COPING WITH ERRORS IN BINARY SEARCH PROCEDURES

(Preliminary Report)

R.L. Rivest, A.R. Meyer, D.J. Kleitman, K. Winklmann*

Massachusetts Institute of Technology, Cambridge, MA, and

J. Spencer

State University of New York at Stony Brook, Stony Brook, N.Y.

Abstract: We consider the problem of identifying an unknown value $x \in \{1,2,\ldots,n\}$ using only comparisons of x to constants when as many as E of the comparisons may receive erroneous answers. For a continuous analogue of this problem we show that there is a unique strategy that is optimal in the worst case. This strategy for the continuous problem is then shown to yield a strategy for the original discrete problem that uses $\log_2 n + E \cdot \log_2 \log_2 n + O(E \cdot \log_2 E)$ comparisons in the worst case. This number is shown to be optimal even if arbitrary "Yes-No" questions are allowed.

We show that a modified version of this search problem with errors is equivalent to the problem of finding the minimal root of a set of increasing functions. The modified version is then also shown to be of complexity $\log_2 n + E \cdot \log_2 \log_2 n + O(E \cdot \log_2 E)$.

INTRODUCTION

Let x be an unknown number which we wish to identify by asking "Yes-No" questions about it. Our objective is to minimize the number of questions required in the worst case, taking into account that some of the answers received may be erroneous.

This problem comes in several versions based on the possible values of x and the nature of the "Yes-No" questions. In the discrete case x is a member of the finite set $\{1,2,\ldots,n\}$; in the continuous case x is a member of the half-open real interval $(0,1]$. In the continuous case we will not in general be able to identify x exactly with a finite number of questions; rather we try to minimize the size, ϵ, of the region which is known to contain x. In either case let U denote the universe of possible values of x. Independent of the choice of U, the allowable questions may in one case be arbitrary "Yes-No" questions about x, i.e. questions of the form "Is $x \in T$?" where T is a specified subset** of U, or in the other case the allowable questions may be restricted to comparisons, i.e. questions of the form "Is $x \leq c$?" where c is a specified element of U.

We carry out a worst-case analysis of the preceding identification problem under the additional assumption that up to E of the answers may be incorrect. This bound E may be a constant or a function of n (the size of the universe) in the discrete case or of ϵ (the desired size of the region found to contain x) in the continuous case. We are mainly interested in strategies that use comparison questions only and consider strategies with arbitrary "Yes-No" questions only in order to derive lower bounds for the complexity of comparison strategies. It is worth noting that in none of the cases we consider does the restriction to comparison questions cause any significant increase in the number of questions needed in the worst case.

In Section II we answer an open problem from [5] by showing that there is a unique optimal strategy for the continuous

* This research was sponsored by NSF Grants MCS77-19754, MCS76-14294, and MCS77-19754A01 and ONR Grant NOOO-14-76-C-0036.

** In the continuous case we restrict the choice of T to measurable subsets of U.

problem with only comparison questions allowed. For a given number Q of questions the strategy minimizes the worst-case size of the region A which is found to contain the unknown x.

This comparison strategy for the continuous case does not translate into a strategy for the original discrete problem just by setting $\epsilon=1/n$. The reason for this is that the region A which is found to contain x may not be an interval. In Section III we show that E extra comparison questions are always sufficient to cut A down to an interval. This then yields a comparison strategy for the discrete problem using no more than $\lg n+E\cdot\lg\lg n + O(E\cdot\lg E)$ questions* in the worst case, which again is optimal even among strategies using arbitrary "Yes-No" questions about x.

In Section IV we consider the problem of finding the minimum root of a set of increasing functions. This problem is considered by Gal et. al. in [3]. It arises in determining how far a point in a feasible region of a mathematical programming problem can be translated in a given direction before striking a boundary. The latter is a familiar subproblem of the common gradient-following approach to nonlinear programming (cf. [3]). We show that this minimum root identification problem (without errors) is equivalent to a variation of the continuous search problem with errors considered in Section II. The variation is that errors may occur in "No"-answers only. The $\lg n+E\cdot\lg\lg n+O(E\cdot\lg E)$ result from Section II then immediately applies as an upper bound and we prove that it also is a lower bound. Establishing the equivalence of the two problems and establishing lower bounds were stated as two open problems in [5].

In Section V we mention some open problems.

