

Temperature Measurement and Equilibrium Dynamics of Simulated Annealing Placements

JONATHAN ROSE, MEMBER, IEEE, WOLFGANG KLEBSCH, AND JÜRGEN WOLF

Abstract—One way to alleviate the heavy computation required by simulated annealing placement algorithms is to replace a significant fraction of the higher or middle temperatures with a faster heuristic, and then follow it with simulated annealing. A crucial issue in this approach is the determination of the starting temperature for the simulated annealing phase—a temperature should be chosen that causes an appropriate amount of optimization to be done, but makes good use of the structure provided by the heuristic. This paper presents a method for measuring the temperature of an existing placement. The approach is based on the measurement of the probability distribution of the change in cost function, $P(\Delta C)$, and makes the assumption that the placement is in simulated annealing equilibrium at some temperature. The temperature of placements produced both by a simulated annealing and a min-cut placement algorithm are measured, and good agreement with known temperatures is obtained. The $P(\Delta C)$ distribution is also used to give an interesting view of the equilibrium dynamics of simulated annealing.

I. INTRODUCTION

THE SUCCESS of the simulated annealing algorithm for automatic placement [18] has been hindered by its excessive computational requirements. Recent work on standard cell placement algorithms [6], [14], [16] has suggested alleviating this by using a two-stage approach: begin with a good, reasonably successful heuristic such as the min-cut algorithm [2], [3] and then follow it with a simulated annealing-based approach for more fine optimization. This can be used to tradeoff quality for execution time: as more of the simulated annealing work is replaced (and this can be just the high temperature, or include both middle and lower temperatures), the quality may decrease but so will the execution time.

A crucial issue in this approach is the starting temperature of the simulated annealing phase. If the temperature is too high, then some of the structure created by the heuristic will be destroyed and unnecessary extra work will 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 [19]. This paper presents a technique for measuring the temperature of a

placement for use in such two-stage systems. We make the assumption that the placement to be measured is in equilibrium for the given simulated annealing process at some temperature.

The temperature determination procedure is based on the measurement of the probability distribution of the change in cost function, $P(\Delta C)$, of a running simulated annealing process. This is a difficult distribution to measure, and we show how it can be approximately measured. Using $P(\Delta C)$, a method is presented for measuring the equilibrium temperature of a placement, and is shown to work both for placements produced by a simulated annealing and a min-cut placement algorithm.

The determination of starting temperature for simulated annealing in two-stage systems has not been seriously addressed before. Both [14], [16], and [6] introduce the question but avoid answering it by using a constant starting temperature based on experience. That approach gives no theoretical foundation for determining starting temperatures and cannot be used if the simulated annealing formulation is changed, or for other problems. Previous work on cooling schedules [1], [7], [10] gives no guidance in choosing a starting temperature other than the very highest. An earlier version of this work appeared in [17].

Section II of this paper defines the terms of simulated annealing, gives a precise definition of the equilibrium temperature of a placement, and the probability distribution $P(\Delta C)$. Section III shows how $P(\Delta C)$ changes at different temperatures. Section IV presents the algorithm for measuring the temperature of a placement and shows that it works on placements produced by the simulated annealing algorithm itself. Section V presents temperature measurements of placements produced by the min-cut placement algorithm.

II. SIMULATED ANNEALING TERMS, DEFINITION OF TEMPERATURE, AND $P(\Delta C)$

To define the temperature of a placement, we first review the terms of simulated annealing [8], [12], and give a few necessary assumptions required on the placement. A concrete definition of the probability distribution $P(\Delta C)$ is given, along with a discussion of how to measure it.

2.1. Simulated Annealing

The purpose of simulated annealing is to choose a good solution to an optimization problem according to some cost function on the state space of possible solutions. A

Manuscript received July 5, 1988; revised December 7, 1988 and July 3, 1989. This work was supported by DARPA under Contract N00014-87-K-0828, a Stanford Center for Integrated Systems Seed Research Grant, and a NSERC Post-Doctoral Fellowship. This paper was recommended by Associate Editor R. H. J. M. Otten.

