Temperature-Aware Global Placement

Bernd Obermeier
Institute of Electronic Design Automation
Technical University of Munich, 80290 Munich
first_name.last_name@ei.tum.de

Frank M. Johannes
Institute of Electronic Design Automation
Technical University of Munich, 80290 Munich
first_name.last_name@ei.tum.de

Abstract— This paper describes a deterministic placement method for standard cells which minimizes total power consumption and leads to a smooth temperature distribution over the die. It is based on the Quadratic Placement formulation, where the overall weighted net length is minimized. Two innovations are introduced to achieve the above goals. First, overall power consumption is minimized by shortening nets with a high power dissipation. Second, cells are spread over the placement area such that the die temperature profile inside the package is flattened.
Experimental results show a significant reduction of the maximum temperature on the die and a reduction of total power consumption.

I. Introduction

Technological advances have provided new challenges for EDA tools during the recent years. The ability to handle more complex designs and a more precise modeling of the physical effects were key factors for successful design tools.
Today, an upcoming issue is power consumption and the temperature distribution over the substrate. Low power consumption is not only a desirable property of modern designs (esp. in mobile applications), but will also reduce the maximum temperature on the die. The problem of minimizing overall power consumption has mostly been addressed on the higher design levels in the past [9, 13]. Up to now, the optimization potential of physical design to reduce power consumption has not yet been exploited intensively.
The temperature distribution on the substrate has various effects on the circuit. Lifetime is heavily dependent on substrate temperature due to the effect of electromigration ([11]). Furthermore, since current drive capability of a CMOS-driver decreases with increasing temperature, timing behavior deteriorates. There are also reports of circuit failure due to large unforeseen thermal gradients on the substrate [8]. Furthermore, high die temperature often requires active cooling.

The temperature distribution on the substrate derives from the power dissipation of the circuitry. This can be classified into four categories ([4]):

- Dynamic power: When the output of a gate switches, the capacitances of the wires and the inputs driven by the gate have to be charged.
- Internal power: Power is consumed when charging the gate-drain capacitances during a value change on an input of the gate.
- Short-circuit power: During transition there might be a DC-path from power source to ground.

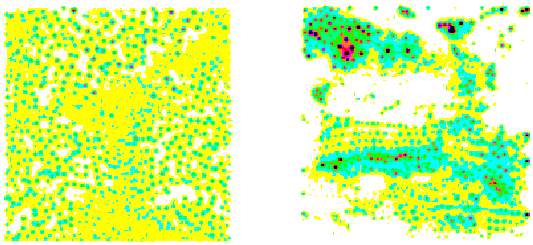

Fig. 1. Die temperature distribution with (left) and without (right) temperature-aware global placement method.

- Leakage power: Even if the gate-source voltage is smaller than the threshold voltage, a CMOS-transistor is not a perfect isolator and a small current always flows.

According to the SIA roadmap [2] dynamic power is responsible for about 85 percent of the total power dissipation. Dynamic power dissipation can be influenced by layout, it is dependent on the capacity C_{load} the cell has to drive, the supply voltage, the clock frequency and the switching rate. C_{load} consists of the input capacity of the driven gates as well as the interconnect capacity of the driven net. Since C_{load} is therefore a function of the net length, shorter nets mean less power dissipation. Thus a method for temperature-aware global placement should feature the following properties:

- minimization of the total power dissipation
- a flattened temperature profile, based on an accurate temperature modeling
- sufficiently small cell overlaps to be removed by a final placer
- ability to handle large modern designs (low numerical complexity, preplaced cells, macro cells)

