AN INVESTIGATION OF PWB LAYOUT BY GENETIC ALGORITHMS TO MAXIMIZE FATIGUE LIFE

Andrew J. Scholand
Robert E. Fulton
Advanced Electronic Packaging Laboratory
Manufacturing Research Center
Georgia Institute of Technology
Atlanta, Georgia

Bert Bras
System Realization Laboratory
Georgia Institute of Technology
Atlanta, Georgia

ABSTRACT

Thermal considerations in Printed Wiring Board (PWB) assemblies are becoming increasingly important as packaging constraints shrink and power use escalates. In this paper, we provide a study on the potential for a Genetic Algorithm-driven PWB layout design tool to improve the thermal performance of such assemblies. As a case study, the thermo-mechanical fatigue of surface mounted leadless chip carriers on an FR4 epoxy board is used. We have found that by utilizing appropriate formula-based engineering approximations, the efficiency of Genetic Algorithms in finding near-optimal and optimal results makes this approach effective as an explorative 'scouting' approach to identify promising board configurations for more computationally expensive evaluations such as finite element method.

1. INTRODUCTION

One of the tasks a designer of a Printed Wiring Board (PWB) faces is placement of electrical components on the board. There are many constraints on possible locations, primarily determined by the electrical circuit design of the board. However, increasing customer demands for more processing power in smaller packages mean that many PWBs are operating at temperatures significantly above ambient (75 - 90°C). This presents a problem for surface mounted components, as the solder connections between the board and the component are stressed by the thermal expansion differences between these two bodies. Repetitive use (i.e. power on, power off cycles) leads to fatigue failure, a phenomena referred to as solder joint thermo-mechanical fatigue. To maximize the fatigue life of a particular surface mounted (SM) component, then, it is desirable that the component be located in a low temperature region of the board. However, since some SM components are more susceptible to fatigue damage than others, and since the temperature gradients across the board may vary as illustrated below, there exists an ordering of components in available locations which will maximize the fatigue life of the entire PWB. (Note that any component failure anywhere on the board constitutes a PWB failure).

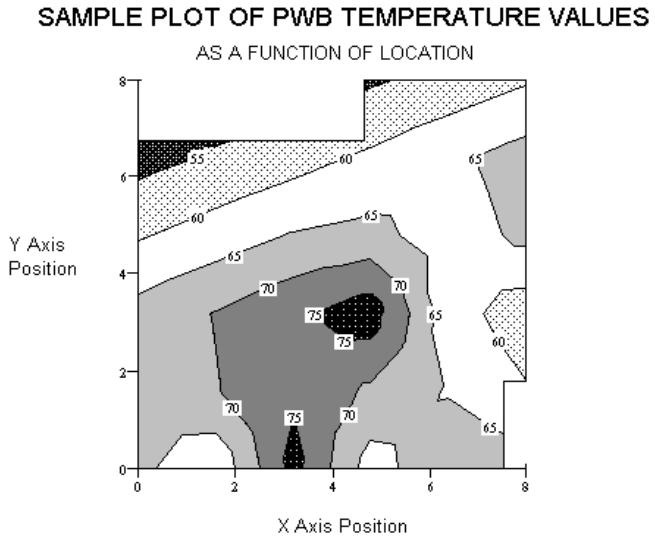

Figure 1-1

Electrical and other constraints on the placement of the components are presumed to limit the number of available positions to a finite number. There are obviously only a finite number of components to place in some or all of these available

slots. Thus, the problem is a combinatorial type of problem, well suited to solution using a genetic algorithm. This article will first provide a brief background in genetic algorithms and solder joint thermo-mechanical fatigue, before presenting the preliminary results obtained by a genetic algorithm approach. Further work is planned to integrate this approach into sophisticated computer based thermal analysis tools.

2. GENETIC ALGORITHM BACKGROUND

2.1 Definition of Genetic Algorithms

Genetic algorithms (GAs) represent a class of adaptive search techniques that combine a Darwinian survival-of-the-fittest among discrete string structures (artificial chromosomes representing possible problem solutions) and a structured, yet randomized, information exchange among these structures. This sets them apart from typical optimization and search processes in the following ways [Goldberg, 1989]:

1. GAs work with a coding of the parameter set, not the parameters themselves.
2. GAs search from a population of points, not a single point.
3. GAs use an evaluation function to determine 'payoff' in directing the search, rather than derivatives or other auxiliary knowledge.
4. GAs use probabilistic transition rules, not deterministic rules.

The search and information exchange functions within a GA are carried out by the following:

GA = Reproduction + Mutation + Selection

The algorithm typically works on a constant-size population of structures (also called strings, or chromosomes). Some portion of that population 'reproduce' - or create new structures based on their own structure. This reproduction is performed through crossover operators that assimilate the characteristics of the parents via different algorithms. Mutation (and possibly inversion, a shuffling operation) operators are also used to increase the diversity of the population, and to avoid premature convergence at local optima. The selection process then removes the least fit members of the population, leaving the better structures to exchange information.