J. Rose was with the Computer Systems Laboratory, Stanford University, Stanford, CA. He is now with the Department of Electrical Engineering, University of Toronto, Toronto, Ont., Canada M5S 1A4.

W. Klebsch and J. Wolf were with the Computer Systems Laboratory, Stanford University, Stanford, CA. They are now with Siemens AG, Munich, Germany.

IEEE Log Number 8932830.

simulated annealing optimization *process* is a procedure that iteratively perturbs a state of the optimization problem, and the perturbations are chosen using a *generation function*. The decision of whether or not to actually make the perturbation is determined by the *acceptance function*, and this is a function of the change in cost function (ΔC) and a control parameter referred to as the temperature, T . A simulated annealing *formulation* for a given problem consists of a generation function, a cost function, and an acceptance function.

A simulated annealing process generates perturbations, or *moves*, at a given temperature until *equilibrium* is achieved. If the process is modeled as a Markov chain [4], equilibrium is said to occur when the probability of being in a particular state becomes constant even as moves are being generated and accepted. That is, a *stationary probability distribution* is achieved.

2.2. Temperature

To define the notion of the temperature of a placement, we need to assume the following: that the placement under consideration is *in equilibrium* under the simulated annealing process at some temperature. This can be strictly true only for placements produced by the same simulated annealing process that are terminated in equilibrium. It is an approximation for any other placement. The effect of this approximation is shown, for one example, in Section V-5.3.

With this assumption, the temperature of a placement is defined as follows: for a given simulated annealing formulation, the temperature of a placement is the temperature for which the simulated annealing process running on that placement (i.e., beginning with that placement) is in equilibrium.

To give this definition concrete mathematical form, we use the observation that at equilibrium, the cost function no longer changes, implying that the expected value of the *change* in cost function is zero:

$$E(\Delta C) = 0. \quad (1)$$

An expression for $E(\Delta C)$ can be formed as follows:

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

where $P(\Delta C)$ is the probability that a simulated annealing process running on the placement at equilibrium will generate a move with cost ΔC . This distribution is defined mathematically below, in Section II-2.3, which includes a discussion on how $P(\Delta C)$ is measured. $P_{\text{Accept}}(\Delta C)$ is the probability that the acceptance function will accept a move with cost ΔC . It commonly has the value 1 for $\Delta C \leq 0$ and $e^{(-\Delta C/T)}$ for $\Delta C > 0$ [9]. This acceptance function will be used in the discussion below, but all of the following can be applied equally well to any acceptance function.

Using this $P_{\text{Accept}}(\Delta C)$ we can split (2) into two parts

and, and at equilibrium from (1) we can equate it to zero:

$$\int_{-\infty}^0 \Delta C P(\Delta C) d\Delta C + \int_0^{\infty} \Delta C P(\Delta C) e^{-\Delta C/T_{\text{eq}}} d\Delta C = 0. \quad (3)$$

Thus the *temperature* of a placement with a given distribution $P(\Delta C)$ is the temperature, T_{eq} , for which (3) is satisfied.

2.3. Definition and Measurement of $P(\Delta C)$

$P(\Delta C)$ is the probability that the generation function of the simulated annealing process running on the given placement generates a move with cost ΔC . It can be defined using the Markov chain terminology used in [9], [12]. Assume that there are N possible states (placements). For a simulated annealing process, the transition probability matrix \mathbf{P} is an $N \times N$ matrix whose elements P_{ij} are the probability of the changing from state i to state j in one step. \mathbf{P} is a function of the temperature, T . A similar $N \times N$ matrix, \mathbf{G} can be defined, whose elements G_{ij} are the probability that a process in state i will generate a move (regardless of whether the move is accepted) that will take it to state j . \mathbf{G} usually depends only on the generation function, and not temperature, unless the generation function is temperature dependent due to the use of a range-windowing device [18]. The effect of range windowing is discussed later in this section.