According to our knowledge, there are three relevant publications dealing with the thermal-driven placement problem. In [12, 4], the authors propose a simulated annealing approach. During a preprocessing step a power distribution on the die which leads to a uniform temperature profile is computed by simulation, whereby the influence of the package on the temperature distribution is respected. Although this approach offers the flexibility to handle most of the necessary properties, it suffers two drawbacks: efficiency of simulated annealing might not be sufficient for large designs and flexibility may not satisfy the demands of modern designs (e.g. placement problems with standard cells together with macro blocks).
The work of [5] proposes a matrix synthesis approach to the thermal placement problem. The placement area is represented by a matrix. They try to find a matrix, such that the power

dissipation of the cells in every $t \times t$ matrix is minimal. Again, this approach leads to a uniform *power*-distribution instead of a uniform *temperature*-distribution, since the influence of the package is not modeled. In addition, no minimization of the global power dissipation takes place.
In [3] the authors propose a partitioning-based approach. They extend a min-cut partitioner to balance the thermal gradient between the partitions. Temperature calculation is performed like in [12], where the temperature grid is equivalent to the partition blocks. This method suffers from the common drawbacks of partitioning algorithms (e.g. no ability to handle largely different cell sizes) and focuses only on minimizing the thermal gradient, not the maximum temperature.
Our algorithm is based on the Quadratic Placement technique presented in [6] which distributes the cells evenly over the placement area using a force field. We will show that this method can also be extended to distribute the cells over the die to achieve a flattened temperature profile.
In a first step a non-uniform power density profile is computed analytically which models the fact that the heat dissipation increases with the distance from the die center. The process of distributing the cells is thus guided by power dissipation rather than by area and leads to a flattened temperature profile. The advantage of this approach is, that we do not have to perform a costly three-dimensional thermal simulation in every iteration step of the algorithm.
The second step is the common Quadratic Placement algorithm where the cost function models the total power dissipation rather than net length.
The following properties distinguish our approach from previous methods:

- Analytical computation of the power density profile which leads to a flat temperature distribution based on package and chip geometry.
- Minimization of the global power dissipation, i.e. the sum of the power dissipations of all cells.
- Suitability to process large designs because of low computational complexity.

A visualization of the capabilities of the proposed algorithm is shown in Figure 1, where darker colors mean higher temperatures. The figure compares the thermal distributions of a benchmark circuit placed with (left picture) and without (right picture) the temperature-aware placement method. The hot spots in the top left corner of the circuit have vanished almost completely.
The rest of the paper is organized as follows: In section II we introduce the temperature-aware placement task. Section III reviews the force-directed placement method. In section IV the new cost function and the thermal model of the package are introduced. The results of section III and IV are used to present the temperature-aware global placement algorithm in section V. Experimental results are given in section VI and section VII concludes the paper.

II. Problem Description

In this section, we define the global placement task with special consideration of the particularities of the temperature-aware placement.
The input to the temperature-aware placement problem consists of a graph which models the connectivity between the cells, the geometric data of the cells and the placement area (master). Furthermore, the switching rates of each net have to be determined beforehand, which could be difficult in practice. The goal of the global placement is to find a configuration of the cells within the placement area such that some cost function is minimized and a set of constraints is satisfied. In our case we try to minimize the total power dissipation and arrange the cells within the placement area in a way such that the resulting temperature profile is flattened. In a final placement step the standard cells are mapped to rows and overlap between the cells gets eliminated.
Details about the cost function and satisfaction of the constraints will be given in the subsequent sections.

III. Force-Directed Placement

This section gives a review of the force-directed placement algorithm [6]. In this introduction we do not pursue a flat temperature profile, this will be discussed later. The force-directed placement method consists of two closely coupled components: The minimization of the cost function and the satisfaction of constraints.
In general, the cost function is composed of a weighted sum of all net lengths. A net is modeled as a clique and length is measured as squared Euclidean distance between two points. Assuming there are n movable cells we introduce the placement vector $\vec{p} = (x_1, x_2, \dots x_n, y_1, y_2, \dots y_n)$. The cost function can then be written in matrix notation:

$$\frac{1}{2}\vec{p}^T C\vec{p} + \vec{d}^T\vec{p} + const$$

Matrix C represents the interconnect topology and $\vec{d}$ models preplaced cells. Net weights influence both C and $\vec{d}$. Minimizing the cost function leads to a unique solution (placement, positions of movable cells) by solving the linear equation system $C\vec{p} = \vec{d}$. If we consider nets as springs, the minimum of the cost function corresponds to the state of equilibrium, i.e. the positions of the movable cells are determined by the positions of the fixed cells. For the rest of the paper we will use this spring analogy when describing the proposed methods. We say, the positions of the cells are determined by forces.
In general, a placement which minimizes the above cost function will not distribute the cells evenly. In order to achieve this the linear equation system is extended with a vector $\vec{e}$:

