On the Density of Solutions in Equilibrium Points for the Queens Problem

Paul Morris
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040
morris@intellicorp.com

Abstract

There has been recent interest in applying hill-climbing or iterative improvement methods to constraint satisfaction problems. An important issue for such methods is the likelihood of encountering a non-solution equilibrium (locally optimal) point. We present analytic techniques for determining the relative densities of solutions and equilibrium points with respect to these algorithms. The analysis explains empirically observed data for the n-queens problem, and provides insight into the potential effectiveness of these methods for other problems.

Introduction

In several recent papers [Minton et al. 1990, Zweben 1990, Morris 1990, Sosic and Gu 1991], iterative improvement methods for solving constraint satisfaction and optimization problems have been studied. These methods work by making local changes that reduce a cost function. This process continues until a configuration or state is reached such that no local change can reduce the cost further. We will call such a configuration an ***equilibrium point***. When an equilibrium point is reached, it is checked to see if it is an acceptable solution to the problem. If not, the algorithm may be restarted or some other action taken to proceed to a new equilibrium point. The papers above provide empirical evidence that such methods may lead to rapid solutions for important classes of problems.

One way of viewing these methods is that they perform a search of equilibrium points looking for solutions. Clearly the effectiveness of such a search is dependent on the density of solutions in equilibrium points. The methods will work particularly well if this density approaches 1. This motivates us to look for a way of analytically determining the density. Such an analysis is useful for predicting when iterative improvement methods are likely to be of value, and complements a quite different analytic approach presented in Minton et al. [1990] which estimates the probability that a single hill-climbing step leads towards a solution.

This paper provides first results in this area. We use n-queens as an illustrative problem since that has been a primary exemplar of the iterative improvement approach. However, the techniques are general, and should be useful elsewhere. In the next section, we review some empirical data on solution density for the n-queens problem. In the section after that, we present analyses that explain these data. In the final section, we discuss the applicability of the techniques to other problems.

Empirical Results

In this section, we describe empirical data reported by other authors, as well as results of our own experiments. For the latter, we estimated the density of solutions in equilibrium points (henceforth, we will call this the solution/equilibrium density) by starting with a random sample of initial states, running the algorithms to their first equilibrium point, and counting the solutions.[1]

As an initial data point, Minton et al. [1990] report that their MinConflicts Hill-Climbing[2] (MCHC) algorithm applied to the n-queens problem never failed to find a solution for $n \geq 100$. We take this as an indication that the solution/equilibrium density for the method tends to 1 as n tends to infinity.

The MCHC algorithm employs several refinements. For example, the algorithm gets a "head start" on hill-climbing by using a preprocessing stage to produce an initial queen configuration with few attacks. Second, the algorithm permits random "sideways" local modifications that leave the number of queen attacks unaltered. To better understand the value of these refinements, we experimented with a simpler algorithm that

[1]This estimate is biased by the relative sizes of the basins of attraction of the equilibrium points. However, the results below concern gross differences in the density (whether it approaches 0 or 1), and it seems reasonable to assume this is not affected by the bias.

[2]Note that here "hill-climbing" means movement to points of *lower* cost.

starts from a random initial configuration, and makes only modifications that strictly reduce the number of queen attacks. We will call this Simple Hill-Climbing (SHC). Experiments with $n = 1000$ yielded no solutions for this algorithm in a sample of 100 equilibrium points. This suggests that the solution/equilibrium density for SHC may tend to 0 as n tends to infinity.

We also experimented with strict hill-climbing that starts close to a solution. This may be called Head-Start Hill-Climbing (HSHC). The initial configurations were obtained by starting with a fixed solution to the 1000-queens problem and randomly mutating the column positions of the queens on the first 20 rows. This yielded 1 solution in a sample of 100 equilibrium points. We take this as evidence for a limiting density of 0.

Sosic and Gu [1991] describe an interative improvement algorithm for the n-queens problem that maintains configurations with one queen per row and one queen per column. Thus, the column positions of the queens on rows 1 through n form a permutation of the integers from 1 to n. Their QS1 algorithm starts with a random permutation and swaps the columns of queens in different rows to reduce the number of attacks. They report that for $n \geq 1000$, the equilibrium position was always a solution. This suggests a limiting density of 1 for QS1.