We do not consider the case where the bound E on the number of erroneous answers is a function of the number Q of questions. The discrete version of this problem where arbitrary "Yes-No" questions are allowed and where E is a linear function of Q is equivalent to the problem of optimal block coding for a noisy channel with a noiseless and delayless feedback channel. This equivalence was pointed out in [2] and [6]. The coding problem has been studied and asymptotic bounds on the "transmission rate" $(\lg n)/Q$ have been given in [1] and [6].

In contrast, we consider the case of strategies that use only comparison questions and where E is independent of Q (but may depend on n).

* We use lg to denote $\log_2$ throughout this paper.

II. THE CONTINUOUS CASE

In this section we give a complete answer to the following questions: Given $\epsilon>0$, how many comparison questions about an unknown $x\in(0,1]$ are necessary in the worst case to determine a subset A of (0,1] of size less of equal to ϵ which contains x? Note that we do not require A to be an interval. (It will, however, be a union of disjoint intervals; the size of such a set is the sum of the sizes of its component intervals.)

To state the following Theorem it is convenient to reverse the situation and let the number Q of questions be given and to determine the smallest ϵ achievable with Q questions. We use $\left(\binom{n}{m}\right)$ to denote $\sum_{i=0}^{m}\binom{n}{i}$.

Theorem 1: For any two nonnegative integers Q and E let $\epsilon(Q,E)$ denote the smallest ϵ such that Q arbitrary "Yes-No" questions about an unknown $x\in(0,1]$, up to E of which may receive erroneous answers, are sufficient in the worst case to determine a subset A of (0,1] with $x\in A$ and $|A|\leq\epsilon$. Then $\epsilon(Q,E) = \left(\binom{Q}{E}\right)\cdot 2^{-Q}$.

This smallest ϵ can be achieved by a strategy using only comparison questions. This optimal comparison strategy is unique.

Proof: The state of knowledge of the questioner at any point during the questioning can be summarized by the number q of questions remaining and the E+1tuple $A_q=(A_q^0,\ldots,A_q^E)$ defined by

$x\in A_q^e$ iff exactly e of the previous answers were incorrect.

Clearly we always have $x\in\bigcup_{e=0}^{E}A_q^e$. Define the weight w of a state A_q when there are q questions remaining by:

$$w(q,A_q) = \sum_{e=0}^{E}\left(\binom{q}{E-e}\right)\cdot|A_q^e|.$$

(This definition is due to Berlekamp [1].)

We now define an "adversary answering strategy" to answer the next question. Suppose the question is "Is $x \in T$?". A "Yes"-answer then results in a state described by

${}^yA_{q-1}=({}^yA^0_{q-1},\ldots,{}^yA^E_{q-1})$ with

$${}^yA^e_{q-1} = (A^e_q \cap T) \cup (A^{e-1}_q - T) \text{ for } 1 \le e \le E, \text{ and}$$

$${}^yA^0_{q-1} = A^0_q \cap T.$$

Similarly, a "No"-answer results in a state described by

${}^nA_{q-1} = ({}^nA^0_{q-1},\ldots,{}^nA^E_{q-1})$ with

$${}^nA^e_{q-1} = (A^e_q - T) \cup (A^{e-1}_q \cap T) \text{ for } 1 \le e \le E, \text{ and}$$

$${}^nA^0_{q-1} = A^0_q - T.$$

It is straightforward to verify that $w(q-1, {}^yA_{q-1}) + w(q-1, {}^nA_{q-1}) = w(q, A_q)$, using the identity $\left(\binom{n+1}{m}\right) = \left(\binom{n}{m}\right) + \left(\binom{n}{m-1}\right)$. The adversary answering strategy consists in answering "Yes" if $w(q-1, {}^yA_{q-1}) \ge w(q-1, {}^nA_{q-1})$ and "No" otherwise, thereby making sure that $w(q-1, A_{q-1}) \ge w(q, A_q)/2$.

The definitions of w and A_q imply that $w(Q,A_Q)=\left(\binom{Q}{E}\right)$ and $w(0,A_0) = \sum_{e=0}^{E} |A^e_0|$. Hence no $\epsilon = \sum_{e=0}^{E} |A^e_0|$ smaller than $\left(\binom{Q}{E}\right) \cdot 2^{-Q}$ is achievable against the adversary answering described above.