$$C\vec{p} = \vec{d} + \vec{e}$$

Vector $\vec{e}$ models a force working on each movable cell. It is chosen such that the cells are evenly distributed and the value of the cost function is kept as low as possible. In order to spread the cells evenly across the placement area, empty regions of the placement area should attract cells, whereas cells repel each other. This postulation (among others, for a detailed description see [6]) constitutes a force field over the placement area which spreads the cells evenly across the placement area. From a placement a two dimensional histogram can be derived: $D(x,y) = D_{\text{master}}(x,y) + D_{\text{cells}}(x,y)$ where $D_{\text{master}}(x,y)$ is a negative constant at all points (x,y) covered by the placement area and 0 otherwise. Every cell adds a constant positive value to all points (x,y) of $D_{\text{cells}}(x,y)$ which are covered by the cell. The force field is determined by the solution of Poisson's equation with $D(x,y)$ as right hand side [6]: $\nabla^2 U(x,y) = D(x,y)$, where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z})^T$. $U(x,y)$ can be efficiently computed using a multigrid solver [10], e.g.. The force field is then $\vec{f}(x,y) = \nabla U(x,y)$. This technique has an

interpretation as supply-demand model: the sink models the area supply offered by the placement area and the stacks are the area demands of the cells.
The complete placement procedure consists of alternating steps of minimizing the cost function and spreading the cells using the force field.
The following section is dedicated to the analysis of the heat sources (power distribution) and the thermal properties of the package. This is necessary to enhance the basic algorithm to temperature-aware placement.

IV. Thermal and power Analysis

This section is divided into two parts: The treatment of the heat sources and the analysis of the temperature distribution inside the package. These efforts lead to the derivation of a distribution of heat sources which ensures a temperature distribution over the die which is as flat as possible.

A. Power Analysis

Since dynamic power dissipation accounts for about 85 percent of the total power dissipation, like [12] we disregard internal power, leakage power and short circuit power for the rest of the paper. The power dissipation of a logic gate can be expressed as

$$P = 0.5 \cdot C_{\text{load}} \cdot V_{\text{DD}}^2 \cdot f \cdot E \tag{1}$$

on average. V_{DD} is the voltage of the power source, f is the clock frequency, and E is the switching rate. The total capacity driven by the gate is composed of two parts: $C_{\text{load}} = C_{\text{pin}} + C_{\text{net}}$, where C_{pin} is the capacity sum of all pins and C_{net} is the interconnect capacity of the net. C_{net} can be influenced by global placement, it depends on the positions of the pins and the routing of the net. When computing the capacity of a net, the routing of a net is approximated by a Steiner tree [7]. It results in $C_{\text{net}} = l_{\text{net}} \cdot c$ where l_{net} is the length of the Steiner tree and c is the capacity per unit length. Shortening tree length (i.e. keeping the pins of a net closer together) therefore results in less power dissipation.

B. Temperature Analysis

The goal of this subsection is to derive the *power* distribution over the die which leads to a constant *temperature* distribution over the die.
The strategy to achieve this consists of two steps: first, we compute the three-dimensional temperature distribution inside the package assuming a constant temperature over the die. Afterwards we employ Poisson's equation to derive the appropriate power distribution, which will then be used to guide the temperature-aware placement process.
Temperature conduction is governed by Poisson's equation $\nabla^2 T(x,y,z) = h(x,y,z)$, where $h(x,y,z)$ describes the source power density inside the medium. We do not introduce a coefficient for thermal conductivity (i.e. we set it to 1), since we assume a homogeneous medium and the constant would only scale temperature. Temperature distribution inside a body depends on the geometry of the body, the thermal boundary conditions and the power sources inside the body as described by $h(x,y,z)$. In order to get an analytical description of the temperature distribution inside the package, we make the following assumptions (see Figure 2):

