal profile.

Also, as shown in Table 3, our floorplanner is able to reduce the standard deviation in the temperatures across the layers. This determines a more homogeneous thermal distribution, which is translated into a reduced reliability risk and diminished leakage currents.

In both scenarios TSVs are deployed to allow inter layer communication. The number of TSVs is enough to guarantee bandwidth restrictions among layers. Chosen the number of TSVs, their location in the chip and the layers that connect each TSV are calculated by our floorplanner in order to minimize wirelength. Despite TSVs are absolutely necessary in 3D technology, they also contribute to minimize communication paths as can be seen in Table 4, where wirelength is compared in optimized floorplans before and after having placed the TSVs.

6. Conclusions

This paper has proposed a novel MILP formulation to cope with the problem of thermal-aware floorplanning in 3D MPSoCs optimizing the location of functional units and through silicon vias. Also, the efficient solver that provides the optimization of the floorplan, interfaces with an accurate thermal model, providing

promising results in the minimization of the main thermal and reliability-related metrics (peak and mean temperature, thermal gradients) with low performance overhead. The experimental results have been conducted with a realistic platform based on the Niagara architecture, outperforming previous results obtained by traditional thermal-aware floorplanner.

References

- [1] D. Cuesta et al., Thermal-aware floorplanning exploration for 3d multi-core architectures, in: *GLSVLSI '10*, 2010, pp. 99–102.
- [2] I.K.Y. Han, C. Moritz, Temperature aware floorplanning, in: *Workshop on Temperature-Aware Computer Systems*, 2005.
- [3] C.C.N. Chu, D.F. Wong, A matrix synthesis approach to thermal placement, in: *International Symposium on Physical Design*, 1997.
- [4] G. Chen, S. Sapanetkar, Partition-driven standard cell thermal placement, in: *International Symposium on Physical Design*, p. 2003.
- [5] K. Sankaranarayanan, S. Velusamy, M.R. Stan, K. Skadron, A case for thermal-aware floorplanning at the microarchitectural level, *Journal of Instruction Level Parallelism* 7 (2005) 8–16.
- [6] M. Ekpanyapong, M.B. Healy, C.S. Ballapuram, S.K. Lim, H. Hsin S. Lee, Thermal-aware 3d microarchitectural floorplanning, *Georgia Institute of Technology Center for, Tech. Rep.*, 2004.
- [7] W.-L. Hung, Y. Xie, N. Vijaykrishnan, C. Addo-Quaye, T. Theocharides, M.J. Irwin, Thermal-aware floorplanning using genetic algorithms, in: *International Symposium on Quality Electronic Design*, 2005.
- [8] Y. Han, I. Koren, Simulated annealing based temperature aware floorplanning, *Journal of Low Power Electronics* 3 (2) (2007) 1–15.
- [9] Y.-C. Chen, Y. Li, Thermal-aware floorplanning via geometric programming, *Mathematical and Computer Modelling* 51 (2010) 927–934.
- [10] J. Cong, J. Wei, Y. Zhang, A thermal-driven floorplanning algorithm for 3d ics, in: *ICCAD*, 2004.
- [11] M. Healy, et al., Multiobjective microarchitectural floorplanning for 2-d and 3-d ics, in: *Trans. on CAD*, vol. 26, 2007, pp. 38–52.
- [12] W.-L. Hung et al., Interconnect and thermal-aware floorplanning for 3d microprocessors, in: *ISQED*, 2006, pp. 98–104.
- [13] J. Cong, M. Sarrafzadeh, Incremental physical design, in: *ISPD*, 2000.
- [14] P.B.J. Creshaw, M. Sarrafzadeh, P. Prabhakaran, An incremental floorplanner, in: *GLSVLSI*, 1999.
- [15] M.L.S. Liao, D. Mehta, Constrained polygon transformations for incremental floorplanning, *ACM Transactions on DAES* 6 (3) (2001).
- [16] R.T.X. Tang, M. Wong, Optimal redistribution of white space for wire length minimization, in: *ASP-DAC*, 2005.
- [17] Y.M.X. Li, X. Hong, A novel thermal optimization flow using incremental floorplanning for 3d ics, in: *ASP-DAC*, 2009, pp. 347–352.
- [18] <<http://www.opensparc.net/pubs/preszo/07/n2isscc.pdf>>, 2010.
- [19] J.L. Ayala, et al., Through Silicon Via-based Grid for Thermal Control in 3D Chips, in *Nano-Net*, ser. *Lecture Notes in Computer Science*, vol. 1, no. 1. Springer, 2009, pp. 1–7.
- [20] L.B.G. Paci, F. Poletti, P. Marchal, Exploring temperature-aware design in low-power mpsoes, *IJES International journal of embedded systems* 3 (1) (2007) 43–51.
- [21] IBM Systems Energy Estimator, <<http://www.ibm.com/systems/support/tools/estimator/energy/>>, 2010.



David Cuesta completed got a MS in Physics and a MS in Electrical Engineering from Complutense University of Madrid, both in 2008. He is currently a Ph.D. student at the same university and his research interests include thermal-aware run-time techniques for multi-processors, thermal simulation and thermal-aware floorplanning.



Jos L. Risco-Martín is Assistant Professor at the Computer Architecture and Automation Department of Complutense University of Madrid (UCM), Spain. His research interests focus on computational theory of modeling and simulation, with emphasis on Discrete Event Systems Specification (DEVS), dynamic memory management for embedded systems and evolutionary computation.



Jose L. Ayala got his MS in Telecommunication Engineering and his Ph.D. in Electrical Engineering in 2001 and 2005, respectively, both from Politecnica University of Madrid. He also has a MS in Physics from the Open University of Spain. Prof. Ayala is currently an Associate Professor in the Complutense University of Madrid (Spain) and a Permanent Visiting Professor at EPFL (Switzerland). His research interests include thermal-aware electronic design and thermal-aware compilation, thermal-modeling and power-efficient computer architectures.