

M. I. Campbell
Graduate Student.

C. H. Amon
Professor.

J. Cagan
Associate Professor.

Department of Mechanical Engineering
and Engineering Design
Research Center,
Carnegie Mellon University,
Pittsburgh, PA 15213-3890

Optimal Three-Dimensional Placement of Heat Generating Electronic Components

This work introduces an algorithm that uses simulated annealing to perform electronic component layout while incorporating constraints related to thermal performance. A hierarchical heat transfer analysis is developed which is used in conjunction with the simulated annealing algorithm to produce final layout configurations that are densely packed and operate within specified temperature ranges. Examples of three-dimensional component placement test cases are presented including an application to embedded wearable computers.

1 Introduction

In the design of multicomponent systems, determining the best placement of components is a difficult task. Designers are challenged with finding solutions to large layout problems especially in computer hardware design where products can contain thousands of components. Issues such as wire routing and container geometry can be resolved by intelligently packing components. Placement of components within compact volumes also has a large impact on a product's success. Systems that are smaller are usually more cost effective and exhibit better performance and endurance. This is especially evident in the computer industry where compactness can often lower production costs and improve operating speeds.

There are, however, limitations as to how densely a system can be packed. In systems having a variety of power-consuming components, poor packaging can lead to high failure rates due to overheating and thermally induced stresses. With electronic components, the problems associated with heat transfer are amplified by the fact that technologically advanced components are more sensitive to high temperatures, despite the current trend to utilize low power CMOS multichip modules. By accounting for thermal performance issues early in the design of component layouts, the need for forced convection cooling might be avoided, the operating temperatures may be decreased, and the resultant system can be more compact.

The layout of electrical systems has traditionally been determined by functional and spatial requirements, independent of thermal constraints. Thermal constraint violations are avoided by modifying substrate materials, adding heat spreaders, optimizing the shape of cooling fins (Osio and Amon, 1996), and improving designs by performing system analyses. This type of thermal design and control has been widely studied as shown in Bar-Cohen (1994) and Agonafer and Linton (1994, 1995). There are, however, few techniques that take thermal issues into account during the placement of the electronic components.

In many problem-specific applications, rules of thumb have been developed for placing electronic components in locations that maximize their cooling. For example, research done by Dancer and Pecht (1989) and Osterman and Pecht (1989, 1990) present methods for placing components in one-dimensional rows according to the components' sensitivities to high temperatures.

Concurrent design methodologies present another approach for laying out large systems. One such concurrent thermal de-

sign methodology, presented by Nigen and Amon (1992), leads the thermal designer through five analysis stages of increasing accuracy during the layout process to arrive at an acceptable final layout. This methodology has been successfully applied to the design of compact electronic devices such as wearable computers (Amon et al., 1995).

Optimal layout of electronic components is subject to different constraints such as electrical, mechanical, thermal, ergonomics, and manufacturing. For example, electronic constraints include minimization of distance between components that interact more frequently to improve speed and transmission requirements. Mechanical constraints include minimization of fatigue induced failure of solder joints. Therefore, to improve layout configurations, optimization algorithms are required to place components in locations that allow for proper cooling. Queipo et al. (1994) use genetic algorithms to determine the location of electronic components using convection thermal analyses while Elias et al. (1990) use Monte Carlo thermal optimization to account for conduction through the substrate. These approaches have shown success in two dimensions but are limited by the discretization of the layout area into a grid of possible component locations. The more general layout problem requires a continuous formulation where components can be packed as close as possible, subject to thermal limitations.

Simulated annealing (Kirkpatrick et al., 1983) is a stochastic optimization algorithm that has demonstrated success in the compact placement of components under spatial constraints. In two-dimensional VLSI systems, simulated annealing has been successfully used to pack electronic components under routability constraints (Sechen and Sangiovanni-Vincentelli, 1985). In three-dimensions, Szykman and Cagan (1996) have used a simulated annealing based algorithm for placement of components under spatial constraints within a variety of container sizes. These algorithms have demonstrated success in compact placement of components but have not taken into account thermal design considerations.

This work presents a layout algorithm using simulated annealing that incorporates thermal design constraints in addition to placement of components in compact volumes. A hierarchical heat transfer analysis is introduced that is specifically tailored to the needs of the simulated annealing algorithm. The hierarchical approach improves the efficiency of the algorithm by selecting a thermal analysis method with a suitable degree of accuracy for different stages of the annealing process. The optimization algorithm, coupled with the hierarchical heat transfer analysis, generates compact and thermally acceptable three-dimensional electronic layouts, as illustrated by several test cases and a three-dimensional layout of a wearable computer. Future work will include electrical routing and other placement constraints.