At equilibrium, the probability of being in any state is given by the stationary probability distribution vector, π whose elements π_i are the probability that the process is in state (placement) i , where $1 \leq i \leq N$. If we define the change in cost from state i to j as

$$\Delta C_{ij} \equiv \text{cost}(\text{state } i) - \text{cost}(\text{state } j)$$

then the distribution $P(\Delta C)$ is given by

$$P(\Delta C) = \sum_{1 \leq i \leq N} \sum_{i \leq j \leq N} \pi_i \times \begin{cases} G_{ij}, & \Delta C_{ij} = \Delta C \\ 0, & \Delta C_{ij} \neq \Delta C \end{cases} \quad (4)$$

We note that if the probability distribution π is stationary, then $P(\Delta C)$ will not vary as moves are made, since \mathbf{G} is constant.

The probability distribution of the states, π , is a function of the equilibrium temperature of the placement. The temperature, however, is the quantity we are seeking to determine and is not known. If $P(\Delta C)$ were measured by running a simulated annealing process and gathering statistics on ΔC at the *wrong* temperature, this would change π , and hence, the measured $P(\Delta C)$ would be incorrect. This is not unlike the Heisenberg uncertainty principle—the act of measuring a quantity can cause that quantity to change. Measuring the $P(\Delta C)$ distribution during a running simulated annealing process is called a *dynamic measurement*. It is impossible to make a correct dynamic measurement without knowing the placement's tempera-

ture. However, it is possible to get arbitrarily close to the right temperature because a measurement process (i.e., gathering statistics on ΔC to measure $P(\Delta C)$) can be restarted any number of times from the initial configuration. This could be used to determine the temperature, in a computationally expensive procedure, by iteratively searching for the temperature at which the cost function does not change at all. This procedure is described in Section V-5.3.

$P(\Delta C)$ can be approximated using a *static* method of measurement, which does not need to know the temperature. Here *virtual* moves are generated on the placement, and statistics are collected on the ΔC , but no moves are actually accepted. $P_{\text{static}}^i(\Delta C)$ then, is the probability of generating a move of cost ΔC from *one* particular state, i , and is defined as follows:

$$P_{\text{static}}^i(\Delta C) = \sum_{1 \leq j \leq N} \begin{cases} G_{ij}, & \Delta C_{ij} = \Delta C \\ 0, & \Delta C_{ij} \neq \Delta C \end{cases} \quad (5)$$

$P_{\text{static}}^i(\Delta C)$ is not a function of π , and hence, is not a function of temperature. Its accuracy as an approximation of $P(\Delta C)$ depends on the simulated annealing formulation, including the underlying problem to be solved. By comparing (4) and (5), it is clear that accuracy depends on how the distribution of $P(\Delta C)$ measured just from one state differs from the distribution measured from all states. For example, suppose that state i is a local minimum. Then $P_{\text{static}}^i(\Delta C)$ would be zero for all $\Delta C \leq 0$, whereas $P(\Delta C)$ would likely be nonzero.

Because the accuracy of the static measurement method is so problem dependent, we can only give quantitative comparisons for a specific problem and formulation. For the placement problem, using the simulated annealing standard cell algorithm similar to the one described in [18] and implemented in [16], the static approximation works very well. To demonstrate this, the simulated annealing process was run to produce placements at various temperatures on several circuits, and then $P(\Delta C)$ was measured using the static and dynamic methods. (The dynamic method can be used when the equilibrium temperature is known.) They were found to be almost exactly the same. For example, Fig. 1 is a plot of $P_{\text{static}}(\Delta C)$ and $P_{\text{dynamic}}(\Delta C)$ for the Primary1 standard cell benchmark circuit [11] produced at temperature 300. It shows that the distributions are very close, and indeed, they are even closer than they appear in the figure, as the dots are more offset from the line than is in fact the case.

This is not enough to guarantee, however, that $P_{\text{static}}(\Delta C)$ will always give a good approximation of $P(\Delta C)$, even for this case. It is likely that a greedy heuristic that uses the same cost function as the simulated annealing process would produce a placement in a local optimum. Note that some heuristics, such as min-cut [2], are unlikely to do so because they only place cells to a specified area and will not have a local optimum placement within that area. It is important then to ensure that the initial heuristic either uses a different cost function from the simulated annealing process, or that it prevents

a local optimum by not using a completely greedy approach. Note that this is how the temperature measurement is intended to be used—as a follow on from a non-optimum (local or global).