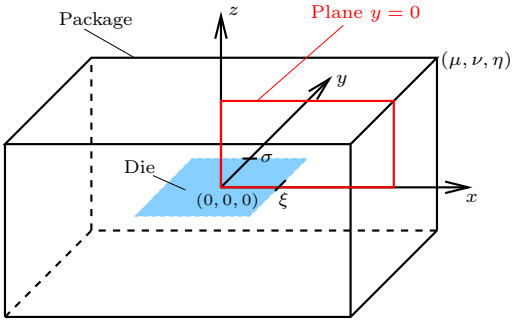

Fig. 2. Package Geometry

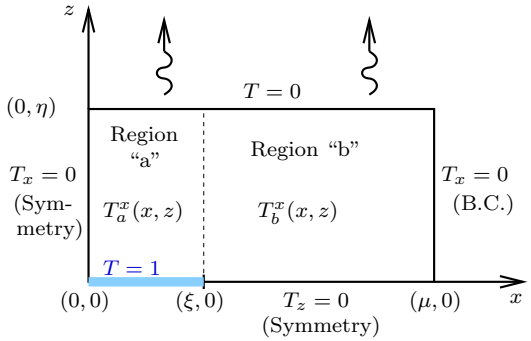

Fig. 3. Plane $y = 0$ of the first quadrant of $T^x(x,z)$

- the package has the shape of a cuboid and the point of origin is at its center of gravity
- the package dimensions in (x,y,z)-direction are $(2\mu, 2\nu, 2\eta)$
- the source of power is a rectangle centered inside the package with dimensions $(2\xi, 2\sigma)$ in (x,y)-direction.
- the boundary conditions on planes $z = \eta$ (top) and $z = -\eta$ (bottom) are $T = const$ (w.l.o.g. $T = 0$)
- the boundary conditions on planes $x = \mu$ and $x = -\mu$ ($y = \nu$ and $y = -\nu$) are $\partial T/\partial x = 0$ ($\partial T/\partial y = 0$) (sides are isolated).

Temperature distribution $T(x,y,z)$ inside the package is a superposition of two temperature distributions $T^x(x,z)$ and $T^y(y,z)$ describing the temperature behaviour in $x-z$ and $y-z$ planes, respectively. The same holds true for the power density $h(x,y,z) = h^x(x,z) + h^y(y,z)$. Due to symmetry we may restrict ourselves to the first quadrant of $T^x(x,z)$ since the derivation of the temperature distribution is identical for $T^x(x,z)$ and $T^y(y,z)$ because of the assumptions made above. Figure 3 shows the first quadrant $T^x(x,z)$ of plane $y = 0$ taken from Figure 2.

For reasons of clarity, we use two-dimensional notation in x and z in Figure 3. The plane shown in the figure is divided into the regions "a" (rectangle from $(0,0)$ to (ξ,η)) and "b" (rectangle from $(\xi,0)$ to (μ,η)) due to different boundary conditions on the x-axis. The temperature distribution within region "a" is $T_a^x(x,z)$ and the temperature distribution within region "b" is $T_b^x(x,z)$. The bold interval on the x-axis ($0 < x < \xi$) depicts a cut through the die. We found that the temperature distributions in region "a" and "b" can be formulated as series, where the coefficients a_n and b_n will be determined from package geometry:

$$T_a^x(x,z) = \sum_{n=1}^{\infty}(a_n \cdot \cosh(\lambda_n \cdot x) \cdot \sin(\lambda_n \cdot z)) + (1 - \frac{z}{\eta}), \text{where } \lambda_n = -\frac{\pi}{\eta} \cdot n \quad (2)$$

$$T_b^x(x,z) = \sum_{n=1}^{\infty}(b_n \cdot \cosh(\delta_n \cdot (x-\mu)) \cdot \cos(\delta_n \cdot z)), \quad (3)$$
$$\text{where } \delta_n = -\frac{\pi}{\eta} \cdot \frac{2n-1}{2}$$