Contributed by the Electrical and Electronic Packaging Division for publication in the JOURNAL OF ELECTRONIC PACKAGING. Revised manuscript received August 1, 1996. Associate Technical Editor: R. Schmidt.

2 Simulated Annealing Approach to Component Placement

Few optimization techniques address highly nonlinear, discontinuous, and multimodal problems such as component layout. Simulated annealing (Kirkpatrick et al., 1983) is a stochastic optimization technique developed to solve problems that exhibit these characteristics. Previous effort has applied simulated annealing to the spatial layout problem. This paper incorporates thermal performance issues into the simulated annealing environment. The following sections develop the thermally constrained layout objective function and present its implementation within the simulated annealing technique.

2.1 Establishing the Objective Function. The objective function is a numerical measure of the overall quality of a layout. Through a large number of iterations, simulated annealing finds layouts that are close to the global optimum. At each iteration, the algorithm evaluates both the design objectives, such as packing density and component temperatures, and the constraint violations, such as overlapping components, components that protrude from their specified container, and components whose temperatures are above critical temperatures for reliability considerations. The objective function is penalized when the design violates these constraints by moving into an infeasible region of the design space.

In this research, thermal performance issues are accounted for by including a coolest packing term and a critical temperature penalty term. The purpose of the coolest packing objective, f_{cool} , is to keep the overall component temperatures to a minimum. This coolest packing term, based on the sum of the squares of the fraction of component maximum temperatures to component critical temperatures, is given by:

$$f_{cool} = \sum_{i=1}^n \left(\frac{T_{i,max}}{T_{i,crit}} \right)^2 \quad (1)$$

where $T_{i,max}$ is the analytically determined maximum temperature of each individual component i ; $T_{i,crit}$ is the specified critical temperature each component cannot exceed; and n is the total number of components.

Although this objective term attempts to minimize the temperature of each component, it does not guarantee that components will not exceed their critical temperatures. Therefore, a penalty term is added to account for violations of these critical temperatures and is included only when a component exceeds its critical temperature (i.e., $T_{i,max} > T_{i,crit}$). This penalty term, p_{hot} , is equal to the sum of the squares of the maximum and critical temperature differences:

$$p_{hot} = \sum_{i=1}^n (T_{i,max} - T_{i,crit})^2 \quad (2)$$

If these new terms are combined with spatial terms, the objective function takes the form

$$F = W_1 f_{size} + W_2 f_{cool} + W_3 p_{overlap-component} + W_4 p_{overlap-container} + W_5 p_{hot} \quad (3)$$

where f_{size} is the objective term for maximizing packing density, f_{cool} is the objective term for minimizing the overall temperature, $p_{overlap-component}$ and $p_{overlap-container}$ are the penalties for component overlap with other components and with the container, and p_{hot} is the penalty for exceeding critical temperatures. The first, third, and fourth terms of the objective function, described in detail by Szykman and Cagan (1995), create layouts with high packing densities while avoiding overlaps, whereas the second and fifth terms attempt to find layouts with electronic components that have temperatures below critical values. Each objective and penalty term also has an associated weight (W_i , $i = 1, \dots, 5$). These weights normalize the terms as well as specify

their relative importance. The weights could be further resolved to adjust the importance of individual components in the thermal design objectives. Finding the global minimum of this function determines the best layout design that is compact, nonoverlapping, and safe from overheating.

2.2 Simulated Annealing Algorithm. The simulated annealing algorithm starts with an initial random layout configuration and makes numerous perturbations to the layout evaluating the objective function with each change. These modifications allow the system to move through the design space in search of optimal solutions. By statistically choosing to retain an old state or adopt a new state, the algorithm is able to converge on an optimum. The algorithm was developed as an analogy to the annealing of metals where an annealing process slowly reduces the temperature of the metal bringing it to a state of minimum energy. In a similar way, simulated annealing "anneals" the layout by bringing the objective function to a minimum. A parameter called the annealing temperature slowly decreases over the annealing process, affecting the probability of accepting or rejecting future design states.