The above analysis also shows that the best questioning strategy is to choose the next question "Is $x \in T$?" such that the two possible weights are equal. Any such strategy does in fact achieve $\epsilon = \sum_{e=0}^{E} |A^e_0|$ smaller than $\left(\binom{Q}{E}\right) \cdot 2^{-Q}$ in the worst case. This can be done with a comparison "Is $x \le c$?" because $w(q-1, {}^yA_{q-1})$ and $w(q-1, {}^nA_{q-1})$ are continuous functions of c, for c=0 we have $w(q-1, {}^yA_{q-1}) \le w(q-1, {}^nA_{q-1})$, and for c=1 we have $w(q-1, {}^yA_{q-1}) \ge w(q-1, {}^nA_{q-1})$. Hence there is at least one c where the two weights are equal. □

We can of course rephrase the Theorem into a statement about the number $Q=Q(\epsilon,E)$ of questions necessary to achieve ϵ against up to E erroneous answers:

$$Q(\epsilon,E) = \min\{Q' \mid \epsilon \ge \left(\binom{Q'}{E}\right) \cdot 2^{-Q'}\}.$$

It is worth pointing out that this defines Q not only for E being a constant but also for E being a function of ϵ.

III. THE DISCRETE CASE

Any strategy for the continuous problem discussed in the preceding section can immediately be used for the discrete problem (by setting $\epsilon=1/n$) if the final set $\bigcup_{e=0}^{E} A^e_0$ known to contain the unknown x, is an interval. Unfortunately, this is not always the case with the optimal comparison strategy, as can be verified with an example as simple as E=1 and Q=4. Since the optimal comparison strategy for the continuous problem is unique, we cannot hope to duplicate its performance in the discrete case. However, we will not do much worse either.

Lemma 1: Consider the optimal comparison strategy for the continuous problem as described in the proof of Theorem 1. Then E additional comparison questions asked after the end of the strategy suffice to reduce the set $\bigcup_{e=0}^{E} A^e_0$ to a single interval.

Proof: The proof is easy and we only give a brief outline.

The set $\bigcup_{e=0}^{E} A^e_0$ consists of a finite number of intervals, say $I_1,\ldots,I_r$ (where $r \le E+1$). Define depth (I_j) to be the number of sets A^e_0, $0 \le e \le E$, that intersect I_j, and define $\mathrm{depth}(A_0) = \sum_{j=1}^{r} \mathrm{depth}(I_j)$. Using the fact that $A_0=(A^0_0,\ldots,A^E_0)$ can described completely by no more than E+1 previous "Yes"-answers and E+1 previous "No"-answers, it is not hard to see that $\mathrm{depth}(A_0) \le E+1$. The proof can then be completed by observing that any comparison questions placed between two intervals reduces $\mathrm{depth}(A_0)$ by at least 1, no matter what the answer is. □

Theorem 2: For any nonnegative integer E and positive integer n, let Q(n,E) denote the number of comparison questions necessary in the worst case to identify an unknown $x \in \{1,2,\ldots,n\}$ when up to E of the questions may receive an erroneous answer. Then

$$\min\{Q' \mid 2^{Q'} \ge n \cdot \left(\binom{Q'}{E}\right)\} \le Q(n,E) \le \min\{Q' \mid 2^{Q'-E} \ge n \cdot \left(\binom{Q'-E}{E}\right)\}.$$

Proof: Immediate from the remark after Theorem 1, Lemma 1, and from the fact that any strategy for the continuous problem that produces an interval as the final set known to contain x can be used for the discrete problem by simply setting $\epsilon=1/n$. □

The theorem holds when E is a constant as well as when E is a function of n. With some tedious manipulations the inequalities given in the Theorem can be shown to imply

$$Q(n,E)=\lg n+E \cdot \lg\lg n+O(E\lg E).$$

IV. FINDING MINIMUM ROOTS

The problem of determining a small interval containing the minimum root in the interval (0,1] of a set of continuous increasing functions g_i, $0\leq i\leq k$, is considered by Gal et.al. [3]. (We assume that $g_i(0)<0$ and $g_i(1)>0$ for all i.) The small interval is to be determined by testing the sign of the functions g_i evaluated at various points, i.e. by asking questions of the form "Is $g_i(c)>0$?" with $0\leq i\leq k$ and $c\in(0,1]$. All questions are to be answered correctly. First we show that this minimum root identification problem is equivalent to a variation of the search-with-errors problem considered in Section II. ("Equivalence" of the two problems is taken to mean that any optimal strategy for one problem translates trivially into an optimal strategy for the other problem, and conversely.) Then we show that its worst-case complexity is essentially the same as that of the search-with-errors problem in the previous section.