This process is essentially a simultaneous optimization of a large population of configurations, which makes the genetic algorithm an extremely efficient optimizer. The search progresses largely in terms of groups of genes which confer advantageous traits, alternately termed co-adapted genes, schema, or hyperplanes. Because the groups of genes which remain in the population do so only because their parents (and earlier ancestors) were 'fit', GAs bias the makeup of new individuals in response to feedback on the fitness of previously generated individuals. Thus GA's exploit knowledge accumulating over time (generations) about hyperplanes within the search space.

GA's are unconstrained search procedures. Constraints can be handled through penalty functions, post recombination/mutation operators which change any constraint-violating structures into valid arrangements, or by crossover operators which only produce valid structures.

There are a variety of terms which are freely used in the genetic algorithm arena, combining both biological terms and computer science terms. The reader not familiar with Genetic Algorithms is encouraged to review one of the very good introductory texts (such as [Goldberg, 1989] or [Davis, 1991]) before attempting to proceed further, as the remainder of this paper assumes a basic level of familiarity with GA jargon.

2.2 Advantages of GAs

Much of the interest in genetic algorithms is due to the fact that they provide a set of efficient search heuristics without the need for incorporating highly domain-specific knowledge. This efficiency stems from the use of selective reproduction and recombination of structures to adjust the area of searching (hyperplanes) according to the average fitness of strings in those hyperplanes. The search for a solution is therefore carried out in the most promising areas of the state space. Because genetic algorithms need not search along the contours of the function being optimized, they tend not to become trapped in local minima (although this is somewhat dependent on the values selected for the GA control parameters). And, as mentioned in the previous section, the recombination of previous solutions means that the algorithm acts as if it had 'learned' from its past states.

The genetic algorithm approach is 'inherently parallel' [Holland, 1975] in that many solution candidates (the multiple hyperplanes within each member of the population) are all interacting at once. GAs are also very well suited to implementation on parallel processing hardware due to the simultaneous evaluation of the entire population. When implemented on parallel hardware, even inefficiently tuned GAs are extremely rapid solution finders.

Finally, "In addition, a number of experimental studies show that GA's exhibit impressive efficiency in practice. While classical gradient search techniques are more efficient for problems which satisfy tight constraints, GA's consistently outperform both gradient techniques and various forms of random search on more difficult (and more common) problems, such as optimizations involving discontinuous, noisy, high-dimensional, and multimodal objective functions." [Grefenstette, 1986]

3. LITERATURE REVIEW

3.1 General Genetic Algorithm Research

John Holland [Holland, 1975] is generally considered to be the pioneer in the field of genetic algorithms in the sense that he was the first to set up a mathematical framework for solving generic problems using methods extracted from biological processes. In addition to defining the terminology explained in the Background section, he also originated much of the theoretical understanding of Genetic Algorithms. A more recent general text in the field has been written by one of John Holland's early students, David Goldberg [Goldberg, 1989]. Similarly, [Davis, 1991] provides an

introductory 'tutorial' on GAs, and then presents chapter-long summaries on current GA applied research.

The computational efficiency of Genetic Algorithms suffers from sensitivity to values used for the tuning parameters. (Although a standard set of parameters do provide a fairly good search for a variety of problems.) These parameters are also quite different than the expected parameters for a given problem (i.e. mutation rate, at first brush, does not seem related to a search for an optimal spring, for example). [Grefenstette, 1986] identified values for Population Size, Crossover Rate, Mutation Rate, Generation Gap, a response surface scaling parameter, and Selection Strategies to produce efficient GAs from a search space of 2^{18} possible GAs for a set of numerical optimization problems (specifically, parabola, Rosenbrock's saddle, step function, quartic with noise, and Shekel's foxholes). Grefenstette concludes that the performance of GAs is a nonlinear function of the control parameters, and discontinuities and local optima cannot be ruled out. However, he also has demonstrated that very good performance can be obtained with a range of control settings.

3.2 Genetic Algorithms Applied to Electronic Component Placement

Genetic algorithms are an extremely topical research topic, and several authors have studied their use for electronic board design. Typically, early configuration design is the focus of these GA design tools, where the main design problem is the location of subassemblies on a larger assembly. These electronic component placement problems have been considered by a number of researchers. Most of the papers ([Cohoon and Paris, 1986], [Shahookar and Mazumder, 1990], [Chan, Heming et al., 1991]) examine placement of modules on a VLSI chip to minimize chip area and interconnect lengths. [Jin and Chan, 1992] extended the same performance metrics to discrete analog (through hole) components.

The increasing importance of good thermal design dictates high rewards for integrating thermal considerations into the early stages of the design cycle these GA-based tools address. Surprisingly, however, no-one has yet incorporated thermal performance criteria into the design systems described in the literature.

An application of genetic algorithms to a related area was published by [Wong and Leu, 1993]. These authors looked at the planning of assembly operations (placement/insertion sequence and machine setup) on printed circuit boards to minimize assembly time. One of the advantages of this approach is its general nature - it is applicable across a wide variety of assembly machines (subject to a taxonomy developed by the authors), whereas traditional approaches had been very machine specific.