The set of allowable perturbations comprises a move set. Since coordinates and orientations are the only variables being manipulated in the layout algorithm, the move set consists of possible translations and rotations of individual components and of swaps of the locations of two components. The current implementation restricts rotations to 90 degree increments.

To improve layout efficiency the algorithm dynamically adapts its annealing parameters based on its performance. The technique of Huang et al. (1986) is used to adjust the annealing temperature while the probability used to select moves is dynamically adjusted, as proposed by Hustin and Sangiovanni-Vincentelli (1987), based on a statistical analysis of its behavior throughout the process.

Simulated annealing has proven to find near-optimal layouts of systems with a variety of components; however, the algorithm has some potential drawbacks. The algorithm has several parameters that are set by the user; poor selection of these parameters can cause the process to converge on suboptimal solutions or continue for an unnecessary number of iterations. Since the algorithm relies heavily on random perturbations and statistical behavior, simulated annealing typically performs on the order of 100,000 iterations to arrive at good solutions. The optimization process is made costly if the evaluation of each modified layout is time consuming; therefore, the computational time of the evaluation must be kept to a minimum.

3 Thermal Analysis

3.1 Approach Overview. Because the simulated annealing process makes on the order of 100,000 changes to a layout, the thermal analysis calculations required at each perturbation must be kept to a minimum. More accurate and time consuming methods, such as those using CFD (Computational Fluid Dynamics) techniques, can take hours to run even for steady-state cases with complex geometries. If each iteration were to take an hour, the entire optimization would last for over eleven years! Because the thermal analysis must be time efficient, a compromise between accuracy and computational speed must be met.

When developing thermal analysis techniques for the simulated annealing process, several approximations must be made to create a workable thermally constrained layout algorithm. The layouts are analyzed at the asymptotically steady-state condition thereby eliminating more costly transient heat transfer calculations. Furthermore, analyzing a system that has come to steady-state best approximates the operating conditions under which the system is most likely to fail. Although the transient temperatures of the system are not accounted for in this analysis with the simulated annealing layout algorithm, the final design

is accurately examined using CFD techniques before building a prototype.

In the simulated annealing process, the high probability of accepting initially inferior layout configurations requires less accurate thermal evaluations for compared designs. As the process continues, more accurate evaluations are required. However, slight inaccuracies in the temperature field can still result in solutions that are close to the optimum. Final designs can remain valid by including safety factors prior to the optimization that account for the inaccuracies present in the analysis. Convection is not considered in these analyses except as a boundary condition. Therefore, within a system, only conduction is used as a means of transferring heat. As described above, at the end of the layout optimization process, the final design can be validated with a more accurate CFD analysis.

In performing the thermal analysis for the layout algorithm, a hierarchical approach is introduced that accesses three different heat transfer analyses, each having a different degree of accuracy as well as speed. The algorithm is made efficient by allowing the choice of the appropriate heat transfer analysis for each stage of the optimization. Although this hierarchical approach greatly reduces computational time, with two-dimensional problems taking as little as two hours, the total run time for three-dimensional problems can still exceed 30 hours on a DECstation 3000 Alpha. The following three sections describe and compare the thermal analysis techniques.

3.2 Resistance Network Method. The Resistance Network Method uses finite-difference-type elements to arrive at a temperature field for the entire system. By discretizing the domain into a series of rectangular elements, the temperatures at key points in the layout are determined.

Several steps are involved in analyzing layouts using this technique. First, a grid of nodes is established. Initially, nodes are placed at the centers of components where the majority of heat is produced and where the maximum temperatures are often located. For example, in a system consisting of two components with centers at (1, 1, 1) and (2, 2, 2), nodes are initially placed at these coordinates. These nodes represent the maximum component temperatures used to calculate the f_{cool} and p_{hot} terms of the objective function (Eq. 3). After the placement of these nodes, additional nodes are placed at points corresponding to all possible combinations of node coordinates. As in the example, nodes are set up at points (1, 1, 2), (1, 2, 1), (1, 2, 2), (2, 1, 1), (2, 1, 2), and (2, 2, 1). These nodes are added to the system to fill out an orthogonal grid of points. Boundary nodes are then added to the container walls at the projection of the interior nodes which include convective resistances to the ambient temperature.

