

AD-A207 230

FILE COPY

1

Temperature Measurement and Equilibrium Dynamics of Simulated Annealing Placements

Jonathan Rose
Computer Systems Laboratory,
The Center For Integrated Systems, Stanford University

Wolfgang Klebsch and Juergen Wolf
Siemens AG, Munich, Federal Republic of Germany

DTIC
ELECTE
APR 28 1989
S Q H D

N00014-87-K-0828
Abstract **1987**

One way to reduce the computational requirements of Simulated Annealing placement algorithms is to use a faster heuristic to replace the early phase of Simulated Annealing. Such systems need to know a starting temperature for the annealing phase that makes the best use of the structure provided by the heuristic, yet does an appropriate amount of improvement. This paper presents a method for *measuring* the temperature of an existing placement. It is based on a view of Simulated Annealing state that differs from previous work - the probability distribution of the *change* in cost function, as opposed to the *absolute* cost function. Using this view a new definition of equilibrium is given and the *equilibrium temperature* of a placement is defined. This also gives rise to an new view of the *equilibrium dynamics* of Simulated Annealing. A measure is developed that quantifies the nearness of a Simulated Annealing placement to equilibrium, and experimental evidence of its ability to detect equilibrium is given. Based on the measure a method is presented for determining the equilibrium temperature of a placement, and it is applied to placements of a real circuit produced both by a Simulated Annealing and a Min-Cut placement algorithm. For the latter an experimental relationship between the Min-Cut *cut area* and the measured temperature is demonstrated.

See Section 1. (ES)

1 Introduction

The success of the Simulated Annealing algorithm for automatic placement [Sech85] has been hindered by its excessive computational requirements. Recent work on standard cell placement algorithms [Rose86a, Grov87, Rose88a] has suggested alleviating this by using a two-stage approach: begin with a good, reasonably successful heuristic such as the Min-Cut algorithm [Breu77, Duni85] and then follow it with a Simulated Annealing-based approach for more fine optimization. Replacement of the early phase of Simulated Annealing with a faster but potentially worse algorithm allows a tradeoff between execution time and quality. A critical issue in this approach is to decide the starting temperature of the Simulated Annealing phase. If the temperature is too high, then some of the structure created by the first phase will be destroyed and unnecessary extra work will have to be done in the Simulated Annealing phase. If the temperature is too low then solution quality is lost, similar to the case of a quenching cooling schedule [Whit84].

This paper presents a technique for *measuring* the temperature of a placement for use in such two-stage systems. The problem is to determine the starting temperature for a Simulated Annealing process

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

089 4 26 077

so that the "best" use of the original structure is made, yet an "appropriate" amount of optimization is done to improve it. To give meaning to the concept of a placement's temperature, a framework is needed in which the notions of "best" and "appropriate" are defined.

Accordingly, we present a new view of Simulated Annealing state different from those articulated in [Laar87, Aart85, Rome84, Whit84]. The principal difference is that we look at probability distributions of the *change* in cost function of a Simulated Annealing state, rather than the absolute cost function. Using this view we give a definition of equilibrium from which follows the notion of the *equilibrium temperature* of a placement. The way in which the probability distribution changes as equilibrium is reached, known as the *equilibrium dynamics* [Laar87], is demonstrated with measurements on a real circuit.

We develop a measure that quantifies the nearness of a Simulated Annealing placement to equilibrium, called the Cost Force Ratio (CFR), and give experimental evidence of the CFR's ability to detect equilibrium. Based on the CFR measure, we present a method for measuring the equilibrium temperature of a placement, and show that it works both for placements produced by a Simulated Annealing and a Min-Cut placement algorithm. For the latter we show an experimental relationship between the Min-Cut *cut area* and the measured temperature.

The determination of starting temperature for Simulated Annealing in two-stage systems has not been seriously addressed before. Both [Rose86a, Rose88a] and [Gro87] introduce the question but avoid answering it by choosing a starting temperature based simply on previous experience. A shorter version of this paper is to appear in [Rose88b].