Theorem 3: The minimum root identification problem for a set $\{g_i|0\leq i\leq k\}$ of strictly increasing functions over (0,1] is equivalent to the problem of identifying an unknown $x\in(0,1]$ using only questions of the form "Is $x<c$?" with $c\in(0,1]$, when up to E=k of the "No"-answers but none of the "Yes"-answers may be erroneous.

Sketch of Proof: The state of the questioning in the minimum root identification problem can be summarized by k+2 real numbers

$$0\leq L_0\leq L_1\leq\ldots\leq L_k\leq R\leq 1$$

and a permutation π of {0,...,k} chosen such that L_i is the largest c such that the question "Is $g_{\pi(i)}(c)>0$?" has received a "No" answer (L_i is 0 if no such answer has been received yet) and R is the smallest c for which a question "Is $g_j(c)>o$?" has received a "Yes" answer, for some j (R is 1 if no such answer has been received yet).

The state of the questioning in the search-with-errors problem described in Theorem 1 can also be summarized by real numbers $L_0,\ldots,L_k,R$ ordered as above. Here the L_i's are the k+1 largest numbers c such that the question "Is $x<c$?" has received a "No" answer ($L_0=L_1=\ldots=L_{k'}=0$ if only k-k′ "No"-answers have been received yet) and R is the smallest c such that the question "Is $x<c$?" has received a "Yes" answer (R is 1 if no "Yes"-answer has been received yet).

Gal et.al. have shown (Proposition 2 in Section 5.1 of [3]) that optimality of any strategy for the minimum root identification problem is preserved if questions of the form "Is $g_{\pi(i)}(c)>0$?" are replaced by "Is $g_{\pi(0)}(c)>0$?". It can easily be verified that the two questions "Is $g_{\pi(0)}(c)>0$?" (for the root identification problem) and "Is $x<c$?" (for the search-with-errors problem) create the same successor configurations, viz.

$$0\leq L_0\leq L_1\leq\ldots\leq L_r\leq c\leq 1$$

if the answer is "Yes", and

$$0\leq L_1\leq L_2\leq\ldots\leq L_r\leq c\leq L_{r+1}\leq\ldots\leq L_k\leq R\leq 0$$

if the answer is "No". □

Because all the "Yes"-answers are guaranteed to be correct in this "half-lie" version, the problem of ending up with several intervals, which complicated the connection between the continuous and the discrete cases of the "full-lie" version in Sections II and III, does not exist here: The continuous and the discrete versions of the "half-lie" problem are trivially equivalent.

We now show that the worst-case complexity of the "half-lie" problem is essentially the same as that of the "full-lie" problem given in the previous Section.

Since an upper bound carries over immediately form the (seemingly harder) full-lie problem, the results from Section III imply an upper bound of $\lg n+E\cdot\lg\lg n+O(E\lg E)$. We now establish this same expression as a lower bound. The proof is somewhat unusual in that it does not describe an adversary answering strategy.

Theorem 4: Let Q(n,E) be the number of comparison questions "Is $x\leq c$?" necessary in the worst case to determine an unknown value $x\in\{1,\ldots,n\}$ when up to E of the "No"-answers, but none of the "Yes" answers, may be incorrect. Then $Q(n,E)\geq\lg n+E\cdot\lg\lg n+O(E\lg E)$.

Proof: Consider an arbitrary optimal (in the worst case) questioning strategy S, kept fixed throughout this proof. We model S by a finite binary "decision tree" T_S. Internal nodes t in T_S correspond to questions "Is $x\leq c_t$?", the left and right sons of an internal node are the questions asked following a "Yes" and "No" answer respectively (unless they are leaves), and leaves are associated with the determined value of the unknown x. We may assume that the strategy S always asks exactly Q questions, e.g. by asking "Is $x \leq n$?" k times if the value of the unknown is already determined after Q-k questions.

With each leaf ℓ of T_S we associate not only the determined value of the unknown, denoted by value(ℓ), but also two subsets of $\{1,...,Q\}$: lies(ℓ) and yes(ℓ). Here lies(ℓ) indicates which questions were answered incorrectly along the path of the leaf ℓ (incorrectly with respect to the correct value which is value(ℓ)), and yes(ℓ) indicates which questions received a "Yes" answer.

