OPTIMAL SEARCHES FROM AND AND OR NODES[1]

Jeffrey A. Barnett

USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, California 90291

ABSTRACT

The problem is to organize search from an AND node in a way that minimizes expected cost. The result is derived as a corollary to earlier work of Simon and Kadane. It is shown that, unless knowledge gained during the search changes the probability or cost estimates of remaining parts of the search, the original a priori strategy remains optimal. The effect of approximating the search statistics used to determine the optimal strategy is examined, and it is found that the impact on expected cost is linearly bounded by the quality of the approximation. Then the case of searching an infinite conjunct is considered. Finally, some related research topics are discussed.

1. INTRODUCTION

> Suppose that an unknown number, possibly zero, of treasure chests have been buried on a random basis at some of n sites, but that neither the sites or the depths of burial are known with certainty. At each site a sequence of one-foot slices can be excavated, and a treasure may be disclosed by the removal of any one of these slices. The probability that a treasure lies just below any specified slice is known, as is the cost of excavating that slice. Develop a strategy with the minimum expected cost to find a treasure.

Simon and Kadane[2] have solved the stated problem. It is a metaphor for satisficing search from an OR node. Each site is a disjunct, and finding a treasure at any site is sufficient to terminate the search. Excavating a slice is analogous to performing the next atomic action in the search at one of the disjuncts, e.g., expanding a node. The given probabilities and costs correspond respectively to the probability of a node expansion producing a goal node and the cost of expanding the node.

This paper develops an optimal search strategy from an AND node. Here, it is necessary to find a treasure at *every* site. If the search succeeds, the order of excavating slices makes no difference because the search at each site terminates exactly when a treasure is found at that site. The search can fail only if the bottom of some site is reached without finding a treasure. Therefore, the best strategy should quickly discover a failure, if any, to preclude further search.

Section 2 presents the best strategy for searching from OR nodes and derives the result for AND nodes as a corollary. The best strategy first digs at one site until either a treasure is found or the bottom is reached. In the first case, another site is selected, etc. In the latter case, the entire search is concluded. It is never optimal to switch back and forth between sites. Section 3 develops notation that is used in the remainder of the paper.

Section 4 considers what happens if knowledge is gained during the search. It is shown that unless the probabilities or costs associated with unexcavated slices are changed, the a priori strategy remains optimal. Section 5 examines the problem of determining approximations to the optimal strategy when the probabilities and costs are not available. Section 6 looks at the problems of a site being infinitely deep. Finally, Section 7 lists some important open questions related to these results

2. THE SIMON-KADANE RESULT

2.1. Optimal Searching from OR Nodes

If s is a slice, let p(s) be the probability that a treasure is just below s and c(s) be the cost of excavating s. Let b = s1..s where the s_1 are slices. Then b is an (admissible) *strategy* when (1) the s are distinct, (2) if s and s_1 are slices from the same site and i<j, slice i is above slice j, and (3) if two slices from a site are in b, all the slices between these two from that site are in b too. A strategy is a proposed order for excavating slices.

The notation is generalized in a natural way so that p(b) is the probability that strategy b will unearth a treasure and c(b) is the *expected* cost of following strategy b. (Recall that execution ceases if a treasure is found, hence some of the work proposed by b may not be done.) Define the benefit-to-cost ratio as <p(b) = p(b)/c(b) with the assumptions that p(s)not =1 and c(s)not=0, where s is a slice. The assumptions entail respectively that no slice or site contains a treasure with certainty and that no slice or site can be excavated for free.

Description of the optimal strategy for searching from OR nodes depends on the idea of *maximal indivisible blocs.* A bloc is any consecutive set of slices from the same site. Let b = $s_1 \ldots s_r$ be all the slices from a single site ordered from top to bottom; thus, b is a strategy. If (Sj)<<p(s_{j+1}), join Sj and s_{j+} as a bloc replacing the original pair. Continue this process until it terminates and b has been divided into a sequence of blocs, b = b.-.b , where each b is

[1] This research is supported by the Defense Advanced Research Projects Agency under Contract No MDA903 81 C 0335. Views and conclusions contained in this report are the author's and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them.