2 Definition of Equilibrium and Temperature

In previous discussions of cooling schedules and convergence [Laar87, Aart85, Rome84, Whit84], the Simulated Annealing state has been represented either as the probability distribution of the absolute cost $P(C)$, or the set of transition probabilities from every state i to every other state j , T_{ij} . We suggest a different view that gives more information about equilibrium dynamics: the probability distribution of the *change* in cost function from the current state. $P(\Delta C)$ is the probability that a given state under a Simulated Annealing process with a particular generation function [Rome84] will generate a move with a change in cost function of ΔC . $P(\Delta C)$ varies with temperature (T) and as moves are made.

We can use this view to give a different perspective on the equilibrium of a Simulated Annealing process. Since at equilibrium the absolute cost function no longer changes, this implies that the expected value of the *change* in cost function is zero:

$$E(\Delta C) = 0 \tag{1}$$

An expression for $E(\Delta C)$ can be formed assuming that $P(\Delta C)$ is known:

$$E(\Delta C) = \int_{-\infty}^{\infty} \Delta C P(\Delta C) P_{Accept}(\Delta C) d\Delta C \quad (2)$$

$P_{Accept}(\Delta C)$ is the probability that the acceptance function will accept a move with cost ΔC [Rome84]. It commonly has the value 1 for $\Delta C \leq 0$ and $e^{\frac{-\Delta C}{T}}$ for $\Delta C > 0$ [Sech85]. We note here that $P(\Delta C)$ in equation (2) must be the distribution measured on a *running* Simulated Annealing process at the equilibrium temperature. This distribution is difficult to measure, as will be discussed further in Section 3.1.

Using this $P_{Accept}(\Delta C)$ we can split equation (2) into two parts and, and at equilibrium from equation (1) we can equate it to zero:

$$\int_{-\infty}^0 \Delta C P(\Delta C) d\Delta C + \int_0^{\infty} \Delta C P(\Delta C) e^{\frac{-\Delta C}{T}} d\Delta C = 0 \quad (3)$$

Thus *equilibrium* can now be defined as the state where, at a given $T = T_{eq}$, the distribution $P(\Delta C)$ satisfies equation (3). Conversely, the *equilibrium temperature* of a placement with a distribution $P(\Delta C)$ is the temperature, T_{eq} , for which equation (3) is satisfied.

We note here that $P(\Delta C)$ in equation (2) must be the distribution measured on a running Simulated Annealing process at the equilibrium temperature. This distribution is difficult to measure, as will be discussed further in Section 3.1.

2.1 Equilibrium Dynamics

The way in which the probability distributions change throughout the process, or the *equilibrium dynamics*, can be explained by observing how $P(\Delta C)$ changes when moving from non-equilibrium to equilibrium. Suppose a system is in equilibrium at temperature T_1 , and its temperature is then lowered to T_2 . Figure 1 is a plot of $P(\Delta C)$ and P_{Accept} versus ΔC for a fictitious system in equilibrium at temperature T_1 . When the temperature is lowered to T_2 the only change is that the positive portion of the accept function becomes uniformly lower because $e^{\frac{-\Delta C}{T_2}} < e^{\frac{-\Delta C}{T_1}}$ for all $\Delta C > 0$.

For this system to regain equilibrium after the temperature change, $P(\Delta C)$ must change to again satisfy equation (3). This means that one or both of the following must happen:

1. The positive portion of $P(\Delta C)$ must either shift right (greater bad moves) or up (more bad moves), increasing the expected positive component of ΔC (E_+) or,
2. The negative portion of $P(\Delta C)$ must either shift right (smaller good moves) or down (fewer good moves), reducing the magnitude of expected negative component of ΔC (E_-).