4. PWB FATIGUE PROBLEM BACKGROUND

4.1 Thermo-Mechanical Solder Joint Fatigue

Surface Mount Technology (SMT), the practice of attaching electrical components to a substrate via a lap solder joint, has moved from the military/aerospace applications of the 1960s to widespread commercial use [Capillo, 1990]. The main drivers for this migration away from conventional through-hole mounting are higher packing densities (fifty to seventy percent higher than possible with conventional designs), better electrical performance, and much more rapid manufacturing processes. However, the advantages of higher packaging densities, miniaturization, and faster circuit speeds, result in elevated operating temperatures. Many commercial applications of SMT utilize components and substrates with large differences in their coefficients of thermal expansion.[1] As mentioned in the introduction, this presents a problem for surface mounted components, as the solder connections between the board and the component are more compliant than the component or the board, and so are stressed in shear by the differences in extension between these two bodies. Repetitive use (i.e. power on, power off cycles) leads to fatigue failure by cracking, a phenomena referred to as solder joint thermo-mechanical fatigue.

Not all components are equally susceptible to fatigue failure, however. Obviously, the thermal expansion equation dictates that the extension is small if the length is small:

$$\Delta L = \alpha * \Delta T * L \qquad (4\text{-}1)$$

where ΔL = Thermal expansion
ΔT = Change in temperature
α = Coefficient of thermal expansion
L = Length of body undergoing expansion

Thus, the solder stresses due to CTE differences are small if the component's area is small. Small SMT components such as chip resistors and capacitors, therefore, seldom experience fatigue failure in practice. The differential expansion between the board and component can be further exacerbated, however, when the two bodies are at different temperatures, as can happen when power is dissipated within the component. The most vulnerable components, in terms of solder joint reliability, are typically the hermetic leadless chip carriers (LCCCs), which usually enclose and protect silicon VLSI chips. These LCCCs are subject to both pressures toward fatigue loading- they are large relative to resistor chips (LCCCs can be as large as 25 mm square, versus the 1.3 mm square typical chip resistor), and they can operate at temperatures substantially different from board temperatures.

Although a description of the underlying phenomena of solder joint fatigue failure makes the problem appear simple, designing a solution is not a trivial matter. Printed circuit boards which more closely match the CTEs of the SMT components are expensive (typically they involve Kevlar composites or copper-Invar-copper restraining cores). Many factors affect the observed fatigue life of a component, including the volume and geometry of the solder joint (and more is not always better, as excess solder makes the joint stiff and brittle), the positioning of the component over the copper traces on the PWB (which in turn affects the solder joint geometry), the formation of brittle intermetallic compounds over time, and material properties of the

[1] The ceramic cores of most SMT components are either beryllium oxide or aluminum oxide, which have CTEs in the range of 5 to 7 ppm/°C, whereas the typical low cost printed circuit board material is epoxy fiberglass laminate, which has a CTE of 12 to 16 ppm/°C.

solder. The section below describes some of the engineering models which have been proposed to facilitate robust design.

4.2 Modeling Solder Joint Fatigue

As described above, the functional reliability of surface mount technology is a complex issue involving many not well understood components. A detailed analysis is typically only possible using finite element methods. For example, [Akay et al., 1992] used FEM to analyze factors affecting the prediction of solder joint cycle life in surface-mount assemblies. He studied several parameters including initial temperature, temperature ramp time and hold time, solder grain size, and constitutive equations for solder alloys. His results indicate that creep and plasticity of the solder joints have a substantial effect on thermal fatigue life. [Ozmat, 1990] reached similar conclusions with two dimensional finite element models. However, in the design of a typical board, it is not possible to analyze each and every solder joint, and furthermore at the initial design stage, many of the details regarding the solder joint may not be known.

Robust design can only be assured, however, by analyzing large numbers of solder joints for each possible board configuration, and the number of board configurations to be considered will also be large. Therefore, a relatively simple predictive model is required to facilitate preliminary design tasks such as component placement while accounting for thermal fatigue issues. One such model is the simple Engelmaier extensional model [Engelmaier, 1983].

This model takes a purely phenomenological approach and relegates second-order effects to a lumped empirical figure. The fatigue model assumes two rods of homogenous material (typically the composite FR4 for the PWB and alumina for the chip component), and homogenous, isotropic, eutectic Pb-Sn solder shear bodies.

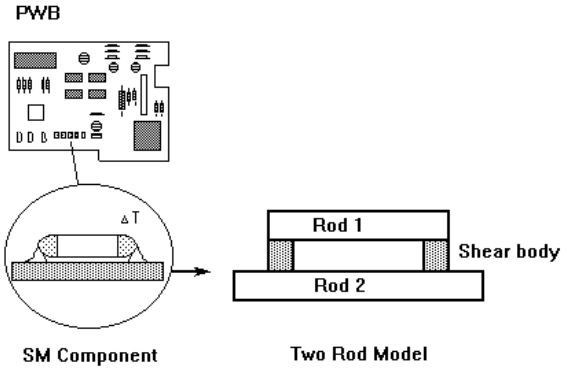

Figure 4.2-1: Modeling a component as a Two Rod Model

The shear strain is calculated from steady-state temperature and TCE (thermal coefficient of expansion) differentials between the package and interconnect board without consideration of elastic deformations. This value represents the average plastic shear strain which is converted into solder joint cycle life. These calculations are illustrated below (adapted from [Engelmaier, 1983]).