Next, the resistances between all adjacent nodes are calculated. These resistances are functions of the conductivities of components and of the board, as well as the distances between nodes. These resistances are then stored in a matrix, \mathbf{R} , which is used to solve for the temperature of each node using the following equation:

$$[\mathbf{R}]\{\mathbf{T}\} = \{\mathbf{Q}\} \quad (4)$$

where \mathbf{Q} contains the heat sources and boundary conditions and \mathbf{T} contains the temperatures of all nodes. \mathbf{R} is an n by n matrix where n is the total number of nodes, and R_{ij} represents the resistance between node i and node j . \mathbf{T} is solved either by direct inversion of the matrix \mathbf{R} or by iteratively solving for the temperatures in \mathbf{T} . As with any finite-difference analysis, \mathbf{R} is a large banded matrix that often contains hundreds of elements. For this analysis the algorithm chooses to solve the matrix equation using either LU-Decomposition for banded systems or Gauss-Seidel iterative method (Chapra and Canala, 1988). The choice is made based on the magnitude of the perturbation performed by the simulated annealing algorithm. If the design to be analyzed is close to the previous design, the algorithm

chooses Gauss-Seidel and uses the previous temperatures as initial guesses to efficiently solve the slightly altered matrix. If large perturbations are the case, the algorithm resolves the matrix using LU-Decomposition.

3.3 Lumped System Method. The Resistance Network Method consumes a large amount of time when used continuously throughout the annealing process. Therefore, a Lumped System Method is implemented to quickly explore design alternatives in the early stages of the simulated annealing algorithm. This method considers the entire system as one large heat producing object that interacts with the ambient air. The lumped system is defined by the dimensions of the volume bounding all components, the average conductivity of the components and board, and the total heat generated by the electronic components.

The average resistance from the center of the volume to the ambient temperature is calculated by using the formulae for conductive and convective thermal resistances which are functions of the distance (L) between nodes, the cross-sectional area (A) the node acts over, and the conductivity (k) or the convection coefficient (h). In the calculation of the bounding volume resistance, this approach is extended to account for heat transfer out of all the container faces. New resistive elements are established based on the surface area (SA) and volume (V) of the container, given by

$$R = \frac{SA}{k(V)} + \frac{1}{h(SA)} \quad (5)$$

By simplifying the system to one node, no matrix solvers are needed. From this resistance model the average temperature of the system can be estimated from the following equation:

$$T_{avg} = K(T_{ambient} + Q_{tot} \cdot R). \quad (6)$$

Q_{tot} is the total heat produced by the system, $T_{ambient}$ is the ambient temperature around the system, and R is found using Eq. (5). To improve the accuracy of this estimated average temperature, T_{avg} is periodically calibrated with the average temperatures found by the Resistance Network Method using the correction factor K . This factor is the ratio of the average temperature of the Resistance Network Method temperature field divided by the average temperature of the Lumped System Method. By recalibrating approximately every hundred iterations, the Lumped System Method can provide useful information. Despite the assumptions made by this method, the results obtained can lead to an overall system shape and size with subsequent stages providing more detailed placement information.

3.4 Subspace Method. The Subspace Method is based on combining attributes of both the Resistance Network Method and the Lumped System Method. This analysis technique provides more information than the Lumped System Method but is not as time consuming as the Resistance Network Method. This method divides the system into equal subspaces just as the Resistance Network Method discretizes the system into nodes. Then, within each subspace the properties are averaged as in the Lumped System Method. The Subspace Method is useful in providing information about the distribution of components within the system and identifying subspaces with large fluxes and high temperatures.

The system is divided into a number of cube-like subspaces. Numerical experiments show that establishing between 100 and 200 cubes produces a good balance between accuracy and time. The properties within each subspace are averaged, and the resultant matrix is solved using LU-Decomposition. This method proves to be quicker than the Resistance Network Method because no time is required in setting up elements based on placement of components, and the matrix to be solved is much smaller. However, like the Lumped System Method, a correc-




Fig. 1 Symbolic representation of lumped system method, subspace method, and resistance network method

tion factor is used to calibrate this method with the Resistance Network Method. Figure 1 shows a symbolic sketch of the three heat transfer analysis techniques from least to most accurate.

4 Incorporating Thermal Results Into Simulated Annealing

The optimization technique and the heat transfer analyses have both been discussed, however, the combination of the two has yet to be addressed. In Section 2.1 the component temperatures were included in the objective function through the coolest packing objective and critical temperature penalty. In these heat transfer analyses, the temperature field, the average temperatures, or the average temperature of the system is calculated. From this data the maximum temperatures of each component are extracted or estimated. Then, the simulated annealing algorithm calculates the objective function value based on these component temperatures.