Lastly, note that $P(\Delta C)$ can be a function of temperature due to such techniques as *range-windowing* [18]. A range-window technique seeks to eliminate highly unlikely moves by preventing the generation function from yielding them. As long as the generation temperature is set to be very high, this will not affect the measured distribution, since at low temperatures the effect of the eliminated moves is negligible by definition of the range window. It does mean that more moves have to be generated because the efficiency of eliminating unlikely moves is lost.

III. 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 nonequilibrium to equilibrium. Define E_- , the magnitude of the expected value of the good (negative) moves, to be the absolute value of the first term of (3):

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

Similarly let E_+ , the expected value of the bad (positive) moves, be the second term of (3):

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

Here T_m is the temperature of the simulated annealing process.

Suppose that now a system is in equilibrium at temperature T_1 , and its temperature is then lowered to T_2 . Fig. 2 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 $e^{(-\Delta C/T_2)} < e^{(-\Delta C/T_1)}$ for all $\Delta C > 0$.

For this system to regain equilibrium after the temperature change, $P(\Delta C)$ must change to again satisfy (3)—i.e., make the magnitude of the expected value of the good moves equal to expected value of the bad moves. 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 E_+ ;
- 2) the negative portion of $P(\Delta C)$ must either shift right (smaller good moves) or down (fewer good moves), reducing E_- .

Experimentally, both these effects are observed. Fig. 3 is a plot of $P(\Delta C)$ versus ΔC for the 833 standard cell Primary1 circuit [11]. It was produced by the SALTOR simulated annealing placement program [15], [16], which is based on the ideas of the TimberWolf standard cell




Fig. 1. Comparison of static and dynamic measurement of $P(\Delta C)$.




Fig. 2. Fictitious probability distribution and acceptance function at temperature change.




Fig. 3. $P(\Delta C)$ versus ΔC on primary1 for temperatures 5000, 300, and 9.

placement program [18]. $P(\Delta C)$ is measured statically by generating 200 000 virtual moves on a placement.

Fig. 3 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 Fig. 3 undergoes a right and upward shift. This occurs because as the placement gets better, there are more moves that will have a worse effect on the placement.

IV. MEASURING TEMPERATURE

As defined in Section II, 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 Equilibrium Binary Search and has the following outline:

- 1) measure $P(\Delta C)$ for the given circuit under the simulated annealing process, using the static method as described in Section II-2.3;
- 2) set the starting search temperature, T_s , arbitrarily;
- 3) based on the current T_s , calculate:

$$P_{\text{Accept}}(\Delta C) = e^{-\Delta C/T_s}, \quad \Delta C > 0$$

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

- 4) calculate E_- and E_+ as defined above, in Section III;
- 5) if $E_- < E_+$, reduce T_s according to a binary search and go to step 3; if $E_- > E_+$, increase T_s according to a binary search and go to step 3; if $E_- = E_+$, T_s is the equilibrium temperature, T_{eq} . Finish.

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

4.1. Temperature Measurements of Simulated Annealing Placements

The Equilibrium Binary Search was used to measure the temperature of a set of Primary1 placements produced by

the SALTOR simulated annealing placement program [14], [16]. Each placement was measured by using $N = 100\,000$ virtual moves to experimentally determine $P(\Delta C)$. Table I gives the temperature at which each placement's simulated annealing process was terminated (while in equilibrium), and the measured temperature using the Equilibrium Binary Search.

The measured temperature is quite accurate at the higher temperature, usually less than a 7-percent 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 the following four effects.

1) Statistical variation caused by measuring $P(\Delta C)$ using a finite number of moves (see below).

2) The cooling schedule used to produce the placement is not perfect, and so the placement is probably not quite in equilibrium.

3) At lower temperatures, there are fewer negative moves, and so the accuracy of E_- decreases, decreasing the accuracy of the temperature measurement.

4) The slight difference, as discussed above, between the static and the (more correct) dynamic measurement of $P(\Delta C)$.