When choosing $T_a^x(x,z)$ and $T_b^x(x,z)$ like that, $\nabla^2 T_a^x(x,y) = 0$ and $\nabla^2 T_b^x(x,y) = 0$ holds and all boundary conditions are satisfied *regardless* of coefficients a_n and b_n. Adaption of the temperature distributions in region "a" and "b" ($T_a^x(\xi,z) = T_b^x(\xi,z)$ and $\partial T_a^x(\xi,z)/\partial x = \partial T_b^x(\xi,z)/\partial x$) leads to a linear equation system. When collecting the coefficients a_n up to some fixed approximation order N in vector $\vec{a} = (a_1, a_2, \ldots, a_N)^T$, $\vec{a}$ is determined by:

$$M \cdot \vec{a} = \vec{r} \quad (4)$$

Entries m_{ij} of matrix $M_{N\times N}$ are

$$m_{ij} = \frac{j\tanh(\delta_i(\xi-\mu))(2i-1)\cosh(-\lambda_j\xi) - 2j^2\sinh(-\lambda_j\xi)}{-4j^2 + (2i-1)^2} \quad (5)$$

and components r_j of vector $\vec{r}$ are

$$r_j = -\frac{\tanh(\delta_j(\xi-\mu))}{\pi(2j-1)} \quad (6)$$

Since the entries of matrix M and vector $\vec{r}$ solely depend on geometry data and indices, we are able to compute an approximation for $T^x(x,z)$. We obtain an approximation for $T^y(y,z)$ by interchanging ξ with σ and μ with ν. The complete power density of the die which leads to a flat temperature profile is given by Poisson's equation:

$$h(x,y,z) = \nabla^2 T(x,y,z) \quad (7)$$

The ability to compute the power dissipation of a given cell on one hand and the target power density for a constant temperature over the die on the other hand will be used in the next section to introduce the temperature-aware placement process.

V. Temperature-aware Placement

In this section we harvest the fruits of labour and combine the temperature-aware placement method with the results of the previous sections. The temperature-aware placement method makes use of two techniques to minimize and flatten the temperature distribution of the die: First, the reduction of heat generation by shortening the specific nets and second, the distribution of the cells within the placement area so that a flat temperature distribution arises.

To reduce heat generation we increase the weights of these nets which contribute most to heat generation. Our analysis of cell heat generation resulted in two observations: Cells featuring a high heat generation either had to drive nets with a high fanout count or a high switching rate or both. The dependence of power dissipation on the switching rate is obvious. Nets with a high fanout count influence power dissipation in two ways: They have usually a high C_{pin} and tend to spread over a larger area (thus resulting in a large C_{net}). These nets should receive a high net weight in order to keep at least C_{net} low. To account for both reasons, the net weight should increase with the product of switching rate (E_n) and pin count (k_n):

$$w_n = 1 + \alpha E_n k_n \quad (8)$$

α is an implementation dependent constant, which is chosen such that the ratio of the highest net weight in comparison to the lowest net weight is about 50:1 for a typical circuit.

Leveling of the temperature distribution will be achieved by modifications to the computation of the force field described in section III. The forces will be used to distribute the cells over the placement area so that their heat generation will lead to a flattened temperature profile. D_{master} represents no longer the *area* supply but the *power density* supply of the placement area and D_{cells} models now the power density instead of the area of a cell. Using the results of section IV D_{master} is the projection of the three-dimensional power density $h(x,y,z)$ to the two- dimensional space:

$$D_{\text{master}}(x,y) = h(x,y,0) \quad (9)$$

On the other hand D_{cell} has to represent the power dissipation of the cells instead of the area demand. Every cell adds a share proportional to its power dissipation (see (1)) to $D_{\text{cell}}(x,y)$ on the points (x,y) covered by the cell.

Thus, cells are spread according to their power dissipation rather than their area demand. The corners of the placement area are more attracting than the center and cells with a great power dissipation get more space than cells with less power dissipation.

In algorithm 1 we summarize the proposed placement method. Effectiveness and efficiency of this proposed method will be demonstrated in the next section.