At each iteration throughout the process the annealer perturbs the design (i.e., alters the component layout), evaluates the objective function, and chooses to accept or reject the new design. After evaluating the placement objective, packing density, and overlaps, the algorithm calls the heat transfer analysis. This analysis begins by choosing which technique to utilize based on how much the process has changed from the initial state as well as the size of the last perturbation.

Early in the process, the annealer chooses the Lumped System Method. At this stage many large moves are being made and the placement of components often dictates the size of the bounding box. The calculation of the heat transfer-related objective function terms provides information about how feasible the current box size and shape are for the final design.

After approximately a third of the process is completed, the algorithm switches to the Subspace Method. The calculation of the heat transfer-related objective function terms now takes into account the distribution of components throughout the system along with the container size. Since the system is going to alter greatly from these early iterations, there is no advantage in performing an extensive thermal analysis of the system. However, an approximate analysis at these early stages in the anneal-

ing process prevents the system from moving toward infeasible regions of the design space. Although both methods rely on a full matrix calculation for recalibration, they are successful at supplying the annealer with the required information in as little time as possible.

Later in the annealing process, the exact placement of components and the container size strongly dictates the heat transferring capabilities of the system, and consequently the component temperatures. Therefore, the algorithm switches to the more accurate Resistance Network Method when the annealing temperature reaches a specified value. Due to the different levels of accuracy, the three methods cannot be used interchangeably from iteration to iteration. When the process switches to a more accurate method, the acceptance and rejection decision of the annealing process is affected; switching back to less accurate methods would cause problems for the final layout. Therefore, upon switching to the Resistance Network Method the full matrix calculation is required at every iteration to maintain consistency of accuracy. The increase in accuracy greatly increases the calculation time of the thermal analysis. Therefore, judicious choice of the thermal analysis type substantially improves the speed of the process.

5 Results

5.1 Test Cases. The thermally constrained layout optimization algorithm is based on the aforementioned simulated annealing approach to three-dimensional placement introduced by Szykman and Cagan (1995). The algorithm is written in C and runs on a DECstation 3000 Alpha. A variety of tests were performed on the algorithm for various packing situations. Two-dimensional examples are shown in Campbell et al. (1995).

Simplified three-dimensional test cases are first examined to demonstrate the algorithm's operation and effectiveness. For this test case we use a system consisting of eight identical cubes all with equal material properties, volumetric heat generation, and temperature sensitivities. The results of the layout algorithm are compared to the predicted optimum. In either one, two, or three-dimensions, identical components are optimally placed in a line configuration when subject to a minimum volume objective and a coolest packing objective. Placing components in a line maximizes the surface area to volume ratio of the system providing a large area to dissipate heat to the environment in a small volume. To perform thermal analyses and spatial calculations, the input provided to the algorithm contains data on the geometric dimensions, thermal conductivities, heat fluxes, and critical temperatures of each electronic component as shown in the caption of Fig. 2.

The results of the layout algorithm are shown in Fig. 2 with the eight components packed in an unconstrained container, a




Fig. 2 Optimized layouts for three-dimensional test cases. Components are eight identical cubes with length = 2 cm, T_{crit} (C) = 70.0°C, $Q = 0.3$ W, $k = 0.0003$ W/cm·C: (a) unspecified container; (b) 10 × 10 (cm) container; (c) 7.5 × 7.5 × 7.5 (cm) container.

10 × 10 × 10 (cm) container, and a 7.5 × 7.5 × 7.5 (cm) container. The dots in the figures correspond to the temperature nodes of the Resistance Network Method. In the unconstrained example (Fig. 2a) the layout compares well to the predicted line configuration despite an initial random three-dimensional arrangement. In Figure 2b, a cube container is imposed which constrains the layout from producing the line configuration. The algorithm produces a solution which still maximizes the surface area to volume ratio by creating a long planar configuration within the constrained space. In Fig. 2c, the components are so constrained that the layout is forced to spread out within the container to prevent high temperatures, while still conforming to the container dimensions. These examples reveal the difficulties of designing layouts that exhibit both minimum volumes and coolest overall temperatures. Next, we address a system of components having unequal heat dissipation rates and temperature sensitivities; this results in a layout problem that is more complex, as shown by Dancer and Pecht (1989).