To summarize these results, the empirical evidence is consistent with a limiting solution/equilibrium density of 1 for MCHC and QS1, and of 0 for SHC and HSHC.

Analytic Methods

The task in the n-queens problem is to place n queens on an $n \times n$ board such that no two queens are on the same row or column or diagonal. By an *assignment* for this problem, we mean a placement of the queens so that there is one queen on each row. Clearly, the solutions form a subset of the assignments. If we assume a uniform probability distribution on assignments, then the solution/equilibrium density can be expressed as the conditional probability that an assignment is a solution given that it is an equilibrium point.

It does not appear feasible to rigorously compute the probability entirely from first principles. Our approach will be to argue for and adopt reasonable assumptions about the distribution that will enable us to derive the empirically observed results. Note that our analysis is only intended to apply for large values of n.

Now consider a random assignment of the queens. The principal assumption we adopt is as follows.

Assumption 1 *The probability that an arbitrary square is attacked along a diagonal by one of the queens is bounded away from 1, i.e., there is a fixed δ (independent of n) such the probability is less than δ and $\delta < 1$. This is also assumed for* conditional *probabilities of this event, unless there is reason to believe the condition implies otherwise.*

We argue that the assumption is reasonable because there are $4n - 2$ diagonals. Since there are only n queens, each of which sits on 2 diagonals, there must be at least $2n - 2$ unoccupied diagonals. Even if these are the shorter diagonals, they will represent a fraction of the board that is bounded away from zero. Thus, provided any condition is such that the distribution remains roughly uniform as n increases, the probability that an arbitrary square is not in this region should be bounded away from 1.

It is worth noting that a similar assumption with respect to column attacks would be false. Suppose we know that the first $n - 1$ queens do not attack each other. Then they lie on separate columns. Hence, the probability that the last queen is placed on a column that is already occupied is $(n-1)/n$, which tends to 1 as n tends to infinity. The difference here is that there are only n columns, as compared to $4n - 2$ diagonals, so columns are ultimately a scarcer resource.

We now proceed to showing that solutions are dense in equilibrium points for Sosic and Gu's QS1 algorithm, and Minton et al.'s MCHC algorithm. This requires analyzing the cost surface for the problem.

In the n-queens problem, the *cost* of an assignment is the number of queen attacks, i.e., the number of queen pairs that share the same column or diagonal in the assignment. Notice that the total cost of an assignment can be broken down into two components, resulting from column attacks and diagonal attacks, respectively. We can study these separately by considering two variations of the queens problem, which we call the *n-rooks* problem, and the *n-bishops* problem.[3] The n-rooks problem is like the n-queens problem except diagonal attacks are ignored. Thus, any assignment that corresponds to a permutation is a solution. In the n-bishops problem, we ignore column attacks, so that any assignment with no diagonal attacks is a solution. The cost surface for the ordinary queens problem will then be a superposition of the surfaces for the component problems. In the following, the term *simple equilibrium point* refers to an equilibrium point with respect to simple hill-climbing. We will also sometimes use "equilibrium point" without qualification where it is clear from the context which algorithm is involved.

We first consider the n-rooks problem. The following result is easily obtained.

Theorem 1 *In the n-rooks problem, every simple equilibrium point is a solution.*

[3] In chess, a rook attacks along rows and columns, while a bishop attacks along diagonals. A queen combines the attacks of a rook and a bishop.

Proof: Suppose an assignment is not a solution. Then there is some column that contains at least two rooks. It follows that there must also be an empty column. Note that moving one of the doubled rooks to the empty column reduces the number of attacks. Thus, the assignment is not an equilibrium point. ∎

Next we consider the n-bishops problem. In this case, it is not true that every simple equilibrium point is a solution. (Consider, for example, a 4×4 board where the bishops are on the columns 1,3,2,2.) However, it turns out that "almost every" (in a well-defined sense) such point is a solution. The basic intuition behind the proof is quite simple and can be expressed as follows. Suppose an assignment is not an n-bishops solution. Then there is some bishop that is attacked. Consider the $n-1$ other squares that are on the same row as this bishop. Since, by assumption 1, the probability of diagonal attack is bounded away from 1, it follows that, *with high probability* (for sufficiently large n), at least one of these squares is not attacked. Thus, with high probability, the assignment is not an equilibrium point. As we see below, presenting a formal proof based on this idea requires some work.

We have the following easy result.