Algorithm 1 Temperature-aware global placement

Compute net weights w_n from pin count and switching rate
Compute $D_{\text{master}}(x,y)$ from die and package geometry
Use net weights to initialize $C, \vec{d}$ from netlist
$\vec{e} = 0$
repeat
  Compute placement $\vec{p} : C\vec{p} = \vec{d} + \vec{e}$
  Compute power dissipation for every cell using placement $\vec{p}$ and therewith $D_{\text{cell}}(x,y)$
  $D(x,y) = D_{\text{master}}(x,y) + D_{\text{cell}}(x,y)$
  Compute force field: $\nabla^2 U(x,y) = D(x,y)$
  $\vec{f}(x,y) = \nabla U(x,y))$
  Update $\vec{e}$ by evaluation $\vec{f}$ at the respective position:
  $\vec{e_i} = \vec{f_x}(x_i, y_i)$; $\vec{e_{i+n}} = \vec{f_y}(x_i, y_i)$
until Convergence

VI. Experimental Results

We conducted three experiments using the MCNC [1] benchmark set. We had to restrict ourselves to benchmarks with timing information since there we can identify the net driver pins. Statistical data about the benchmark circuits used are summarized in Table I. Two experiments compare a conventional placement run (as described in section III) with the

circuit	number of cells	number of nets	original core size [$\mu \times \mu$]	enlarged core size [$\mu \times \mu$]
fract	125	147	648 × 696	726 × 696
struct	1888	1920	2340 × 2436	2415 × 2436
biomed	6417	5742	4897 × 5336	5000 × 5336
avq.small	21854	22124	9160 × 9164	9259 × 9164
avq.large	25114	25384	9536 × 9744	9631 × 9744

TABLE I
BENCHMARK SET STATISTIC DATA

temperature-aware placement. The third experiment compares our method with the work of [3]. The experiments were conducted with the following assumptions:

- wire capacitance was 242 pf/m
- switching rate for each net was randomly generated between 0.1 and 0.9
- package $x-y$ dimension is three times the $x-y$ dimension of the die. Package height is $\frac{1}{10}$ of die width.
- for practical relevance, all results given refer to legal placements, i.e. after applying a final place step.

The absolute values for supply voltage V_{DD} and clock frequency f are not relevant since power dissipation of every net is proportional to the constant $V_{\mathrm{DD}}^2 \cdot f$, hence these values only scale the temperature. The same holds true for the thermal conductivity of the package. As stated in section IV, top and bottom surface temperatures of the package are assumed to be 0 degrees.
The result tables consist of several columns which are explained below. "Max temp." is the maximum die temperature over the package surface temperature and has been normalized so that for every benchmark circuit the maximum die temperature equals 20 degrees for the conventional placement run. "Tot. power" is the power dissipation sum of all cells computed by formula (1). A similar scaling as for maximum temperature has been employed for the total power dissipation: power dissipation of every benchmark circuit has been scaled to 100 watts for the conventional placement run. "Hrpm" is the sum of the half rectangle perimeter of all nets and "cpu" is the cumulative computation time needed for global placement and final placement. Furthermore, the relative improvements of the temperature-aware placement method over the conventional placement method are given for these results.
Unfortunately we can not compare ourselves to the results presented by [12] since they use a proprietary variation of the MCNC benchmark set which does not match the statistical data (e.g. number of cells) of the published benchmark set. Similarly, the authors of [5] used the ISCAS'89 benchmark suite which do not contain any geometry information.
Authors [3] modified the MCNC benchmark set to obtain zero row spacing, which matches modern design styles. A comparison to the results of [3] is given in table IV. After global placement the cells have been mapped to the standard cell rows by a final placement step. All experiments have been conducted on Intel Pentium 4 computers with 2.4 Ghz and 512 Mb ram.
The effect of the proposed method is demonstrated in Figure 1. The left picture of the figure shows the temperature distribution over the die after global placement using the temperature-aware placement algorithm. The right picture shows the temperature distribution after a conventional placement run. The hot spots in the upper left corner of the bottom picture have vanished almost completely in the top picture.