Description	Equation
Manson-Coffin Fatigue Model for number of cycles to failure	$N_f = \frac{1}{2}\left\{\frac{\Delta\varepsilon}{2\varepsilon_f'}\right\}^{\frac{1}{c}}$
Material constant (fatigue ductility exponent) for 60 Sn-40 Pb Solder	$c = -0.442 - 0.0006T_{sj} + 0.0174\ln(1+f)$
Approximate temperature of the solder joint	$T_{sj} = \frac{1}{4}(2T_o + T_s + T_c)$
Material constant (fatigue ductility coefficient) for 60 Sn-40 Pb Solder	$\varepsilon_f' = 0.325$
Corrected shear strain	$\Delta\varepsilon = \Delta\gamma = F\lvert\gamma\rvert$
Raw shear strain	$\gamma = \frac{L\Delta(\alpha\Delta T)}{2h}$
Expansion mismatch due to differences in CTEs and body temperatures	$\Delta(\alpha\Delta T) = \alpha_s(T_s - T_o) - \alpha_c(T_c - T_o)$

Table 4.2.1: Engelmaier Thermo-mechanical Fatigue Model

Engelmaier claims the assumption that all shear deformation is plastic is valid because, " ...the Manson-Coffin plot for solder is completely dominated by plastic deformation and does not have any significant elastic strain component...". However, it is important to note that the lumped empirical figure of merit (*F* in the above table) is valid only for an approximate temperature range.

Since the model can only be assumed to be close to experimental results in the neighborhood where the 'tweaking' parameter was set, the range of board temperatures will be kept within the range 85°C to 92°C, with the ambient temperature kept fixed at 20°C.

4.3 Modeling Board/Component Temperatures

The temperature profiles, both transient and steady state, that develop across printed circuit wiring boards can be extremely complex. They are affected by physical conditions such as internal cyclic power dissipation, external thermal cycling of forced air stream, thermal gradients with mounting enclosures, mounting hardware, etc. Most commercial packages use finite element or finite difference methods to calculate the temperature profile of a board under 'worst case' conditions. An example of such a commercial package is AutoTherm by Mentor Graphics.

Because the temperature distribution on a PWB is determined in part by the location of the heat-dissipating chips on the board, ideally each new placement of components should be analyzed using a package such as AutoTherm. Genetic Algorithms require thousands, or even tens of thousands, of such evaluations, however. The computational expense of performing the number of calculations required effectively rules out this option. An

alternate approach is to identify similar groupings of chip configurations, and have separate thermal analyzes run for each configuration. The analysis of any individual layout would then consist of determining which category or group of layouts it belonged to, and then applying the thermal profile corresponding to that group. The simplest implementation of this approach is to have only one group, of course, and this approach has been adopted initially.

The changes in the board's temperature distribution brought about by the placement of chips on the board was accounted for in a primitive way. This approach started with an initial temperature distribution across the circuit board, which was then modified depending on the locations of the LCCC chips- each chip making its immediate environment a few degrees warmer. This first order model (described in greater detail in Section 6.2 below) represents some of the component placement / board temperature interactions which characterize real systems, without the computational expense of finite element methods.

The component temperatures are modeled for the purposes of their contribution to the plastic strain of the solder as fixed values. Therefore, the transient strains caused by the component and circuit board reaching their steady state values are lumped into the correction factor *F*. Engelmaier claims that for the high CTE mismatch case of ceramic components on FR4 substrates, these transients are not significant contributors to the overall strain. (He suggests a correction factor between 1.2 to 1.4 to account for both transient strains and board warpage effects.)

The actual component temperatures used come from Engelmaier's 1983 paper and [Peak, 1992]. Engelmaier's LCCC temperature of 96°C assumes a ceramic chip carrier, 16.5 mm square, dissipating 1.25 Watts, with forced convection of 65°C air at 1.5 m/s.

5. PRELIMINARY COMPUTER IMPLEMENTATION

5.1 General Approach

The models described above were implemented in the GENEtic Search Implementation System (GENESIS), version 5.0 ([Grefenstette, 1990]). GENESIS is a genetic algorithm for function optimization available in the public domain. It is written in the programming language C, and requires only that the user write a problem specific evaluation function. GENESIS provides several advanced GA features including scaling, two point crossover, the option to heuristically chose the initial population, schema tracking, and an optional elitist selection strategy. This code was compiled and run on a Unix SparcServer 1000 using Sun OS 5.