For	<input checked="" type="checkbox"/>
	<input type="checkbox"/>
	<input type="checkbox"/>
on	
by	
ty Codes	
Avail and/or	
Special	



A-1




Figure 1 - Fictitious Probability Distribution and Acceptance Function at Temperature Change

Experimentally, both these effects are observed. Figure 2 is a plot of $P(\Delta C)$ versus ΔC for the 833 standard cell Primary1 circuit from the Preas-Roberts standard cell benchmark suite [Prea87]. It was produced by the SALTOR Simulated Annealing placement program [Rose86b,Rose88a], which is based on the ideas of the Timberwolf standard cell placement program [Sech85]. $P(\Delta C)$ is measured by generating 200,000 moves on a placement without actually accepting those moves (these are called *virtual* moves). In this way the placement is not changed, and a "point" measurement of the distribution in time is obtained. As discussed in Section 3.1, this *static* measurement is very close to the *dynamic* one, where the measurement is made on the Simulated Annealing process running in equilibrium.

Figure 2 gives $P(\Delta C)$ for three temperatures: very high ($T = 5000$), medium ($T = 300$) and low ($T = 9$). As the temperature decreases, the negative portion of $P(\Delta C)$ undergoes a dramatic shift to the right, and is much smaller than the positive portion of $P(\Delta C)$. This relates to the placement process in that all of the large good moves are used up, and only a few relatively small improvements are possible.

As temperature decreases, the positive portion of $P(\Delta C)$ in Figure 2 undergoes a right and upward shift. This occurs because as the placement gets better, there are more moves that will have a greater bad effect on the placement.

2.2 An Equilibrium-Nearness Measure

Using equation (3) we can invent a measure of the nearness of a given Simulated Annealing state to equilibrium. Define E_- to be the absolute value of the first term in the equation, that is

$$E_- = \left| \int_{-\infty}^0 \Delta C P(\Delta C) d\Delta C \right|$$




Figure 2 - $P(\Delta C)$ versus ΔC on Primary1 for Temperatures 5000, 300 and 9

Similarly let E_+ be the second term of equation (3):

$$E_+ = \int_0^{\infty} \Delta C P(\Delta C) e^{-\frac{\Delta C}{T_m}} d\Delta C$$

Where T_m is the temperature of the Simulated Annealing process. We can now define the *Cost Force Ratio*, (CFR) as:

$$CFR = \frac{E_-}{E_+ + E_-} \times 100 \quad (4)$$

The closer CFR is to 50% (the expected value of the good moves being equal to the expected values of the bad moves, $E_- = E_+$) the closer the system is to equilibrium.

Figure 3 is a plot of CFR versus generated move number for a Simulated Annealing process running on circuit Primary1, as it goes from non-equilibrium to equilibrium at temperature 400 changing to 300. CFR is determined by keeping a window of ΔC values multiplied by the P_{Accept} function and using this to calculate E_+ and E_- . In this figure the CFR comes down from an initial value of 55% and hovers around 50%. This shows that the CFR indicates when equilibrium has been achieved. It varies about the 50% point due to the stochastic nature of the algorithm and the approximation of measuring the CFR in a finite window.




Figure 3 - CFR versus Move Number as Process Achieves Equilibrium

3 Measuring Temperature

As defined in Section 2, the temperature of a placement is the temperature at which the Simulated Annealing process running on the placement is in equilibrium. In this section we present a method for measuring the temperature of an arbitrary placement.

The method is called the CFR Binary Search and has the following outline:

1. Measure $P(\Delta C)$ for the given circuit under the Simulated Annealing process. This is discussed in detail in Section 3.1.
2. Set the starting search temperature, T_m , arbitrarily.
3. Based on the current T_m , calculate:

$$P_{Accept}(\Delta C) = e^{\frac{-\Delta C}{T_m}} \quad \Delta C > 0$$

$$= 1 \quad \Delta C \leq 0$$

4. Calculate the effective probability distribution: $P_{eff}(\Delta C) = P(\Delta C) \times P_{Accept}(\Delta C)$.
 $P_{eff}(\Delta C)$ is the probability that a move with cost ΔC will be both generated and accepted.
5. Calculate the Cost Force Ratio, CFR, using $P_{eff}(\Delta C)$ and

$$E_- = \int_{-\infty}^0 \Delta C P_{eff}(\Delta C) d\Delta C$$

$$E_+ = \int_0^{\infty} \Delta C P_{eff}(\Delta C) d\Delta C$$

and equation (4).