A. Temperature-aware placement

The first experiment compares the conventional placement method (see section III) with the temperature-aware placement method (see section V). To obtain a flat temperature profile over the die, power density at the corners of the die is greater than power density at the die center. Since on average power density is correlated with cell density, the temperature-aware placement method as introduced in section V needs some spare area for a legal placement. Furthermore, cells with the highest power dissipation need a little amount of empty space around them to avoid a too high temperature rise. Since area utilization of the MCNC benchmarks is 100 percent, we enlarged the original benchmarks by 3 times the average width of a standard cell along the standard cell row direction (since we did not want to introduce new standard cell rows). Otherwise the global placement would be quite congested at the borders of the placement area. The resulting core dimensions can be found in Table I. Table II shows the comparison of the conventional placement method (as described in section III) with the temperature-aware placement method (see section V) using the enlarged benchmark circuits. Temperature-aware placement decreases the maximum die temperature on average by 24 percent and the total power dissipation on average by 5.7 percent on the expense of total net length. Please note that targeting on a flat temperature profile only does not necessarily decrease power consumption since total net length could be enlarged.

B. Minimization of power dissipation

The second experiment demonstrates the impact of power minimization alone on temperature distribution. This is useful if there is no spare area on the master to level the temperature distribution. In table III we compare the conventional placement method with the temperature-aware placement method with the exception that D_{master} is not computed according to the results of section IV but is a constant. The experiment is conducted with the original MCNC benchmark circuits since these feature 100 percent area utilization.
As expected, maximum temperature decreases less than above since hot spots can not be eliminated by spreading the responsible cells further apart. Note that when there is no excess white space the degree of freedom for optimization is restricted. Thus, the temperature distribution can not get maximal flat and the power dissipation can be higher. The average maximum temperature reduction is 19 percent in this experiment. The relative increment of cpu-time is independent of the circuit size, the numerical complexity of the temperature-aware placement method is identical to the conventional method.

C. Comparison with previous approaches

This experiment was conducted in order to enable a comparison with the results of [3]. They have modified the benchmark circuits in order to obtain zero row spacing, therefore the core height of each benchmark for this experiment is only half of the height given in table I, the core width remains unchanged. The placement method is identical to the method described in section B. The columns "max temp." and "temp. grad." in Table IV give the relative improvement of the maximum temperature and the maximum temperature gradient for placement runs with and without thermal constraints. The column "hrpm" has the same meaning as in the previous tables.
Since our implementation exhibits a 36 percent smaller total

	conventional method				temperature-aware method				rel. improvement [%]			
circuit	max. temp.	tot. power	hrpm [m]	cpu [s]	max. temp.	tot. power	hrpm [m]	cpu [s]	max. temp.	tot. power	hrpm	cpu
fract	20	100	0.0388	0.2	19	96.4	0.0396	0.47	4.95	3.56	-2.12	-135
struct	20	100	0.424	1.08	19.9	99.2	0.449	2.67	0.682	0.797	-6.04	-147
biomed	20	100	2.16	5.33	15.3	88.1	2.59	13.5	23.3	11.9	-20	-153
avq.small	20	100	6.34	38.6	12.9	94.4	7.52	71.8	35.4	5.62	-18.6	-85.8
avq.large	20	100	7.04	48.3	8.51	93.2	8.3	85.5	57.4	6.81	-17.8	-77

TABLE II
TEMPERATURE-AWARE PLACEMENT

	conventional method				power minimization				rel. improvement [%]			
circuit	max. temp.	tot. power	hrpm [m]	cpu [s]	max. temp.	tot. power	hrpm [m]	cpu [s]	max. temp.	tot. power	hrpm	cpu
fract	20	100	0.0354	0.27	18.3	98.7	0.038	0.33	8.71	1.31	-7.11	-22.2
struct	20	100	0.4	1.55	19.4	100	0.432	1.77	3.2	0	-8.01	-14.2
biomed	20	100	2.03	7.72	15.5	88	2.48	9.44	22.5	12	-21.9	-22.3
avq.small	20	100	5.61	49.7	16.5	98.5	7.37	56.3	17.6	1.48	-31.4	-13.2
avq.large	20	100	6.44	54.7	10.6	96.9	8.23	62.1	47.1	3.08	-27.9	-13.4