[2] Simon, H. A., and J B. Kadane. "Optimal Problem-Solving Search. All or None Solutions," *Artificial Intelligence* 6 (1975), 235-247.

a bloc such that $\varphi(b_i) \geq \varphi(b_{i+1})$. These b_i are the maximal indivisible blocs for the site and depend only on the statistics of the slices at that site.

Let b = $b_1 \ldots b$, where the b. are now all the maximal indivisible blocs from all the sites. Simon and Kadane prove that strategy b is *optimal* if and only if $\varphi(b_i) \geq \varphi(b_{i+1})$. Thus, the optimal strategy is unique up to permutations of maximal indivisible blocs with equal φ's. Note that even though the maximal indivisible blocs from a single site must be ordered in b from top to bottom, they need not be contiguous in the optimal strategy.

2.2. Applying the Result to Searches from AND Nodes

Searching from AND nodes is the dual of searching from OR nodes. If the search is to succeed, *every* site must disclose a treasure. Further, if the search is successful, the total amount of work is independent of the order in which the slices are excavated. (Since the constraint of excavating top slices before bottom slices is still in force, each site is excavated until it discloses its first treasure.) A *true treasure* in an AND search is to reach the bottom of a site without finding anything: Since the AND search can no longer succeed, the search remaining at the other sites can be abandoned.

From this viewpoint, it is easy to see that a true treasure can be discovered only by excavating the *bottom* slice of some site. The probability of finding a true treasure at any other slice is 0. For the bottom slice, the a priori probability of success is the probability that the entire site is free of ordinary treasures. Thus, AND search is just like OR search where we are looking for true treasures.

It is obvious that the maximal indivisible blocs are entire sites since only blocs containing the last slice can have nonzero expectations of success, hence positive φ. Therefore, the optimal strategy orders excavation (of entire sites) by decreasing values of the ratio of expected failure at the site (the probability of finding a true treasure) to the expected cost of excavating the site. The probabilities associated with the slices affect both the probability of search termination at the site, hence its expected cost, and the probability that the site contains a true treasure.

The next sections give explicit expressions for all the quantities involved.

3. FORMALITIES

If s is a slice, then p(s) is the probability that excavating s discloses a treasure and q(s) = 1-p(s). Further, c(s) is the cost of excavating the slice. Let b be an entire site, where b = $s_r \ldots s$ and the s1 are the slices of b in top-to-bottom order. We define p(b) as the probability that site b contains at least one treasure and q(b) = 1-P(b). It is obvious that

$$q(b) = \prod_{1 \leq i \leq n} q(s_i)$$

Now define E(b) as the expected cost of excavating site b. Evidently,

$$E(b) = \sum_{1 \leq i \leq n} c(s_i) \prod_{j<i} q(s_j)$$

because slice *s.* is excavated and the cost $c(S_i)$ paid only if no slice above s. discloses a treasure; the probability of this is represented by the product in the formula for E.

Since a true treasure in an AND search is to reach the bottom of a site without finding an ordinary treasure, the probability of success to cost ratio is $\varphi(b) = q(b)/E(b)$. Therefore, the optimal strategy orders site excavations in terms of decreasing φ values.

Let $b_1 \ldots b_n$ be the strategy of first digging at site b_1 if a treasure is found, start digging at site b_2, etc. If the bottom of some site is reached without disclosing a treasure, search ceases. Then it is evident that E, the expected work for this strategy, is

$$E(b_1 \ldots b_n) = \sum_{1 \leq i \leq n} E(b_i) \prod_{j<i} q(b_j)$$

It is assumed below that p(s)*I and $c(s)>0$, i.e., no slice discloses a treasure with certainty and no slice can be excavated without cost. Further,, $q(b) \neq 1$ is assumed, meaning that no site is known in advance not to contain a treasure.

4. DYNAMIC KNOWLEDGE

So far, we have insisted that a strategy be formed before the start of excavation and not modified afterwards. Knowledge gained during the search has only two effects: (1) the remainder of the strategy is abandoned if the bottom slice of any site is excavated without disclosing a treasure, and (2) if a treasure is found at a site, the remaining slices at that site are not excavated.