Lemma 1 *For the n-bishops problem, suppose an assignment is a simple equilibrium point. Then, for each row, either the bishop on that row is not attacked, or else every square on the row is attacked.*

Proof: Immediate. ∎

Now let E_i denote the event that the bishop on the i-th row is not attacked by any of the other bishops, while F_i denotes the event that every square on the i-th row is attacked. Note that E_i and F_i are mutually exclusive. Set $E = E_1 \wedge \ldots \wedge E_n$.

It is easy to see that an assignment is a solution if and only if E holds. Furthermore, if an assignment is a simple equilibrium point, then by Lemma 1, $E_i \vee F_i$ must hold for every i. Thus,

$$\Pr[(E_1 \vee F_1) \wedge \ldots \wedge (E_n \vee F_n)]$$

provides an upper bound on the probability that an assignment is a equilibrium point. We have the following lemma.

Lemma 2 *In the n-bishops problem, there exists a* λ *with* $0 < \lambda < 1$ *such that* $\Pr(F_i) < \lambda^n \Pr(E_i)$ *for every* i *and sufficiently large* n.

Proof: Consider any i. Note that F_i implies a diagonal attack on each of the n squares of row i. Recall that, by assumption 1, the probability of diagonal attack is less than δ, for some $\delta < 1$. Thus, $\Pr(F_i) < \delta^n$ and $\Pr(E_i) > (1-\delta)$. It follows that

$$\frac{\Pr(F_i)}{\Pr(E_i)} < \frac{\delta^n}{1-\delta}.$$

Now choose λ such that $\delta < \lambda < 1$. Clearly

$$\frac{\delta^n}{1-\delta} < \lambda^n$$

for sufficiently large n. The result follows. ∎

Corollary 1 *Suppose* $A = A_1 \wedge \ldots \wedge A_n$, *where for each* i *either* $A_i = E_i$ *or* $A_i = F_i$. *Then, under the conditions of the lemma, there exists* λ, *with* $0 < \lambda < 1$, *such that*

$$\Pr(A) < \lambda^{nk} \Pr(E)$$

where k *is the number of values of* i *for which* $A_i = F_i$.

Proof: Let $\{i_1, \ldots, i_k\}$ be the values of i for which $A_i = F_i$. By an argument similar to that of the lemma, we get $\Pr(F_{i_1} \wedge \ldots \wedge F_{i_k}) < \delta^{nk}$ and $\Pr(E_{i_1} \wedge \ldots \wedge E_{i_k}) > (1-\delta)^k$. It follows that

$$\Pr(F_{i_1} \wedge \ldots \wedge F_{i_k}) < \lambda^{nk} \Pr(E_{i_1} \wedge \ldots \wedge E_{i_k})$$

for a suitable λ and sufficiently large n. Since assumption 1 allows us to ignore irrelevant conditioning on probabilities, we can add $\bigwedge_{i \notin \{i_1,\ldots,i_k\}} E_i$ to the conjuncts on both sides, giving

$$\Pr(A) < \lambda^{nk} \Pr(E).$$

∎

We are now ready to prove the following result.

Theorem 2 *In the n-bishops problem, the density of solutions in simple equilibrium points approaches 1 for large n.*

Proof: Note that if $(E_1 \vee F_1) \wedge \ldots \wedge (E_n \vee F_n)$ is converted to disjunctive normal form, the number of conjuncts with k occurrences of the F_i propositions will be $\binom{n}{k}$. Thus, by Corollary 1,

$$\begin{aligned} &\Pr[(E_1 \vee F_1) \wedge \ldots \wedge (E_n \vee F_n)] \\ &< \Pr(E)[1 + \ldots + \tbinom{n}{k}(\lambda^n)^k + \ldots + (\lambda^n)^n] \\ &= \Pr(E)(1+\lambda^n)^n \end{aligned}$$

Since $(1+\lambda^n)^n$ tends to 1, we have

$$\lim_{n\to\infty} \frac{\Pr(E_1 \wedge \ldots \wedge E_n)}{\Pr((E_1 \vee F_1) \wedge \ldots \wedge (E_n \vee F_n))} = 1.$$

The result follows. ∎

It may be remarked that empirical testing produces results consistent with Theorem 2.