The value of the evaluation function for each member of the population (a specific component layout on the circuit board) will be determined by finding the minimum fatigue life of all the placed components. A thermal profile of the board will allow determination of the temperature at each position, initially by using a look up table. A second, more sophisticated algorithm will determine the temperature at each position based on the initial board temperature profile and the location of the chips themselves. It should be pointed out that this evaluation function does not provide much information on the full ordering of the components - in effect, it only provides information about the worst placed component, and it does not indicate which component is the worst component. This did not prove to be a problem in practice, as the results below will indicate.

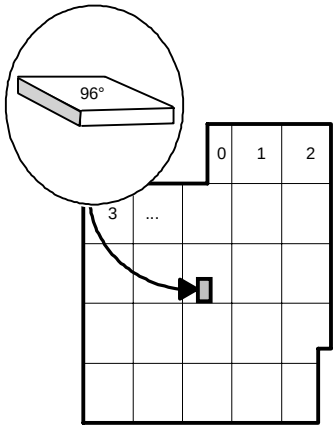

Figure 5.1-1: Identifying chip placement locations

This implementation provides 32 available 'slots' for the chips to be positioned in.[2] These slots may be conceptualized as broad regions of the PWB, as shown above, or as specifically located available board space on a more tightly constrained PWB layout.

No explicit constraint has been placed on the number of components which may cluster in one region (square patch in the figure above). If, conversely, a one-to-one mapping of components to slots is desired, a penalty function could be added to the evaluation function. For example, the locally variable board temperature problem provides a mild penalty to the grouping of components together, because each component in a region raises the temperature by 2 degrees, making it a less suitable place for additional components.

Since each chip can then be located in any of the regions, or slots, on the board, and the total number of regions on the board was set at 32, the size of the possible solution space is:

$$N = 32^5 \approx 33.6 \text{ million possible solutions.}$$

A simple way to encode the information required is to have a structure of 5 genes, each corresponding to one of five LCCC chips to be placed on the PWB. Each gene can take a value from 0 to 31, corresponding to the region that chip is located in. This is illustrated below.

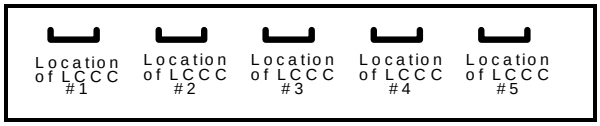

Figure 5.1-2 : Chromosome/String Format

[2]The actual structures manipulated by GENESIS are binary strings. The range of values that each gene may assume, therefore, must be a power of 2.

Several features of the GENESIS code were utilized to improve the performance of the GA. Elitist selection strategy was employed to insure the best string from each generation was not lost (removed from the next generation) due to mutation or replacement. The location values were represented in binary form using a Gray code rather than the usual base 2 binary form. (Gray codes have the property that adjacent integer values differ at exactly one bit position.) Finally, a scaling function was applied to prevent early emergence of an inferior solution.

The results obtained from these implementations are discussed below.

6. RESULTS

6.1 Initial Results and Trend Analysis of GA Parameters

As mentioned in the literature review, Genetic Algorithms suffer from a high degree of sensitivity to the values of the 'control parameters'. These parameters- population size, crossover rate, mutation rate, replacement rate- drastically affect how quickly the GA finds the optimal solution. (It has been shown that GAs will always find the optimal solution in a finite period of time, although this time may be transcomputational. [Rudolph, 1994])

Generic recommendations exist for 'good' control settings, based on observations of performance on a wide range of different functions. We conducted a series of experiments on a simplified problem to tune the parameters for the specific nature of our function space. (Note that the values obtained here are for the evaluation function described above, and may not be quite as optimal for a non-linear FE evaluation function.)

The simplified problem chosen was placement of 5 chips on a board with an arbitrary fixed temperature distribution which remains unaffected by chip placement. Thus, the ideal solution to this problem can be found by solving 5 independent placement problems. (The GA sees a coupled problem, however, because the only information it receives is the worst performance of all 5 chips, with no information identifying which chip is failing first. We force the GA to solve the coupled problem because we are trying to tweak the parameters for searching the coupled solution space.) No restrictions are made on placement, so the problem solution is easily visible by inspection- all of the chips must be placed in the coolest part of the board.

Factorial analysis of the control settings revealed a high degree of interaction. These interactions and possible explanations for them are produced in the table below.

Effect of	and:	and:	is	Why?
High Replacement (Generation Gap) (100% vs. 50%)			−	Replace 'good' structures faster than you can make them
High Mutation (5% vs. 1%)			−	Disrupts inheritance
Larger Population (500 vs. 100)			−	Slows convergence
More trials (10,000 vs.1,000)			+	More 'time' to look
More trials	Larger Population		+	When a big population finally converges, it's a very good solution.
High Mutation	Larger Population		−	Worsens a slow convergence
More trials	High Mutation	Larger Population	−	2 Strong divergent forces on population
More trials	High Crossover	Larger Population	−	For disruption, high crossover ≈ mutation

Table 6.1.1: Short description of GA control parameter effects

The optimal settings thus obtained are reproduced in Table 6.1.2 below, along with typical generic 'recommended' settings.