Point 1) can be seen experimentally: Fig. 4 is a plot of the percentage standard deviation of the measured temperature as a function of the number of virtual moves, N_v , for temperatures 28, 153, and 405. The standard deviation was calculated from five runs at each number of virtual moves. The variation is a decreasing function of N_v , as would be expected. The higher percentage variation at lower temperatures and fewer virtual moves is illustrated in the plot.

V. TEMPERATURE MEASUREMENT OF MIN-CUT PLACEMENTS

The reason for measuring the temperature of a placement is to be able to switch from a nonannealing algorithm to an annealing-based one, and to begin at an appropriate temperature. In this section we first define a few relevant terms, then discuss the feasibility of measuring nonannealing placements, and finally measure a set of placements produced by the min-cut placement algorithm [2], [3].

5.1. Definition of Terms

Fig. 5 depicts the cut lines typically used in a min-cut placement algorithm. Min-cut placement is characterized by, among other things, the order and spacing of the cut lines applied. In Fig. 5, 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 $A = V \times H$.

5.2. Feasibility and Matching of Algorithms

We must take into account a mismatch between min-cut placement and the simulated annealing *move set* used in TimberWolf [18] and SALTOR [16]. This move set allows cells to overlap and penalizes that overlap. The min-




Fig. 4. Variation of temperature with \log_{10} (number virtual moves).




Fig. 5. Definition of cut-area.

TABLE I
TEMPERATURE MEASUREMENT OF SIMULATED ANNEALING PLACEMENTS

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

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 formulations, such as [5], 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.

5.3. Measurements

Using the Equilibrium Binary Search method we measured the temperature of several min-cut placements with a range of cut areas. These placements were produced by the ALTOR standard-cell placement program [13]. Table II gives the measured temperature for each placement and its cut area.

TABLE II
TEMPERATURE MEASUREMENT OF MIN-CUT PLACEMENTS

Cut Area $\mu\text{m}^2 \times 10^4$	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
25.24	40	32	+8
12.60	34	30	+4
7.697	29	27	+2
3.139	28	26	+2

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 a dynamic 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 percent is taken as 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 Equilibrium Binary Search method. Table II shows the temperatures determined by the delta method, and the difference between 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.

The Equilibrium Binary Search method consistently overestimates the equilibrium temperature. The principal reason is that a min-cut placement is not in equilibrium, which is one of the assumptions on the input placement, as discussed in Section II-2.2.

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 II. This intuition comes from the observation that at higher temperatures, simulated annealing moves cells over large distances and determines a coarse placement, while at lower temperatures a more fine placement is determined by small moves [19]. The first few cuts of min-cut placement correspond to a large cut area, and hence, a coarse placement, while smaller cut areas result in a fine placement.

VI. CONCLUSIONS

We have presented a method for determining the *temperature*, in the simulated annealing sense, of a placement. It is based on the measurement of the probability distribution of the change in cost function, which is difficult to measure exactly. We describe a method of approximately measuring this distribution and show that it works well for a particular instance of a placement prob-

lem and algorithm implementation. We outline conditions under which the approximation will work.

The temperature of several simulated annealing placements have been measured with good accuracy, and the temperature of a set of min-cut placements has been measured with reasonable accuracy. These measurements are useful for determining the starting temperature when switching from a nonannealing based placement strategy to an annealing-based one.

ACKNOWLEDGMENT

The authors are grateful to the reviewers for insightful comments and suggestions for improving this paper, and to John Hennessy for the support and encouragement of this work.