TABLE III
MINIMIZATION OF POWER DISSIPATION

	improvement of Chen & Sapatnekar			improvement of our method		
circuit	max. temp.	temp. grad.	hrpm [m]	max. temp.	temp. grad.	hrpm [m]
struct	-2 %	25 %	0.501	1.82 %	12.6 %	0.29
biomed	20 %	58 %	2.60	13.5 %	18.2 %	1.59
avq.small	19 %	17 %	6.67	22.5 %	32 %	4.27
avq.large	35 %	39 %	7.46	31.5 %	42.1 %	4.5

TABLE IV
MINIMIZATION OF POWER DISSIPATION WITH ZERO ROW SPACING

wire length on the average of the benchmark circuits, our approach features remarkable lower power consumptions and absolute temperatures.

VII. Conclusion

In this work we presented a temperature-aware placement method. It is based on two foundations: The power minimization by shortening nets which account for most of the power dissipation and the arrangement of the cells so that their power dissipation leads to a flattened temperature distribution over the die.

The conducted experiments proved the effectiveness and efficiency of the proposed method: Both power consumption and die temperature could be reduced significantly.

This technique improves the development of low-power designs with various impacts on the technical realization of systems: Fan cooling might get obsolete for a range of devices and battery lifetime may increases for mobile devices. Since thermal stress is reduced on the die, robustness of the devices will increase. Beside that, the range of operation temperature will also increase.

References

[1] Benchmarks. http://www.cbl.ncsu.edu/.

[2] International technology roadmap for semiconductors: 2000 edition. International SEMATECH, 2000.

[3] G. Chen and S. Sapatnekar. Partition-driven standard cell thermal placement. In *ACM/SIGDA International Symposium on Physical Design (ISPD)*, volume 1, pages 75–80, Apr. 2003.

[4] Y.-K. Cheng, C.-H. Tsai, C.-C. Teng, and S.-M. S. Kang. *Electrothermal Analysis of VLSI Systems*. Kluwer Academic Publishers, Boston, June 2000.

[5] C. C. Chu and D. Wong. A matrix synthesis approach to thermal placement. *IEEE Transactions on Computer-Aided Design of Circuits and Systems*, 17(11):1166–1174, Nov. 1998.

[6] H. Eisenmann and F. M. Johannes. Generic global placement and floorplanning. In *ACM/IEEE Design Automation Conference (DAC)*, pages 269–274, June 1998.

[7] M. Hanan. On Steiner's problem with rectiliner distance. *SIAM Journal of Applied Mathemetics*, 14(2):255–265, 1966.

[8] S.-M. S. Kang. On-chip thermal engineering for peta-scale integration. In *ACM/SIGDA Int. Symposium on Physical Design*, page 76, Apr. 2002.

[9] T. Karnik, Y. Ye, J. Tschanz, L. Wei, S. Burns, V. Govindarajulu, V. De, and S. Borkar. Total power optimization by simultaneous dual-vt allocation and device sizing in high performance microprocessors. In *ACM/IEEE Design Automation Conference (DAC)*, page 486, 2002.

[10] M. Kowarschik and C. Weiß. DiMEPACK — A Cache-Optimized Multigrid Library. In H. Arabnia, editor, *Proc. of the Int. Conf. on Parallel and Distributed Processing Techniques and Applications (PDPTA 2001)*, pages 425–430. CSREA Press, June 2001.

[11] R. H. Otten, R. Camposano, and P. R. Groeneveld. Design automation for deepsubmicron: present and future. In *Design, Automation and Test in Europe (DATE)*, pages 650–657, 2002.

[12] C.-H. Tsai and S.-M. S. Kang. Cell-level placement for improving substrate thermal distribution. *IEEE Transactions on Computer-Aided Design of Circuits and Systems*, 19(2):253–266, Feb. 2000.

[13] W. Wang, A. Raghunathan, G. Lakshminarayana, and N. K. Jha. Input space adaptive design: A high-level methodology for energy and performance optimization. In *ACM/IEEE Design Automation Conference (DAC)*, page 738, 2001.