We now return to consideration of the ordinary n-queens problem. One might expect some fraction of the n-rooks and n-bishops solutions to intersect, giving solutions to the full queens problem. Note, however, that the n-rooks solutions are isolated, i.e., each of them is surrounded (within one step) by non-solution assignments. It thus seems reasonable to expect many of the n-rooks solutions to generate non-solution equilibrium points in the full problem. As we will see later, this makes simple hill-climbing ineffective for generating solutions.

One way of dealing with this problem might be to somehow "factor out" column attacks by performing hill-climbing while sticking to assignments that are already solutions to the n-rooks problem, i.e., permutations. In this case, the restricted cost surface might be expected to resemble that for the n-bishops problem. It is well-known that the space of permutations can be traversed by 2-step transpositions. This suggests generalizing the notion of simple hill-climbing to that of hill-climbing with bounded lookahead. A *k-step* hill-climbing algorithm is allowed to search k steps from the current assignment looking for one of lower cost. Each equilibrium point of a k-step hill-climbing algorithm will be called a *k-step equilibrium point.* In the queens problem, we are particularly interested in 2-step equilibrium points. We have the following results.

Lemma 3 *In the n-queens problem, the probability that a 2-step equilibrium point is a permutation approaches 1 for sufficiently large n.*

Proof: Informally, the idea of the proof is as follows. Suppose an assignment is not a permutation. Then there is some column c_1 that contains at least 2 queens, and some other column c_2 that is free of queens. Take one of the queens on c_1, and move it to some other column j. If there is already one or more queens on that column, choose one such arbitarily and move it to c_2. Since there are $n-1$ possible choices for j, it follows from assumption 1, that with arbitrarily high probability (for sufficiently large n), we can choose j so that the moved queen(s) will now be free of diagonal attacks. Thus, we have not increased the count of diagonal attacks. But the count of column attacks has been decreased. Thus, the original assignment was not a 2-step equilibrium point.

A more formal proof is similar to that of Theorem 2. In this case E_i would represent the event that the queen on row i is not attacked along a column by any of the other queens, while F_i would represent the event that it is so attacked and, moreover, none of the possible choices for j above frees the moved queen(s) of diagonal attacks. It is not hard to see that the probability of this F_i declines exponentially with n, as required in the proof. (In this case the probability of E_i declines linearly with n, but this does not impede the proof.) ▮

Lemma 4 *In the n-queens problem, the probability that a 2-step equilibrium point is free of diagonal attacks approaches 1 for sufficiently large n.*

Proof: We present the proof informally. Suppose the assignment is not free of diagonal attacks. Then there must be some queen that is attacked along a diagonal. Take that queen, and move it to some other column j. If there is already one or more queens on that column, choose one of them arbitarily and move it to the column vacated by the first queen. Since there are $n-1$ possible choices for j, with arbitrarily high probability, we can choose j so that the moved queen(s) will now be free of diagonal attacks. Note that the column attack count has not been increased. But the count of diagonal attacks has been decreased. Thus, the original assignment was not a 2-step equilibrium point.

Again, the proof can be made formal along the lines of Theorem 2. ▮

Theorem 3 *In the n-queens problem, the density of solutions in 2-step equilibrium points approaches 1 for large n.*

Proof: Immediate from Lemmas 3 and 4. ▮

We can use these results to explain the performance of the QS1 algorithm. Recall that this algorithm performs hill-climbing that swaps columns to reduce the number of diagonal attacks, i.e., it moves from permutation to permutation in 2-step jumps. Essentially the same proof as that of Lemma 4 can be used to show that, with arbitrarily high probability, the equilibrium permutation is free of diagonal attacks, i.e., it is a solution. This explains the high density of solutions encountered by this algorithm.

In order to understand the behavior of MCHC, we divide simple equilibrium points that are not solutions into two categories: we will say such a point is a *pit* if every path from the point to a region of lower cost must pass through a region of higher cost; otherwise, the point is a *plateau.* (In the case of a plateau, we can reach a region of lower cost by passing through points of equal cost.) It is easy to see that MCHC will eventually escape from a plateau because of its random sideways movements. Note that MCHC either reaches a solution, or ends up cycling randomly among a fixed group of points with equal cost. Since MCHC escapes from plateaus, every such point must be a pit. We have the following lemma.

Lemma 5 *In the n-queens problem, every pit is a 2-step equilibrium point.*

Proof: We will show that if a point is not a 2-step equilibrium point, then it cannot be a pit.