6. If $CFR < 50$, reduce T_m according to a binary search and go to step 3;
If $CFR > 50$, increase T_m according to a binary search and go to step 3.
7. When $CFR = 50$, T_m is the equilibrium temperature, T_{eq} . Finish.

Each iteration of the CFR Binary Search requires only the recalculation of the positive portion of the acceptance function probability, $P_{Accept}(\Delta C)$, and subsequently E_+ and CFR since E_- does not change with T_m . Note also that $P(\Delta C)$ need only be generated once. This is important since it takes many moves (10^4 to 10^5) to get an accurate picture of the probability distribution.

3.1 Measurement of the Probability Distribution

A key and difficult step in the CFR Binary Search temperature measurement procedure is the measurement of the distribution $P(\Delta C)$. There are two potential methods:

1. **Static Measurement.** $P(\Delta C)$ is measured by generating *virtual* moves in the Simulated Annealing process on the placement, and recording the frequency with which each cost occurs. That is, moves are generated in the usual manner, but none are accepted, and so the placement does not change.
2. **Dynamic Measurement.** $P(\Delta C)$ is measured by generating *and* accepting moves on the placement. Here the placement does change as the measurement is made.

For the general case of any Simulated Annealing application a static measurement will not give the correct distribution. This is because a static measurement of $P(\Delta C)$ could be taken when the system was at a local (but not global) optimum. In this case there would be no good (negative) moves generated and since E_- would appear to be 0, the temperature would also appear to be 0, which is incorrect in the case of a local optimum. This is an example of an extreme case, but similar problems can occur when the state is at or near discontinuities in the energy landscape.

It is not possible, however, to measure dynamically the distribution while running a Simulated Annealing process at the placement's equilibrium temperature because that temperature is what we are seeking. If $P(\Delta C)$ is measured at the wrong temperature, then the placement's temperature will actually

change, and the $P(\Delta C)$ will reflect the temperature of the measuring process rather than the true temperature. This is not unlike the Heisenberg uncertainty principle - the act of measuring the temperature can cause the temperature to change.

An alternative is to measure $P(\Delta C)$ using the static method, and to determine how accurate this is as an approximation. The accuracy of this approach, with respect to dynamic measurement is entirely problem dependent - it depends on the energy landscape of the underlying Simulated Annealing formulation. We have experimented to determine the accuracy for the standard cell placement problem and have found that the static measurement of $P(\Delta C)$ is almost exactly the same as the dynamic measurement. Figure 4 shows a plot of a static distribution and a dynamic distribution measured on circuit Primary1 at temperature 300. Measurements and numerical comparisons on this and several other circuits at various temperatures have shown very small differences between the static and dynamic measurements. Thus we will use the static measurement of $P(\Delta C)$ in the temperature measurement algorithm.

Note that this can only be done because of the nature of the placement problem and the specific Simulated Annealing formulation - it is not a general result for all Simulated Annealing problems.




Figure 4 - Comparison of Static and Dynamic Measurement of $P(\Delta C)$

3.2 Temperature Measurements of Simulated Annealing Placements

The CFR Binary Search was used to measure the temperature of a set of Primary1 placements produced by the SALTOR Simulated Annealing placement program [Rose86,Rose88a]. Each placement was measured by using $N = 100,000$ virtual moves to experimentally determine $P(\Delta C)$. Table 1 gives the temperature at which each placement's Simulated Annealing process was terminated (while in equilibrium), and the measured temperature using the CFR Binary Search.