5.2 Application to Embedded Wearable Computers. In this section the layout algorithm is applied to the design of wearable and mobile computer systems. Wearable computers often involve more severe mechanical, physical, and thermal constraints than other electronic devices (Amon et al., 1995). It is often necessary to make these computers compact, lightweight, durable, and shielded from environmental conditions. This often limits the heat transferring abilities of electronic components since the possibility of cooling with fans or through vents is eliminated. Therefore, intelligent layout of components and use of heat spreaders dictates how well a system will dissipate heat. Also, the six to twelve month production cycles of most wearable computers makes the use of design tools a necessity in creating robust systems. Finally, new technologies in material deposition allow electronic devices to be embedded in a solid substrate allowing components to be placed in three dimensions. The three-dimensional placement as well as the conduction-only cooling of embedded devices makes this algorithm a suitable design tool for wearable computers.

The Vu-Man 3R wearable computer, designed at the Engineering Design Research Center and manufactured at the Shape Deposition Laboratory of Carnegie Mellon University, is shown in Fig. 3. The Vu-Man 3R is a small state-of-the-art augmented reality device worn at the hip for maintenance applications. Size, durability, and isolation from the environment makes laying out the components within this device a difficult task. Although still using printed circuit boards, the components are completely embedded in a polyurethane-nylon resin (Egan and Amon, 1996). The original placement of components within the system was performed using traditional methods; the thermally constrained layout algorithm is used to generate possible layout alternatives for comparison. This is a more complex layout




Fig. 3 Vu-Man 3R wearable computer

Table 1 Data on Vu-Man 3R wearable computer

Component	Dimensions (cm)			T_{crit} (C°)	Q (W)	k (W/cm·C)
CPU Intel 84 pin	0.8	3.6	3.6	70.0	0.54	0.064
EPROM(128Kx8)	0.45	1.7	1.95	70.0	0.15	0.25
Voltage Regulator	2.3	2.25	0.6	60.0	0.15	0.118
Altera 25rs PLD	0.8	3.6	3.6	70.0	0.43	0.039
SRAM TC551001	2.0	1.05	0.25	70.0	0.4	0.082
SRAM TC551001	2.0	1.05	0.25	70.0	0.4	0.082
AA Battery 1	1.47	1.47	3.00	100.0	0.00	0.01
AA Battery 2	1.47	1.47	3.00	100.0	0.00	0.01
AA Battery 3	1.47	1.47	3.00	100.0	0.00	0.01
AA Battery 4	1.47	1.47	3.00	100.0	0.00	0.01

situation in which the electronic components differ in size, thermal sensitivity, heat generation, and thermal conductivity as it is shown in Table 1. Figure 4 depicts the original placement of the ten components of the wearable computer within a 12.0 × 4.14 × 12.7 (cm) container.

First, the algorithm is used to find the optimal coolest packing of components within the same container size of the original design. The data of the components shown in Table 1 is given to the algorithm. Figure 5 shows the optimized coolest layout generated by the algorithm. Components with higher heat generation are placed far from each other and on the sides of the container, allowing for more efficient heat dissipation to the surroundings. This placement enables the overall average temperature to be kept to a minimum. It can also be seen that the four AA batteries which dissipate no heat and are fairly insensitive to temperature are arbitrarily placed towards the middle of the container. They have little effect on the coolest packing objective and therefore we constrain them to be placed together




Fig. 4 Original layout of Vu-Man 3R wearable computer and temperature field across PLD and CPU component plane




Fig. 5 Optimized layout for coolest packing and temperature fields on the side faces

for ease of accessibility. The placement algorithm can be expanded to address issues of accessibility and other placement constraints as it is shown in Szykman and Cagan (1996).

The temperature distribution of both layouts, the one generated with the thermally constrained layout algorithm in Fig. 5 and the original design in Fig. 4, are next obtained numerically with a spectral element CFD technique (Patera, 1984; Amon, 1993). The fishnet temperature grids included in these figures depict the temperature profiles through critical planes in the layouts as obtained by the direct numerical simulations. The first two columns of Table 2 compare the CFD analysis of the original layout to the newly generated optimized layout. All maximum component temperatures in the original layout are larger than those of the optimized layout. It is important to note that with more components or higher heat fluxes, these differences would have been greater.

The temperatures predicted by the resistance network method in the third and fourth columns of Table 2 can be compared to the temperatures predicted by the CFD analysis. There is a 14 percent average error between the CFD analysis and the resistance network method; however, in most cases the resistance network conservatively over-approximates the temperatures.