Suppose a double queen movement leads to a lower cost. If neither queen move singly lowers the cost, this can only be because the two queens attack each other after one is moved first. But there is at most one attack between two queens. Thus, the intermediate state, at worst, has equal cost to the original state. ▮

By theorem 3, for large n, almost every 2-step equilibrium point is a solution; thus, there must be few pits relative to solutions. It follows that the frequency with which MCHC terminates in a solution should approach 1 for large n.

We now consider the negative data described in the section on empirical results. These data suggest that for SHC (simple hill-climbing) the solution/equilibrium density tends to 0. This is borne out by the following result.

Theorem 4 *In the n-queens problem, the solution/equilibrium density for SHC tends to 0 as n tends to infinity.*

Proof: Suppose an assignment is a solution to the n-queens problem. Let q be a fixed queen. Consider the $n-1$ possible new positions arrived at by swapping columns between q and some other queen. The probability of diagonal attack on the squares to which the queens are moved is bounded away from 1 by assumption. We can also assume it is bounded away from 0, since at least $n-2$ of the diagonals are occupied by queens. It follows that the probability that the swap leads to *exactly one* diagonal attack is bounded away from 0.

Now let k be some arbitrary number. For sufficiently large n, with arbitrarily high probability, at least k of the possible swaps must lead to a position involving exactly one diagonal attack. Any such position is an equilibrium point under simple hill-climbing because a single queen movement from the position necessarily produces a column attack.

The above shows that for every solution, with high probability, there are at least k non-solution equilibrium points. However, there remains the possibility that these overlap for different solutions. We can deal with this consideration if we assume that solutions are distributed roughly uniformly across assignments. Since the density of solutions in assignments declines rapidly with n, and the equilibrium points exhibited above are within a bounded distance of the corresponding solution, it follows that the amount of overlap is ultimately not significant. Therefore, the solution/equilibrium density is less than $1/k$ for sufficiently large n. But k is arbitrary. Thus, the solution/equilibrium density tends to 0. ▮

An examination of the proof of Theorem 4 shows it applies equally well to HSHC.[4] These results suggest that "sideways" local change—not the "head start" preprocessing algorithm—is the important factor leading to the high density of solutions for MCHC in the queens problem.

Discussion

In this section, we will try to place the results above in perspective, and see how they might apply to other problems. We are particularly interested in constraint satisfaction problems (CSPs), of which the n-queens problem is an example. Informally, a CSP consists of a set of variables, each of which is assigned a value from a set called the *domain* of the variable. The possible assignments are restricted by a set of *constraints*, which mandate relationships or restrictions between the values of different variables. The reader is referred to Dechter [1990] for the formal definition of a CSP.

We require some additional terminology. A simple equilibrium point that is not a solution will be called a *basin*. The *radius* of a basin is defined to be the number of steps required to escape from it, i.e., the minimum number of steps needed to reach a region of lower cost. Notice that saying solutions have high density in k-step equilibrium points is equivalent to saying that basins with radius greater than k are rare compared to solutions.

We can summarize the results of the previous section as follows. In the n-queens problem, for large n, basins of radius greater than 2, and hence pits, are sparse relative to solutions, but plateaus of radius 2 are quite common. This suggests that, visually, each solution appears surrounded by "stairs," with steps of width 2, rather like in an amphitheatre. Notice that the cost surface appears smooth when viewed at a coarse resolution.

The proof method used in the queens problem can be generalized to support the following observation: *If a problem area has the property that an arbitrary constraint violation can be removed within k steps with low probability of introducing a fresh violation, then hill-climbing with k-step lookahead will be an effective solution technique.*

The n-bishops example demonstrates one way of showing that the low probability criterion is met. If a CSP has the property that the domain size of each variable is at least comparable to the number of variables, and the probability that any one value is in conflict is bounded away from 1, then resetting a single variable will generally suffice to eliminate one conflict, in large problems. Actually, the condition on the probability can be relaxed somewhat: instead of being bounded away from 1, it is enough to suppose it does not approach 1 too quickly. Recall that the crucial property used in the proof of Theorem 2 is that $(1+\lambda^n)^n$ tends to 1. It turns out that this will be true even if λ increases with n provided at least that $\lambda < 1 - 1/n^\alpha$ for some $\alpha < 1$.