	PWB Settings	Generic Settings
Crossover Rate:	0.60	0.60
Generation Gap:	0.50	0.50
Mutation Rate:	0.01	0.001
Population Size:	500	50
Number of Trials:	25,000	1,000

Table 6.1.2: Experimentally tuned vs. 'Plain Vanilla' GA Parameter settings

The parameters differ from recommended generic values mainly as a tradeoff of efficiency versus maximum exploration of the state space. We use a higher mutation rate to lower the chances of getting stuck in a local optimum. We use a higher population to start with a higher sampling of the solution space.

With an optimally tuned GA, we are now ready to tackle the more challenging problem where chip placement affects local board temperature.

6.2 Coupled Component Placement/Board Temperature Results

The component placement problem addressed in the previous section, although reasonably difficult from a GA point of view, given the number of possible solutions and the information content of the evaluation function, is fairly simple to solve manually. A simple examination of the board reveals the coolest area, and the best solution is simply to place all the chips there. Unfortunately, the problem of placement of heat generating chips is not that simple in the real world, because very placement of the chips affects the temperature distribution we are trying to exploit to find the coolest part of the board. To better model this real world process, a slightly more sophisticated evaluation function reflects the localized increase in temperature due to a chip placement, as illustrated below.

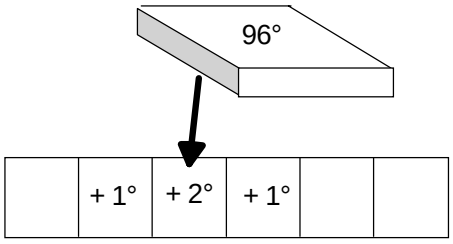

Figure 6.2-1: Localized PWB warming by LCCC chip

The temperature of the board is increased by 2 degrees Celsius in the board region in which the chip is placed, and by 1 degree in adjacent numbered regions. Heat is allowed to spread to adjacent numbered regions rather than geographically adjacent regions to simplify the coding involved and to prove the concept. The approach taken here can be generalized to a more elaborate geography-based heat distribution scheme of arbitrary complexity (accounting for the direction of forced convection, for example).

Note that this approach, adopted to minimize computational effort in this feasibility study, is linearly additive- there are no nonlinear effects due to side by side placement of chips. This linear additive property means other optimization algorithms could be used to solve this specific model. (For example, placing each chip sequentially in the coolest remaining area of the board.) We have used a GA, however, for a number of reasons. First, the GA approach has reliably found the optimal placement for several thermal distributions and different chip/board local heating patterns. Second, the component placement research undertaken by the authors cited in Section 3.2 above indicates additional performance metrics associated with interconnect lengths and PWB area utilized can be usefully incorporated into a Genetic Algorithm approach. Finally, a GA will permit a non-linear interaction model, where all the chips must be placed on the board before a final temperature distribution can be computed. GAs excel at solving this type of problem, where placements must be known before an evaluation can be computed.

Even this simplistic coupling of the temperature profile to the component placement, however, makes this problem difficult to solve by examination. This problem approximates a realistic optimization problem in printed wiring assemblies, and so is a useful application case to test the full power of GAs.

GENESIS was able to solve this problem extremely quickly, handling the 25,000 scheduled trials in approximately 4 seconds on the SparcServer 1000 described earlier. One of the best placement solutions for fatigue is illustrated below. First, the original board, shown both numerically and as a patch plot where elevation indicates temperature (cooler patches are darker):

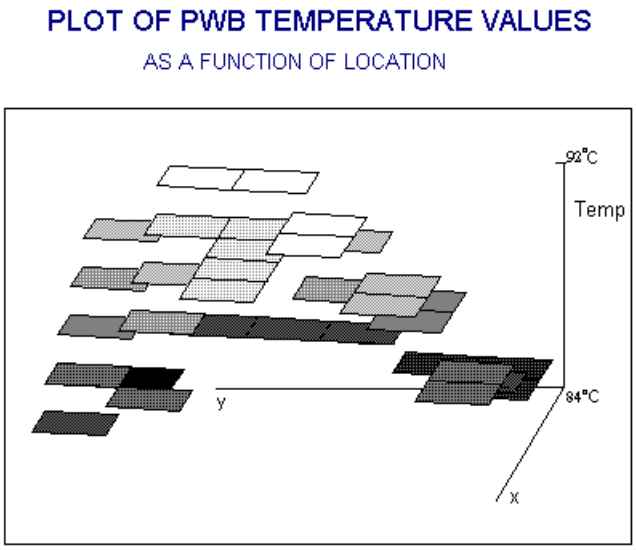

BOARD TEMPERATURES

$$\begin{bmatrix} 84 & 86 & 86 & 86 & 85 & 85 \\ 90 & 92 & 92 & 90 & 88 & 85 \\ 89 & 91 & 91 & 89 & 88 & 86 \\ 88 & 90 & 91 & 92 & 90 & 87 \\ 87 & 89 & 91 & 92 & 90 & 87 \\ 86 & 87 & & & & \end{bmatrix}$$