The main idea of this proof is this: If t is an internal node of T_S with question "Is $x \leq c_t$?" whose left son ("left" branches correspond to "Yes" answers) has a leaf ℓ with value(ℓ)=x where $x \leq c_t$ and $|\text{lies}(\ell)|<E$ among its successor, then the right son must also have at least one leaf ℓ' with value(ℓ')=x (and, in fact, $|\text{lies}(\ell')|=|\text{lies}(\ell)|+1$) among its descendants. This is just saying that if the "oracle" has not yet given E incorrect answers and the correct answer to the current question is "Yes", then the oracle may choose to answer (incorrectly) "No". If $|\text{lies}(\ell)|$ was less than E-1 then any later "Yes" answer along the path from the right son of t to the leaf ℓ' with value(ℓ')=x and $|\text{lies}(\ell')|=|\text{lies}(\ell)|+1$ (ℓ' is unique among the successors of the right son of t) in turn gives rise to another leaf ℓ'', with value x. This successive "creation" of leaves with value x then forces the number of leaves of T_S to be large enough to imply the desired lower bound on the depth of T_S. Since these "new" nodes are only caused by "Yes" answers we have to show that there are plenty of "Yes" branches in the tree.

Definition: A path from the root of T_S to a leaf ℓ is called regular if for all $i \in \{0,...,E-1\}$ at least 1/4 of the Q/E answers between (and including) the $((Q/E)\cdot i)+1^{st}$ and the $(Q/E)\cdot(i+1)^{th}$ answers are "Yes" answers. A path that is not regular is irregular.

Claim 1: For all large enough n, the number of irregular paths in T_S is less than n/2.

Proof: The number of "Yes"-"No" sequences of length Q/E with no more then (Q/E)/4 "Yes" answers is bounded by $2^{(Q/E)\cdot H(1/4)}$ where H is the "Entropy function" defined by $H(p)=-p\cdot \lg p-(1-p)\cdot \lg(1-p)$. Hence the number of "Yes"-"No" sequences of length Q with less than 1/4 of the answers between the $((Q/E)\cdot i)^{st}+1$ and the $(Q/E)\cdot(i+1)^{th}$ answer being "yes" answers is bounded by $2^{(Q/E)\cdot H(1/4)+(Q-(Q/E))}$ and consequently the number of regular paths in T_S is bounded by $E\cdot 2^{(Q/E)\cdot H(1/4)+(Q-(Q/E))}$. (For simplicity we are assuming that Q is divisible by E.) Let $d=1-(1/E)+(1/E)\cdot H(1/4)$. Then $d=1-(.18.../E)<1$ and $E\cdot 2^{(Q/E)\cdot H(1/4)+(Q-(Q/E))}=2^{Q\cdot d+\lg E}$. Since $Q \leq \lg n+E\cdot \lg\lg n+O(E\lg E)$ and $d<1$, we have for all large enough n: $Q\cdot d+\lg E \leq (\lg n+E\cdot \lg\lg n+O(E\lg E))\cdot d+\lg E<(\lg n)-1$, which implies the Claim.

Claim 2: For any $x \in \{1,...,n\}$, if all paths from the root of T_S to leaves with value x are regular then there are at least $(Q/4E)^E$ leaves with value x.

Proof: For each $x \in \{1,...,n\}$ and $i \in \{0,...,E\}$ define

$N_x(i)$ = {t|t is a node in T_S on level $((Q/E)\cdot i)+1$

t has among its descendants a leaf ℓ with value(ℓ)=x and $|\text{lies}(\ell)|=i$ but no leaf ℓ' with value(ℓ')=x and $|\text{lies}(\ell')|<i$}

For any x, $N_x(0)$ consists of the root of T_S, hence $|N_x(0)|=1$. For all i<E, $|N_x(i+1)| \geq 1/4\cdot(Q/E)\cdot|N_x(i)|$ because each "Yes" branch on the path from a node $t \in N_x(i)$ to a leaf ℓ with $|\text{lies}(\ell)|=i$ and value(ℓ)=0 gives rise to a node in $N_x(i+1)$, and by assumption there are at least $1/4\cdot(Q)$ "Yes" branches between levels (Q/E)i+1 and $(Q/E)\cdot(i+1)$. This is explained above in the description of the main idea of this proof and is illustrated in Figure 1. By induction, $|N_x(Q)| \geq (Q/4E)^E\cdot|N_x(0)|=(Q/4E)^E$, which proves the Claim.

Since $N_x(Q) \cap N_{x'}(Q)=\emptyset$ for $x \neq x'$, at least n/2 values x have only regular paths to leaves with value=x, and since $Q \geq \lg n$, we get

number of leaves in $T_S \geq (n/2)\cdot(Q/4E)^E \geq (n/2)\cdot(\lg n/4E)^E$

which implies that the depth of T_S is at least $\lg n+E\cdot \lg\lg n+O(E\lg E)$.