SA Produced Temperature	CFR Binary Search Measured Temp	Difference
500	496	-4
405	420	+15
294	285	-11
213	215	+2
153	164	+11
99	97	-2
57	60	+3
28	28	0
9	15	+6
2	4	+2

Table 1 - Temperature Measurement of Simulated Annealing Placements

The measured temperature is quite accurate at the higher temperature, usually less than 7% error. The lower temperature measurements are *proportionately* less accurate, but since their absolute values are small this is not surprising. The error is due to three effects:

1. The cooling schedule used to produce the placement is not perfect, and so the placement is probably not quite in equilibrium.
2. The slight difference, as discussed above, between the static and the (more correct) dynamic measurement of $P(\Delta C)$.
3. At lower temperatures, there are fewer negative moves, and so the accuracy of E_- decreases, decreasing the accuracy of CFR and hence the temperature measurement.

This last point can be seen experimentally: figure 5 is a plot of the percentage standard deviation of the measured temperature as a function of the number of virtual moves, N , for temperatures 28, 153 and 405. The standard deviation was calculated from five runs at each number of virtual moves. The variation is a




Figure 5 - Variation of Temperature with \log_{10} (Number Virtual Moves)

decreasing function of N , as would be expected. The increase in percentage variation at lower temperatures is illustrated as described above.

4 Temperature Measurement of Min-Cut Placements

The reason for measuring the temperature of a placement is to be able to switch from a non-annealing algorithm to an annealing-based one, and to begin at the correct temperature. In this section we first define a few relevant terms, then discuss the feasibility of measuring non-annealing placements, and finally measure a set of placements produced by the Min-Cut placement algorithm [Breu77,Dunl85].

4.1 Definition of Terms

Several terms first need to be defined for Min-Cut placements, as shown in Figure 6. A Min-Cut placement algorithm is characterized by, among other things, the order and spacing of the cut lines applied. In Figure 6, the rectangle represents the entire placement, over which is laid a set of vertical and horizontal cut lines. If the spacing of the vertical cut lines is V and of the horizontal cut lines is H , then the cut area, A , is given by $A = V \times H$.

4.2 Feasibility and Matching of Algorithms

One difficulty with measuring the temperature of non-annealing produced placements is that the definition of temperature presented in Section 2 depends on the associated Simulated Annealing process being in equilibrium. It is clear, however, that a placement produced by the non-annealing algorithm is not




Figure 6 - Definition of Cut-Area

in equilibrium. Thus we must make an approximation and assume that a min-cut placement can be thought of as being in equilibrium at some temperature. The effect of this approximation is measured in the next section where we compare the CFR Binary Search method with a more direct method.

Next we must take into account a mismatch between Min-Cut placement and the Simulated Annealing *move set* used in Timberwolf [Sech85] and SALTOR [Rose88a]. This move set allows cells to overlap and penalizes that overlap. The Min-Cut placement, however, has no overlap. Thus the first moves made on the Min-Cut placement during a Simulated Annealing process are more likely to be bad until a basic amount of overlap occurs, since almost every move will create some overlap where there was none before. This will shift the $P(\Delta C)$ distribution to the right and give erroneous results for a measured temperature. On the other hand, some Simulated Annealing algorithms, such as [Gro86], do not use overlap and would not have this problem. To avoid it here, we used a simplified circuit in which all cells were set to be of equal size and only exchange moves are made in the Simulated Annealing process. This prevents any overlap from occurring. Experimentally, we have seen that reasonable results are still obtained if overlap is allowed to occur, since the wire length portion of the cost function dominates the overlap.

4.3 Measurements

Using the CFR Binary Search method we measured the temperature of several Min-Cut placements with different cut areas. These placements were produced by the ALTOR standard-cell placement program [Rose85]. Table 2 gives the measured temperature for each placement and its cut area.

To check if the temperature measurements were correct, we measured the temperatures of the placements in a different way, called the *Delta Method*. The Delta Method finds the temperature of a placement by running an annealing process on the placement at a range of temperatures. It is run for 100 move generations per cell, for each temperature, and the percentage difference in absolute cost function is measured, called the delta. The temperature at which the absolute value of the delta is less than 2% is

the equilibrium temperature of the placement. This is a direct way of experimentally finding the temperature at which the change in cost function is near 0. The Delta Method requires much more computation than the CFR Binary Search method, and thus is of no practical use. Table 2 shows the temperatures determined by the delta method, and the difference between the the binary search method and the Delta method. The binary search temperature measurement of Min-Cut placements is not as accurate as those for Simulated-Annealing produced placements, yet it does track the temperature reasonably well.