However, the knowledge that the slices already excavated did not disclose a treasure becomes available during the search. Assume that the portion of the optimal strategy not yet executed is $s_1 \ldots s_n b_2 \ldots b_m$, where the s are the remaining slices at site 1 and tne b. are the sites that have not yet been searched. Since this is an optimal strategy, $\varphi(b_i) \geq \varphi(b_{i+1})$ for $i \geq 2$. Assume further that $\varphi(s_1 \ldots s_n) \geq \varphi(b_2)$; this is surely true initially.

If s_1 is excavated and discloses a treasure, then $b_2 \ldots b_m$ is obviously the optimal strategy for the rest of the work. If s_1 is a bottom slice, there is no problem in any event. If s_1 does not disclose where to insert the "site" $s_2 \ldots s_n$ into the sequence $b_2 \ldots b_m$. If it can be shown that $\varphi(s_2 \ldots s_n) \geq \varphi(s_1 \ldots s_n)$, the original strategy will still be the best because $\varphi(s_2 \ldots s_n) \geq \varphi(b_i)$ for $2 \leq i \leq n$. We have

$$\begin{aligned}\varphi(s_1 \ldots s_n) &= q(s_1 \ldots s_n) / E(s_1 \ldots s_n) \\ &= \prod_{1 \leq i \leq n} q(s_i) \Big/ \sum_{1 \leq i \leq n} c(s_i) \prod_{1 \leq j<i} q(s_j) \\ \varphi(s_2 \ldots s_n) &= q(s_2 \ldots s_n)/E(s_2 \ldots s_n) \\ &= \prod_{2 \leq i \leq n} q(s_i) \Big/ \sum_{2 \leq i \leq n} c(s_i) \prod_{2 \leq j<i} q(s_j).\end{aligned}$$

Notice that $\varphi(s_2 \ldots s_n)$ is computed as if s_1 had never existed. If we assume that $\varphi(s_2 \ldots s_n) < \varphi(s_1 \ldots s_n)$, then

$$\prod_{2 \leq i \leq n} q(s_i) \Big/ \sum_{2 \leq i \leq n} c(s_i) \prod_{2 \leq j<i} q(s_j) < \prod_{1 \leq i \leq n} q(s_i) \Big/ \sum_{1 \leq i \leq n} c(s_i) \prod_{1 \leq j<i} q(s_j)$$

$$\sum_{1 \leq i \leq n} c(s_i) \prod_{1 \leq j<i} q(s_j) < q(s_1) \sum_{2 \leq i \leq n} c(s_i) \prod_{2 \leq j<i} q(s_j)$$

$$\sum_{1 \leq i \leq n} c(s_i) \prod_{1 \leq j<i} q(s_j) < \sum_{2 \leq i \leq n} c(s_i) \prod_{1 \leq j<i} q(s_j)$$

$$c(s_1) < 0.$$

But it has been assumed that no slice can be excavated without cost; thus, $c(s_1)>0$. Therefore, a contradiction has been reached. The immediate conclusion is

> If excavating slices does not alter the probabilities and costs associated with unexcavated slices, the remainder of the optimal a priori strategy is still the best

strategy for the remainder of the search, i.e., the optimal a priori strategy is the optimal dynamic strategy.

A similar result is true for OR searches because of the nature of the construction of the maximal indivisible blocs; the φ's within the joined blocs do not decrease.

5. APPROXIMATIONS

For both the optimal AND and OR searches described above, it appears necessary to know the p and c values for the individual slices. For the case of an OR search, estimation is a difficult process. A reasonable tactic is to estimate p and c, hence φ, for each indivisible bloc. However, this begs the question of how to identify these blocs in the first place.

The estimation problem for AND searches is relatively easier because the blocs are entire sites. If the φ's for the sites can be approximated, the optimal strategy based on the approximations should not be very far from optimal. In fact, the exact φ values are not needed; only their relative magnitudes make any difference.