This lower bound result also holds for strategies with arbitrary "Yes"-"No" questions since the above proof does not make use of the restriction to comparisons.

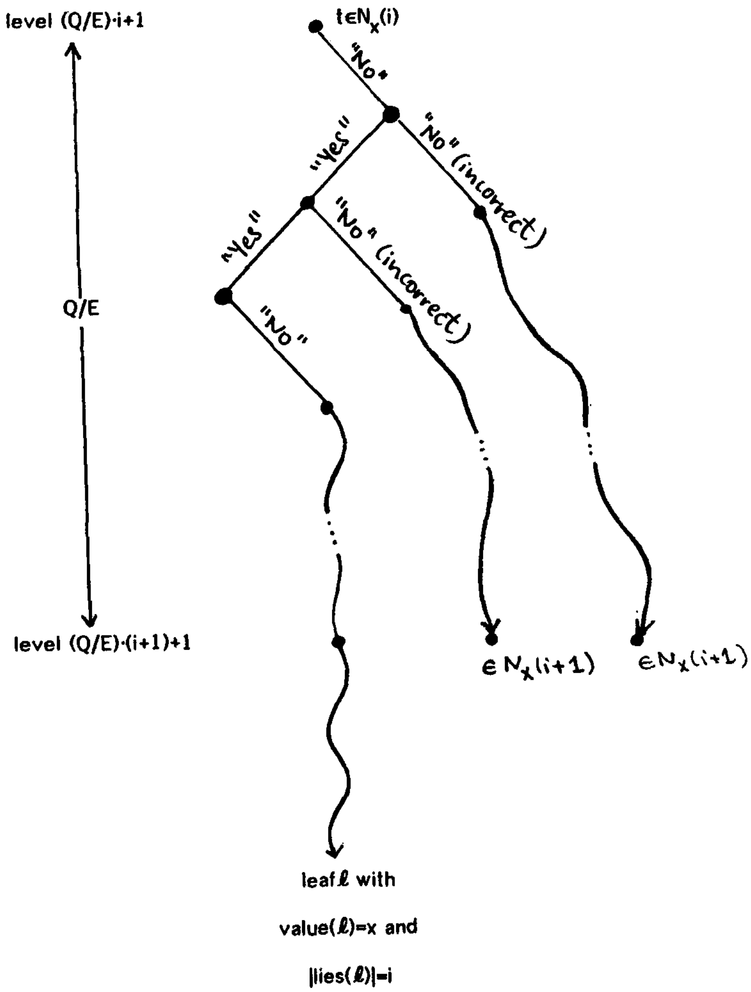

Figure 1.

V. OPEN QUESTIONS

We conjecture that the gap of E questions between the continuous and discrete problems discussed in Theorems 1 and 2 can be decreased.

Clearly the issue of coping with erroneous answers could also be studied in the context of other familiar search problems such as sorting, finding medians, finding maxima of unimodal functions, etc. (cf.[4]).

VI. REFERENCES

[1] Berlekamp, E.R., "Block Coding for the Binary Symmetric Channel with Noiseless, Delayless Feedback," in Error-Correcting Codes (Wiley, 1968), pp. 61-85.

[2] Brylawski, T.H., "The Mathematics of Watergate," (Unpublished paper)

[3] Gal, S., B. Bacherlis, and A. Ben-Tal, "On Finding the Maximum Range of Validity of a Constrained System," SIAM J. Control and Opt., to appear.

[4] Katona, G.O.H., "Combinatorial Search Problems," in A Survey of Combinatorial Theory (Chapter 23), (North-Holland, 1973).

[5] Rivest,R.L., A.R. Meyer, D.J. Kleitman, and J. Spencer, "Binary Search Using Unreliable Comparisons," in Proc. 15th Annual Allerton Conf. on Communication, Control, and Computing, Sept. 28-30,1977.

[6] Schalkwijk, J.P.M., "A Class of Simple and Optimal Strategies for Block Coding on the Binary Symmetric Channel with Noiseless Feedback," in IEEE Trans. on Information Theory, 17, 3(1971), pp.283-287.

[7] Ulam, S.M., Adventures of a Mathematician (Charles Scribner's Sons, New York, 1976), p. 281.

ACKNOWLEDGEMENTS

We wish to thank Shimon Even, Michael J. Fischer, and Gene Lawler for helpful conversations.

We also wish to cite Ulam's memoirs [7] which originally stimulated our interest in this problem