Cut Area $\mu m^2 \times 10^4$	Temperature Measured		Difference
	Binary Search	Delta Method	
2021	398	374	+24
1011	234	200	+34
505.3	162	132	+30
252.6	124	96	+28
126.3	91	67	+24
63.22	73	50	+23
31.58	49	40	+9
25.74	40	32	+8
12.60	34	30	+4
7.697	29	27	+2
3.139	28	26	+2

Table 2 - Temperature Measurement of Min-Cut Placements

The CFR Binary Search method consistently overestimates the equilibrium temperature. due to the fact that a min-cut placement is not in equilibrium, as discussed in section 4.2. A more specific reason for this is that the Min-Cut placement leaves several particularly good moves possible, because of its lesser hill-climbing ability (we used [Fidu82] as the partitioning algorithm). A Simulated Annealing process would quickly correct these, but they result in an overestimation of E_c and hence a temperature that is too high.

4.4 Comments

Intuitively, one would expect the measured temperature of a Min-Cut placement to be an increasing function of the cut area, and this is observed in Table 2. This intuition comes from the notion that at higher temperatures, Simulated Annealing moves cells over large distances which determines a coarse placement. The first few cuts of Min-Cut placement, corresponding to a large cut area, also determines a coarse placement. At lower temperatures, Simulated Annealing makes moves that are much smaller in scope [Whit84] corresponding to the much smaller cut area of Min-Cut placement. The results of the

measurements shown in Table 2 bear out this intuition, as it is clear that the measured temperature is an increasing function of the cut area.

It is interesting to note the relationship between cut area and measured binary search temperature. We have found that the measured temperature is close to a linear function of the square root of the cut area (\sqrt{A}), as shown in Figure 7. The square root of the cut area is roughly equivalent to either the H or the V shown in Figure 6. This makes sense under the following line of reasoning: bad moves will move a distance from proportional to \sqrt{A} since Min-Cut only places cells to an "accuracy" of \sqrt{A} . Assume that the cost of those moves is proportional to the move distance, as an approximation. The temperature that is likely to accept bad moves of cost $k \times \sqrt{A}$ is also proportional to \sqrt{A} because moves are accepted with probability $e^{-\frac{k \times \sqrt{A}}{T}}$. Hence the temperature is an approximate linear function of the distance \sqrt{A} .

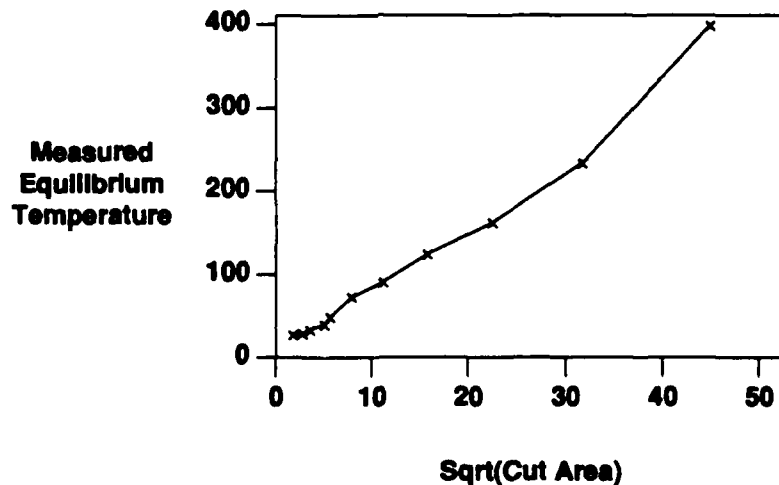


Figure 7 - Plot of Measured Temperature vs. Sqrt Cut Area

5 Conclusions

We have presented a method for determining the *temperature*, in the Simulated Annealing sense, of an arbitrary placement. It uses a new view of Simulated Annealing state that is based on the probability distribution of the change in cost function. This view provides a new definition of equilibrium, a measure of the nearness of a Simulated Annealing state to equilibrium, and an interesting perspective on equilibrium dynamics.