In order to show the effect of approximating φ's, assume for some r>l that $\varphi(b_1) = r\varphi(b_2)$, where b_1 and b_2 are two sites. Obviously b_1 ought to be excavated before b_2. But what if $\varphi(b_2)$ is over-estimated by a factor of r so that the sites are excavated in the wrong order? With a little algebra, it can be shown that the relative amount of extra work is

$$[E(b_2b_1) - E(b_1b_2)]/E(b_1b_2)$$
$$= (r-1)q(b_1)q(b_2)/[q(b_1) + rq(b_2) - rq(b_1)q(b_2)]$$

and the maximum of this expression occurs as $q(b_1)$ and $q(b_2)$ approach 1, because the partial derivatives with respect to $q(b_1)$ and $q(b_2)$ are strictly positive—recall that q's are always less than 1. The limit value is r-1. Therefore, if $r = 1 + \varepsilon$, the maximum fraction of expected extra cost is just e.

Whether or not approximations are used, knowledge may be gained during the search. If the knowledge does not alter the measures associated with unexcavated slices (or blocs or sites, for that matter), the previous section showed that the optimal strategy is unaffected. On the other hand, when the measures are approximated, it is likely that excavation will change some of them, particularly the measures for the remainder of the site of current excavation. If the estimate of φ increases, there is no need to reconsider the ongoing strategy—it is still optimal. The strategy need be revised only if the estimate for φ of the remainder of the site drops below φ for the next site in the strategic order. Little more work is necessary if the search at one site affects the estimates of φ for other sites.

6. BLACK HOLES AND TAR PITS

So far it has been assumed implicitly that a site has only a finite number of slices. This entails a site having a bottom slice. Though this is an obvious pragmatic assumption, it is amusing to consider a case of an infinite site, B, with slices s_i, where $i \geq 1$. As an example, define p and c as

$$p(s_n) = (n+1)^{-2}$$
$$c(s_n) = 2n^{-1}(n+1)^{-1}.$$

It is easy to see that

$$\prod_{1\le i\le n} (1 - p(s_i)) = \prod_{1\le i\le n} (1 - (i+1)^{-2})$$
$$= (n+2)/(2n+2)$$

and the limit of the product as n gets large is ½. By stretching the definitions of φ, q, and E to the infinite case, we have

$$q(B) = \Pi(1 - p(s_i)) = ½$$

$$E(B) = \Sigma c_i \prod_{j<i}(1 - p(s_j))$$
$$= \Sigma 2n^{-1}(n+1)^{-1}(n+1)/2n$$
$$= \Sigma n^{-2}$$
$$= \pi^2/6.$$

Therefore, $\varphi(B) = q(B)/E(B) = 3/\pi^2$.

Is it lunacy to expend any effort during an AND search at an infinite site before all the finite sites have been tried? This is not a trap for OR searches; there it makes as much sense to explore a portion of an infinite site as it does a portion of a finite site—in fact, the right amount to search is a maximal indivisible bloc. But is an infinite site a black hole that soaks up all our resources in an AND search? Actually, no. If $\varphi(B) \neq 0$, then $q(B) > 0$ and $E(B) < \infty$. But in that case, $\Sigma c(s_i) < \infty$ is necessarily entailed; for the example here, the value of the sum is 2. Thus, even though there may be an infinite number of slices, the total maximum work at the site must be finite or will be 0, so that the site will be searched last in any event.

7. FUTURE DIRECTIONS

In addition to the work described in Section 2.1, Simon and Kadane solved the problem of finding the optimal strategy to search OR graphs. An obviously important generalization is to develop optimal search strategies for AND/OR graphs. Several of us have tried this problem and failed to produce any noteworthy results. The dilemma appears to be that the expected cost function depends in a nontrivial way on the exact execution results of all prior excavations.

A companion problem is what to do when only approximations of the cost and probability measures are available. For AND/OR graph searches, it is not yet clear even what needs to be estimated. Do we need estimates for nodes as a whole or only for the next maximal indivisible bloc; if for the next maximal indivisible bloc, then what does one look like? Also, is it necessary to have separate estimates for the expected costs of working below a node if the search succeeds and if it fails?

Another outstanding problem concerns the use of dynamic knowledge: If we are told in advance how the excavation of one slice might affect the measures associated with unexcavated slices, what is the optimal strategy? Further, the definition of what a treasure is may depend on the site being searched. It is even possible that a treasure unearthed at one site can simultaneously satisfy the requirements for a treasure at another site, hence eliminating the need to search that site, and this knowledge can affect the optimal order of excavation.