Table 2 Comparison of maximum temperatures ($^{\circ}\text{C}$) for components of original layout and optimized layout of the Vu-Man 3R

Component	CFD Analysis		Final Resistance Method Approximation	
	Original Layout	Optimized Layout	Original Layout	Optimized Layout
CPU Intel 84 pin	39	38	50	40
EPROM(128Kx8)	37	33	45	34
Voltage Regulator	35	31	38	32
Altera 25rs PLD	38	37	46	38
SRAM TC551001	41	36	63	42
SRAM TC551001	41	36	66	44
AA Battery 1	32	30	31	31
AA Battery 2	32	30	32	31
AA Battery 3	32	30	30	31
AA Battery 4	32	30	30	31

Despite the approximations of the resistance network method, information about differences in layout quality is enough to produce noticeable improvements in final designs, as validated by a thorough CFD technique. In addition, the resistance network method can analyze one layout in less than a second while the CFD technique can take several hours.

The layout algorithm is then used to maximize the packing density of the system while satisfying the temperature constraints. This problem is the same as above except it is now independent of container size and constrained by the critical temperatures of components. Figure 6 shows a new layout for the wearable computer that is $5.10 \times 5.10 \times 3.15$ (cm)—less than one-seventh the volume of the actual Vu-Man 3R. In this layout, components are operating close to their critical temperatures due to the high packing density. The layout again shows that the components with higher heat generation are placed at the corners to promote better cooling and heat transfer to the environment—thereby resulting in a higher packing density. Here, the batteries are unconstrained to enable a maximization of the packing density.

6 Conclusions

This paper introduces a layout algorithm using a hierarchical heat transfer analysis and simulated annealing to address concurrent design issues of compact electronic devices. Problems such as overheating often develop when densely packing components within small containers. Therefore, by concurrently analyzing the thermal performance of designs throughout the layout optimization, high packing densities can be maintained while preventing overheating.

Incorporating heat transfer analyses into a simulated annealing layout algorithm often requires large computational times. By developing a hierarchical heat transfer analysis approach which allows for quicker evaluations of the objective function, computational time is greatly reduced enabling the optimization algorithm to produce near-optimal layouts.

Several three-dimensional test cases demonstrate the algorithm's ability to find optimally directed solutions. Applications to wearable computers are model candidates for study since container size and cooling methods often make component layout challenging. The proposed algorithm finds an optimally directed solution for three-dimensional component placement that proves to be cooler than the original layout for electronic components with unequal heat dissipation rates, conductivities, and




Fig. 6 Optimized layout for dense packaging under thermal constraints as compared to the original layout of the system

temperature sensitivities. For a conduction internal cooling environment, the thermal analysis of the layout algorithm provides conservative approximations of the actual temperatures of components in a fraction of the computational time. The layout algorithm is also able to propose a more compact layout of the system that is one-seventh of the volume of the original layout, while avoiding problems of overheating.

Acknowledgments

The authors would like to thank Simon Szykman for his advice and assistance throughout this project, Eric Egan and Kyongsoo Chae for their help on the heat transfer simulations, and Lee Weiss for the manufacturing of the embedded wearable computer. The authors would also like to acknowledge the financial support of the National Science Foundation under grants MIP-9420372 CTS-9630801 and EID-9256665, and the Engineering Design Research Center (EDRC), an Engineering Research Center of the National Science Foundation at Carnegie Mellon University, for supporting this work under Grant EEC-8943164.

References

Agonafer, D., and Linton, R. L., 1994, "Solution of a System-Level Benchmark Problem in Electronic Cooling," ASME 94-WA/HTD-Vol. 292, New York, NY.