The temperature of several Simulated Annealing placements have been measured with good accuracy. The temperature of a set of Min-Cut placements has been measured with reasonable accuracy, and we have demonstrated an experimental relationship between cut area and temperature. These measurements are useful for determining the starting temperature when switching from a non-annealing

based placement strategy to an annealing-based one.

6 References

Aart85

E.H.L. Aarts, P.J.M. van Laarhoven, "A New Polynomial-Time Cooling Schedule," Proc. ICCAD 85, November 1985, pp. 206-208.

Breu77

M.A. Breuer, "Min-Cut Placement," Journal of Design Automation and Fault-Tolerant Computing, Oct 1977, pp 343-362.

Dunl85

A.E. Dunlop, B.W. Kernighan, "A Procedure for Placement of Standard-Cell VLSI Circuits," IEEE Transactions on CAD, Vol. CAD-4, No. 1, January 1985, pp 92-98.

Fidu82

C.M. Fiduccia, R.M. Matheyses, "A Linear Time Heuristic for Improving Network Partitions," Proc. 19th Design Automation Conference, June 1982, pp 175-181.

Grov86

L.K. Grover, "A New Simulated Annealing Algorithm for Standard Cell Placement," Proc. ICCAD 86, November 1986, pp. 378-380.

Grov87

L.K. Grover, "Standard Cell Placement Using Simulated Sintering," Proc. 24th DAC, June 1987, pp. 56 - 59.

Laar87

P.J.M. van Laarhoven and E.H.L. Aarts, *Simulated Annealing: Theory and Applications*, D Reidel Publishing Co., Dordrecht, Holland, 1987.

Prea87

B.T. Preas, "Benchmarks for Cell-Based Layout Systems," Proc. 24rd Design Automation Conference, June 1987, pp. 319-320.

Rome84

F. Romeo, A. Sangiovanni-Vincentelli, "Probabilistic Hill Climbing Algorithms: Properties and Applications," Memorandum No. UCB/ERL M84/34, March 1984, Electronics Research Laboratory, University of California, Berkeley.

Rose85

J.S. Rose, W.M. Snelgrove, Z.G. Vranesic, "ALTOR: An Automatic Standard Cell Layout Program," Proc. Canadian Conference on VLSI, November 1985, pp. 168-173.

Rose86a

J.S. Rose, "Fast, High Quality VLSI Placement on an MIMD Multiprocessor," Ph.D. Thesis, Department of Electrical Engineering, University of Toronto 1986; also Computer Systems Research Institute Technical Report # 189.

Rose86b

J.S. Rose, D.R. Blythe, W.M. Snelgrove, Z.G. Vranesic, "Fast, High Quality VLSI Placement on an MIMD Multiprocessor," Proc. ICCAD 86, November 1986, pp. 42-45.

Rose88a

J.S. Rose, W.M. Snelgrove, Z.G. Vranesic, "Parallel Standard Cell Placement Algorithms with Quality Equivalent to Simulated Annealing," IEEE Transactions on CAD, Vol. 7, No.3, March 1988, pp. 387-396.

Rose88b

J.S. Rose, W. Klebsch, J. Wolf, "Temperature Measurement of Simulated Annealing Placements," to appear in Proc. ICCAD 1988.

Sech85

C. Sechen, A. Sangiovanni-Vincentelli, "The Timberwolf Placement and Routing Package," IEEE JSSC, Vol. SC-20, No. 2, April 1985, pp 510-522.

Whit84

S.R. White, "Concepts of Scale in Simulated Annealing," Proc. Int. Conf. on Computer Design, October 1984, pp. 646-651.