Figure 6.2-2: PWB temperatures before warming by LCCC chips

Second, the temperature of the same board, after the LCCC components have been positioned to yield the maximum fatigue life, is illustrated below (arrows and circles indicate sites):

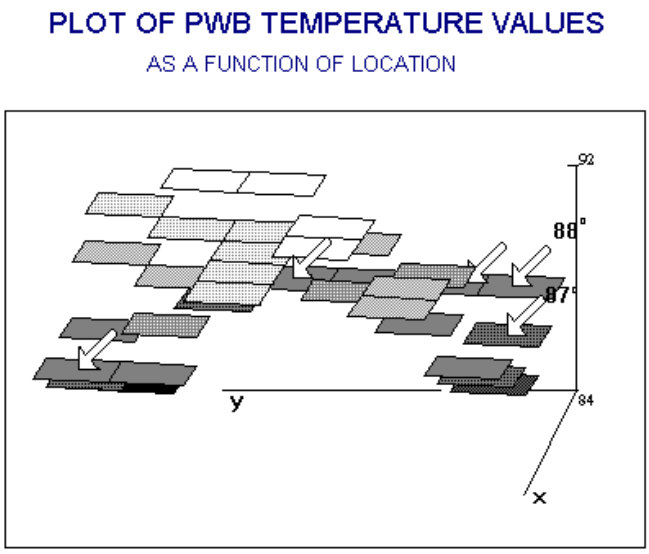

ADJUSTED BOARD TEMPERATURES

84	87	(88)	88	(88)	(88)
91	92	92	90	89	(87)
90	91	91	89	88	86
88	90	91	92	90	87
87	89	91	92	90	88
(88)	88				

Figure 6.2-3: PWB temperatures after warming by LCCC chips (arrows and circles indicate chip placement)

Note the overall increase in board temperature illustrated in Figure 6.2-3 (indicated by the increased number of white patches), a result of the localized warming effected by the chips. The actual value obtained for the fatigue life of the board was 847 cycles- bounded by an LCCC placed over a section of PWB with a local temperature of 88°C. This fatigue life was verified as an uppermost limit possible with the given temperature profile by re-running the GA with an iteration limit of 100,000 trials. The fatigue life obtained is consistent with the values obtained by Engelmaier, as he states, "Thus for an application requiring 1000 power on/off cycles per year, solder joint failures will occur in less than one year with [large LCCC components and] epoxy/glass substrates..." [Engelmaier, 1983].

The placement illustrated above is not the only possible arrangement. One of the benefits of a GA approach is that it generates multiple solutions (if multiple solutions are possible). The specifics of the problem described above allowed the GA to generate 8 alternative placements with equal fatigue lives. These alternate solutions can be evaluated with other PWB performance metrics in a lexicographic fashion, if desired.

7. DIRECTIONS FOR FURTHER IMPROVEMENTS AND CLOSURE

In this paper, we discussed the use of Genetic Algorithms for exploring the vast combinatoric solution space of alternate PWB layouts. We have provided an overview of genetic algorithm research and their application to electric component placement. We applied a GA code (Genesis) to solve a combinatorial thermo-fatigue placement problem using a well-established formula based approximation for fatigue life (the Engelmaier model) and simple heuristics for heat transfer effects. Our GA implementation was able to find superior PWB layouts. In our opinion, a GA driven PWB layout design tool holds promise for design explorations.

Open issues to be addressed are the computational demand of GAs versus the accuracy and detail needed at different stages in the PWB design process. The high computational demand of GAs on serial computer hardware will realistically restrict a GA-based PWB thermal design tool to using formula expressions and numeric approximations. Industrial applications will, of course, require more accurate thermal models and methods for calculating the changing thermal gradient across the PWB than those explained here. Also, other thermal management objectives, such as reducing the temperature of the hottest part of the board, or separating heat sensitive components from heat sources need to be investigated. At some point, a trade-off has to be made between searching a vast design space with lower accuracy versus a small search with higher accuracy.

In our opinion, however, the application of GAs to the preliminary design of PWBs holds great promise in improving the thermal performance of electronic packages. We believe that sufficiently sophisticated approximations will allow the discovery of beneficial layouts, possibly far removed from a starting configuration, which can then be thoroughly examined by extensive finite element analysis. We believe, at present, that a GA based design tool is best envisioned as a scout, exploring the length and breadth of the solution space for the most globally promising PWB layouts, where more exhaustive methods can be focused. Our future work will be focused on establishing the boundaries for such explorative design tools.

REFERENCES

Akay, H.U.; Tong, Y.; Paydar, N., "Thermal fatigue analysis of an SMT solder joint using nonlinear FEM approach, " International Electronics Packaging Conference, Proceedings of the Technical Conference, Austin, TX, USA. vol.2 1992 pp.980-998.

Capillo, Carmen. Surface Mount Technology Materials, Processes, and Equipment. New York: McGraw Hill, 1990.