REFERENCES

- [1] E. H. L. Aarts and P. J. M. van Laarhoven, "A new polynomial-time cooling schedule," in *Proc. ICCAD'85*, pp. 206-208, Nov. 1985.
- [2] M. A. Breuer, "Min-cut placement," *J. Design Automation Fault-Tolerant Computing*, pp. 343-362, Oct. 1977.
- [3] A. E. Dunlop and B. W. Kernighan, "A procedure for placement of standard-cell VLSI circuits," *IEEE Trans. Computer-Aided Design*, vol. CAD-4, pp. 92-98, Jan. 1985.
- [4] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3rd Ed. New York: Wiley, June 1970.
- [5] L. K. Grover, "A new simulated annealing algorithm for standard cell placement," in *Proc. ICCAD'86*, pp. 378-380, Nov. 1986.
- [6] L. K. Grover, "Standard cell placement using simulated sintering," in *Proc. 24th DAC*, pp. 56-59, June 1987.
- [7] M. D. Huang, F. Romeo, and A. Sangiovanni-Vincentelli, "An efficient general cooling schedule for simulated annealing," in *Proc. ICCAD'86*, pp. 381-384, Nov. 1986.
- [8] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671-680, May 13, 1983.
- [9] P. J. M. van Laarhoven and E. H. L. Aarts, *Simulated Annealing: Theory and Applications*. Dordrecht, The Netherlands: Reidel, 1987.
- [10] J. Lam, and J-M. Delosme, "Performance of a new annealing schedule," in *Proc. 25th DAC*, pp. 306-311, June 1988.
- [11] B. T. Preas, "Benchmarks for cell-based layout systems," in *Proc. 24th DAC*, pp. 319-320, 1987.
- [12] F. Romeo and A. Sangiovanni-Vincentelli, "Probabilistic hill climbing algorithms: Properties and applications," Memo. UCB/ERL M84/34, Elect. Res. Lab., Univ. Calif., Berkeley, Mar. 1984.
- [13] J. S. Rose, W. M. Snelgrove, and Z. G. Vranesic, "ALTOR: An automatic standard cell layout program," in *Proc. Canadian Conf. on VLSI*, pp. 168-173, Nov. 1985.
- [14] J. S. Rose, "Fast, high quality VLSI placement on an MIMD multiprocessor," Ph.D. dissertation, Dep. Elect. Eng., Univ. Toronto, 1986; also Computer Systems Res. Institute Tech. Rep. 189, 1986.
- [15] J. S. Rose, D. R. Blythe, W. M. Snelgrove, and Z. G. Vranesic, "

Fast, high quality VLSI placement on an MIMD multiprocessor," in *Proc. ICCAD'86*, pp. 42-45, Nov. 1986.

- [16] J. S. Rose, W. M. Snelgrove, and Z. G. Vranesic, "Parallel standard cell placement algorithms with quality equivalent to simulated annealing," *IEEE Trans. Computer-Aided Design*, vol. 7, pp. 387-396, Mar. 1988.
- [17] J. S. Rose, W. Klebsch, and J. Wolf, "Temperature measurement of simulated annealing placements," in *Proc. Int. Conf. Computer-Aided Design*, pp. 514-517, Nov. 1988.
- [18] C. Sechen and A. Sangiovanni-Vincentelli, "The Timberwolf placement and routing package," *IEEE J. Solid-State Circuits*, vol. SC-20, pp. 510-522, Apr. 1985.
- [19] S. R. White, "Concepts of scale in simulated annealing," in *Proc. Int. Conf. on Computer Design*, pp. 646-651, Oct. 1984.

*



Jonathan Rose (S'79-M'86) received the B.A.Sc. degree in engineering science in 1980, and the M.A.Sc. and Ph.D. degrees in electrical engineering in 1982 and 1986, respectively, from the University of Toronto.

During 1983, he was with Bell-Northern Research Ltd., Ottawa, in the Integrated Circuits CAD/CAM group. From 1986 to 1989, he was a Research Associate in the Computer Systems Laboratory at Stanford University. In 1989, he joined the faculty of the University of Toronto, where he

is currently an Assistant Professor of Electrical Engineering. His research interests include CAD and architecture for programmable gate arrays, automatic layout, and parallel CAD algorithms.



Wolfgang Klebsch received the M.S. and the Ph.D. degrees from the University of Cologne.

Since 1986, he has been with Siemens AG, Munich, Germany, where he works in the Department of Integrated Components, designing CMOS IC's in the telecommunications field. In 1987, he was a Guest Scientist at Stanford University, Stanford, CA.

*



Jürgen Wolf received the M.S. degree in computer science from the University of Bonn in 1986.

As an employee of Siemens AG, Erlangen, Germany, he was sent as a visiting scholar to Stanford University in 1987. He has been working with Siemens AG, Munich, Germany, since 1988.