- Agonafer, D., and Linton, R. L., 1995, "Coarse and Detailed CFD Modeling of a Finned Heat Sink," *IEEE Transactions on Components, Packaging, and Manufacturing Technology*, Vol. 18, pp. 517-520.
- Amon, C. H., 1993, "Spectral Element-Fourier Method for Transitional Flows in Complex Geometries," *AIAA Journal*, Vol. 31(1), pp. 42-48.
- Amon, C. H., Nigen, J. S., Siewiorek, D. P., Smailagic, A., and Stivoric, J., 1995, "Concurrent Design and Analysis of the Navigator Wearable Computer System: The Thermal Perspective," *IEEE Transactions on Components, Packaging, and Manufacturing Technology*, Vol. 18, No. 3, pp. 567-577.
- Bar-Cohen, A., 1994, "Thermal Design and Control," *Physical Architecture of VLSI Systems*, R. Hannemann et al., eds., John Wiley and Sons, Inc., New York, NY, pp. 541-605.
- Campbell, M. I., Amon, C. H., Cagan, J., and Szykman, S., 1995, "Electronic Component Placement Using Simulated Annealing Under Thermal Constraints," *Proceedings of the 1995 International Mechanical Engineering Congress and Exposition*, ASME, HTD-Vol. 319/EEP-Vol. 15, New York, NY, pp. 155-162.
- Chapra, S. C., and Canala, R. P., 1988, *Numerical Methods for Engineers*, McGraw-Hill, Inc., New York.
- Dancer, D., and Pecht, M., 1989, "Component-Placement Optimization for Convectively Cooled Electronics," *IEEE Transactions on Reliability*, Vol. 38, No. 2, pp. 199-205.
- Egan, E., and Amon, C. H., 1996, "Cooling Strategies for Embedded Electronic Components of Wearable Computers Fabricated by Shape Deposition Manufacturing," *ITHERM '96, Fifth ASME/IEEE/ISHM/CPMT/NIST Conference on Thermal Phenomena in Electronic Systems*, pp. 13-20.
- Elias, R., Elperin, T., and Bar-Cohen, A., 1990, "Monte Carlo Thermal Optimization of Populated Printed Circuit Board," *IEEE Transactions on Component Hybrids, and Manufacturing Technology*, Vol. 13, No. 4, pp. 953-960.
- Huang, M. D., Romeo, F., and Sangiovanni-Vincentelli, A., 1986, "An Efficient General Cooling Schedule for Simulated Annealing," *ICCAD-86: IEEE Interna-*

tional Conference on Computer-Aided Design—Digest of Technical Papers, Hoboken, NJ, pp. 381–384.

Hustin, S., and Sangiovanni-Vincentelli, A., 1987 "TIM, a New Standard Cell Placement Program Based on the Simulated Annealing Algorithm," *IEEE Physical Design Workshop on Placement and Floorplanning*, Hoboken, NJ.

Kirkpatrick, S., Gelatt, Jr., C. D., and Vecchi, M. P., 1983, "Optimization by Simulated Annealing," *Science*, Vol. 220, No. 4598, pp. 671–679.

Nigen, J. S., and Amon, C. H., 1992, "Concurrent Thermal Designs of PCB's: Balancing Accuracy With Time Constraints," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. 15, No. 5, pp. 850–859.

Osio, I. G., and Amon, C. H., 1996, "An Engineering Design Methodology with Bayesian Surrogates and Optimal Sampling," *J. Research in Engineering Design*, Vol. 8, pp. 184–206.

Osterman, M. D., and Pecht, M., 1989, "Component Placement for Reliability on Conductively Cooled Printed Wiring Boards," *ASME JOURNAL OF ELECTRONIC PACKAGING*, Vol. 111, pp. 149–155.

Osterman, M. D., and Pecht, M., 1990, "Placement for Reliability and Routability of Convectively Cooled PWB's," *IEEE Transactions on Computer-Aided Design*, Vol. 9, No. 7, pp. 734–744.

Patera, A. T., 1984, "A Spectral Element Method for Fluid Dynamics: Laminar Flow in a Channel Expansion," *Journal of Computational Physics*, Vol. 54, pp. 468–477.

Queipo, N., Devarakonda, R., and Humphrey, J. A. C., 1994, "Genetic Algorithms for Thermosciences Research: Application to the Optimized Cooling of Electronic Components," *International Journal of Heat and Mass Transfer*, Vol. 37, No. 6, pp. 893–908.

Sechen, C., and Sangiovanni-Vincentelli, A., 1985, "The TimberWolf Placement and Routing Package," *IEEE Journal of Solid-State Circuits*, Vol. 20, No. 2, pp. 510–522.

Szykman, S., and Cagan, J., 1995, "A Simulated Annealing Approach to Three-Dimensional Component Packing," *ASME Journal of Mechanical Design*, Vol. 117, No. 2(A), pp. 315–321.

Szykman, S., and J. Cagan, 1996, "Constrained Three-Dimensional Component Layout Using Simulated Annealing," *ASME Journal of Mechanical Design*, in press.