A further point to observe is that the constraints in the queens problem fall into two categories. Constraints corresponding to diagonal attacks meet the low probability criterion discussed above. However, the column attack constraints do not have this felicitous property. We will refer to constraints that do not meet the low probability criterion as *tight* constraints. The analysis in the queens problem shows that if the space is reformulated so that tight constraints are somehow "factored out," then hill-climbing can be made effective. We can draw an analogy to game-playing programs where search continues until a *qui-*

[4] However, the density may approach zero at a slower rate with HSHC.

escent position is reached at which evaluation takes place. This may, for example, involve following down all possible sequences of captures. In the case of CSP solving using the techniques discussed here, quiescence would require that none of the tight constraints are violated. For example, in the QS1 algorithm, a potential queen move must be accompanied by movement of a second queen so that a situation involving no column attacks is maintained. In terms of our observation about hill-climbing with k-step lookahead, we remark that the lookahead need not involve a full-width search; instead, the search may be tailored to the structure of a particular problem. The queens problem also suggests that examining variant problems that ignore one or more categories of constraints may provide a useful analysis tool for formulating the search strategy.

The distinction between tight and non-tight constraints suggests a piece of practical advice in organizing an environment for scheduling or other activities requiring constraint satisfaction. It may be wise to allow a certain amount of slack in providing resources for the tasks. For example, we can predict from the analysis here that simple hill-climbing should be effective, for large n, in placing $n/2$ queens without attack on an $n \times n$ board (because then the probability of both column and diagonal attacks would be bounded away from 1). Similarly, suppose n tasks requiring an exclusive resource of a particular type need to be performed within m time slots, so that at least n/m copies of the resource are required. The results here indicate that the allocation problem may be simplified if, say, cn/m instances of the resource are available, where $c > 1$.

An additional observation is that certain categories of constraints appear to be associated with basins of a definite size. For example, constraints associated with a fully subscribed resource (such as columns in the n-queens problem) tend to require swaps to make progress, i.e., they produce basins with a characteristic radius of 2. Similarly, equality constraints would generate basins of radius 2. Of course, basins resulting from separate constraints may, by random coincidence, occur next to each other and coalesce into a basin of larger size. The Central Limit Theorem of probability theory suggests that for problems with sufficiently randomized constraints, the basin sizes should occur in a normal distribution with a mean that depends on the relative prevalence of constraints. Note that this does not necessarily increase in proportion to the size of the problem.

This raises the possibility that for important classes of problems, hill-climbing with bounded lookahead might perform well. In [Morris, 1991], a hill-climbing algorithm is proposed that fills in basins as it goes along, so that it always reaches a solution if one exists. Roughly speaking, the algorithm simulates k-step hill-climbing for *every* k at a cost bounded by V steps per simulated k-step. Here V is the "volume" of a basin of radius k, i.e., the number of steps required to fill it. This indicates that the algorithm should perform well in problems where the average basin size is small.

Conclusion

Using the n-queens problem as an illustrative example, we have shown that, under certain conditions, hill-climbing algorithms can enjoy the property that almost all equilibrium points will be solutions. The analysis explains the success of some algorithms previously reported in the literature. Furthermore, it predicts circumstances under which more general constraint problems may be amenable to similar approaches.

Acknowledgement The author is grateful to Rina Dechter for suggestions concerning this work.

References

R. Dechter. Enhancement schemes for constraint processing: backjumping, learning, and cutset decomposition. *Artificial Intelligence*, 41(3), 1990.

S. Minton, M. D. Johnston, A. B. Philips, and P. Laird. Solving large scale constraint satisfaction and scheduling problems using a heuristic repair method. In *Proceedings of AAAI-90*, Boston, 1990.

P. Morris. Solutions Without Exhaustive Search. In *Proceedings of AAAI-90 Workshop on Constraint-Directed Reasoning*, Boston, 1990.

P. Morris. *An Iterative Improvement Algorithm With Guaranteed Convergence.* Technical Note TR-M-1991-1, IntelliCorp, Mountain View, California, 1991.

R. Sosic, and J. Gu. 3,000,000 Queens in Less Than One Minute. *Sigart Bulletin*, 2(2), 1991.

M. Zweben. A Framework for Iterative Improvement Search Algorithms Suited for Constraint Satisfaction Problems. In *Proceedings of AAAI-90 Workshop on Constraint-Directed Reasoning*, Boston, 1990.