Chan, Heming, Mazumder, P., Shahookar, K., "Macro-cell and module placement by genetic adaptive search with bitmap-represented chromosome," Integration, the VLSI Journal v 12 n 1 Nov 1991 pp. 49-77.

Clech, Jean-Paul M., Engelmaier, Werner, Kotlowitz, Robert W., Augis, Jacques A., "Reliability figures of merit for surface-soldered leadless chip carriers compared to leaded packages," IEEE Transactions on Components, Hybrids and Manufacturing Technology, vol. 12 no. 4 Dec 1989 pp. 449-458.

Cohoon, James P., Paris, William D., "GENETIC PLACEMENT," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems v CAD-6 n6 Nov 1987.

Davis, Lawrence, ed. Handbook of genetic algorithms. New York: Van Nostrand Reinhold, 1991.

De Jong, K.A., Spears, W.M., Schwefel, H.-P., Manner, R., "An analysis of the interacting roles of population size and crossover in genetic algorithms", Parallel Problem Solving from Nature. 1st Workshop, PPSN 1 Proceedings, Dortmund, West Germany, 1991, pp.38-47.

Engelmaier, Werner. IEEE Transactions on Components, Hybrids, and Manuf. Technol., vol. CHMT-6, no. 3, Sep 1983.

Engelmaier, W., 1989, "Thermal-Mechanical Effects," in Electronics Material Handbook. Volume 1. Packaging, Minges, M. L., ed., ASM International, Materials Park, OH, pp 740-753.

Fine, M.E., Jeannotte, D.A., Shieh, W.T., "Soft solder-general principles and fatigue lifetime estimation," Microelectronic Packaging Technology, Materials and Processes. Proceedings of the 2nd ASM International Electronic Materials and Processing Congress, Philadelphia, PA, 24-28 April 1989. pp.147-57.

Goldberg, David E. Genetic Algorithms in Search, Optimization, and Machine Learning. Reading: Addison-Wesley Publishing Company, Inc., 1989.

Goodman, T.; Snyder, D.; Weiss, J., "Predicting circuit board hot spots," Computer-Aided Engineering, vol.9, no.9 Sept 1990. pp.71-72.

Grefenstette, John J., "OPTIMIZATION Of CONTROL PARAMETERS FOR GENETIC ALGORITHMS", IEEE Transactions on Systems, Man and Cybernetics, v. SMC-16 n. 1 Jan-Feb 1986 pp. 122-128.

Grefenstette, John J. "A User's Guide to GENESIS Version 5.0." 1990. Unpublished documentation provided with source code.

Holland, John H. Adaptation in Natural and Artificial Systems. Ann Arbor: The University of Michigan Press, 1975.

Jin, Lin-Ming, Chan, Shu-Park, "Analogue placement by formulation of macro-components and genetic partitioning," International Journal of Electronics v 73 n 1 Jul 1992 pp. 157-173.

Knight, Leslie R., Wainwright, Roger L., "HYPERGEN - a distributed genetic algorithm on a hypercube," Scalable High Performance Computing Conference Proc (SHPCC 92), Williamsburg, VA, USA. Publ by IEEE, Computer Society, Los Alamitos, CA, USA (IEEE cat no 92TH0432-5). pp. 232-235.

Leu, Ming C., Wong, Hermean, Ji, Zhiming, "Genetic algorithm for solving printed circuit board assembly planning problems," Proceedings of the 1992 Japan - USA Symposium on Flexible Automation, San Francisco, CA, USA. Part 2 (of 2). Publ by ASME, New York, NY, USA. pp. 1579-1586.

Mutalik, Pooja P., Knight, Leslie R., Blanton, Joe L., Wainwirght, Roger L., "Solving combinatorial optimization problems using parallel simulated annealing and parallel genetic algorithms," Applied Computing: Technological Challenges of the 1990's - Proceedings of the 1992 ACM/SIGAPP Symposium on Applied Computing SAC '92, Kansas City, KS, USA. Publ by ACM, New York, NY, USA. pp. 1031-1038.

Ozmat, B., "A nonlinear thermal stress analysis of surface mount solder joints," 1990 Proceedings 40th Electronic Components and Technology Conference p.959-72 vol.2.

Peak, R. S. and Fulton, R. E., "Integrating Analysis and Design Information in Electronic Packaging Using Product Model-based Analytical Models," Computer Aided Design in Electronic Packaging, American Society of Mechanical Engineers Winter Annual Meeting, EEP Vol. 3, Anaheim CA, 1992, pp.41-56.

Rudolph, G., "Convergence analysis of canonical genetic algorithms," IEEE Transactions on Neural Networks, vol.5, no.1, Jan. 1994 pp.96-101.

Shahookar, Khushro, Mazumder, Pinaki, "Genetic approach to standard cell placement using meta-genetic parameter optimization," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems v 9 n 5 May 1990 pp. 500-511.

Wong, H., Leu, M.C., "Adaptive genetic algorithm for optimal printed circuit board assembly planning", CIRP Annals v 42 n 1 1993. pp. 17-20.