AN EFFICIENT ALGORITHM FOR COMPUTING AN OPTIMAL (*r*, *Q*) POLICY IN CONTINUOUS REVIEW STOCHASTIC INVENTORY SYSTEMS

AWI FEDERGRUEN
Columbia University, New York, New York

YU-SHENG ZHENG
University of Pennsylvania, Philadelphia, Pennsylvania

(Received May 1990; revisions received November 1990, April 1991; accepted May 1991)

The reorder point/reorder quantity policies, also referred to as (r, Q) policies, are widely used in industry and extensively studied in the literature. However, for a period of almost 30 years there has been no efficient algorithm for computing optimal control parameters for such policies. In this paper, we present a surprisingly simple and efficient algorithm for the determination of an optimal (r^*, Q^*) policy. The computational complexity of the algorithm is linear in Q^*. For the most prevalent case of linear holding, backlogging and stockout penalty costs (in addition to fixed order costs), the algorithm requires at most $(6r^* + 13Q^*)$ elementary operations (additions, comparisons and multiplications), and hence, no more than 13 times the amount of work required to do a single evaluation of the long-run average cost function in the point (r^*, Q^*).

In this paper, we derive a simple and efficient algorithm for computing optimal (r, Q) inventory control policies. These policies are also known as reorder point/order quantity policies. We will restrict ourselves to the case when demands arise on a unit-by-unit basis. Under an (r, Q) policy, the inventory position (= inventory on-hand + orders outstanding-backorders) of the item in question is *continuously reviewed*, and an order of fixed quantity Q is placed as soon as the inventory position drops to a reorder point r.

Such (r, Q) policies are widely used in inventory systems with uncertain demands and lead times. For single item inventory systems under standard assumptions, it is well known that an optimal policy exists within the class of (r, Q) policies. Many multi-item or multilocation systems are designed such that each item's (facility's) inventory is governed by an (r, Q) policy. Other planning models consist of a large number of single item (r, Q) systems, tied together by aggregate inventory constraints. These models are decomposed into single item models via Lagrangian relaxation. Highly efficient solution methods are essential here. Similarly, Atkins and Iyogun (1988) propose a decomposition method to derive a tight lower bound for stochastic joint replenishment models in which optimal (r, Q) policies need to be computed repeatedly for each of the items involved; (r, Q) policies are also optimal in many (generalized) stochastic clearing systems with point arrival processes that arise in other settings than those involving physical inventories. See Federgruen and Zheng (1988) for details.

The use of (r, Q) policies has been propagated since the seminal paper of Galliher, Morse and Simmond (1959), and the classical textbook by Hadley and Whitin (1963) appeared 30 years ago. Nevertheless, and as mentioned in Browne and Zipkin (1991), "until recently, there was no reliable, straightforward method for computing an optimal (r, Q) policy, even in the simple case of Poisson demand processes." Instead, a large number of heuristics have been proposed (see Lee and Nahmias 1989). The only existing algorithm, to our knowledge, was presented in Zipkin's (1988) classnotes. This procedure is based on a result in Sahin (1982); see also Sahin (1990).

Our algorithm is based on the observation that the long-run average cost $C(r, Q)$ of an (r, Q) policy is of the form:

$$C(r, Q) = \left[\kappa + \sum_{y=r+1}^{r+Q} G(y)\right] \Big/ Q. \tag{1}$$

Here $\kappa > 0$ is a given constant and $-G(\cdot)$ is a unimodal function with $\lim_{|y|=\infty} G(y) = \infty$. Our approach is

Subject classification: Inventory/production: stochastic policies.
Area of review: Stochastic Processes and Their Applications.

based on the following observations: The unimodality of $-G(\cdot)$ implies: for fixed Q, $C^*(Q) \overset{\text{def}}{=} \min_r C(r, Q)$ is achieved when the sum in (1) consists of the Q smallest values of this function; and these values are achieved in Q contiguous points and the optimal corresponding reorder level r is trivially identified. Next it is easy to verify that $-C^*(\cdot)$ is unimodal as shown in Sahin (1982), and Q^*, the optimal order size, is obtained as the largest value of Q for which $C^*(Q - 1) > G_Q$ with G_Q the Qth smallest $G(\cdot)$ value, so that

$$C^*(Q + 1) = [QC^*(Q) + G_{Q+1}]/(Q + 1) \geq C^*(Q),$$

and

$$C^*(Q) = [(Q - 1)C^*(Q - 1) + G_Q]/Q < C^*(Q - 1)$$

(see Figure 1). These observations may be exploited in an efficient algorithm whose complexity is linear in Q^*. In the case of linear holding and backlogging costs the computational complexity of the algorithm is no larger than $(6r^* + 13Q^*)$ elementary operations (additions, comparisons and multiplications) when $r^* > 0$, and hence, no more than 13 times the amount of work required to do a single evaluation of the $C(\cdot, \cdot)$ function in the point (r^*, Q^*). (Similar complexity counts apply when $r^* \leq 0$ or when a more general one-step expected cost function is used.)

The (r, Q) policies are a special case of (s, S) policies, under which the item's inventory position is ordered to S whenever it is observed to have fallen *to* or *below* the level s $(s < S)$. This more general structure arises when demands occur in batches of random size. In a related paper (Zheng and Federgruen 1991) we develop an efficient algorithm for finding optimal (s, S) policies. Since the cost function of an (s, S) policy fails, in general, to be quasiconvex, except under a restrictive assumption on the demand size distribution (see Stidham 1977, and Sahin 1982), that algorithm has to use a different and more complex search procedure. Its complexity is at least *quadratic* in $(r^* + Q^*)$, even when applied to models in which (r, Q) policies are optimal, i.e., where the cost function is given by (1).

Another commonly used generalization of the (r, Q) policy, to accommodate for random demand sizes, is the (r, nQ) policy: Here as soon as the inventory position drops to or below r, an (integer) *multiple* of a fixed quantity Q is ordered to raise the inventory position back to the interval of $[r + 1, \ldots, r + Q]$. As for the simpler (r, Q) policies, the steady-state distribution of the inventory position is uniformly distributed under standard assumptions (see Hadley and Whitin, and Richard 1975). The cost function of an (r, nQ) policy is therefore of a form similar to (1). An extension of the algorithm in this note may thus be employed (see Zheng and Chen 1990 for details).

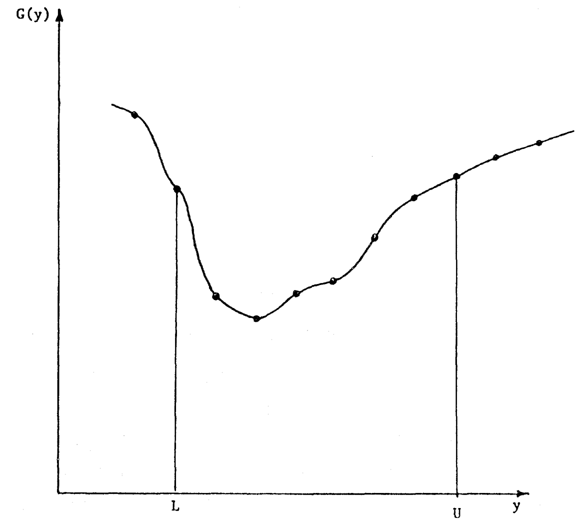

Figure 1. $G(\cdot)$ function. If $G(L), G(L + 1), \ldots, G(U)$ represent the q smallest values of the $G(\cdot)$ function, then the $q + 1$st smallest value is found for $y = L - 1$ or $U + 1$.

In Section 1 we introduce the notation and preliminaries. The proposed algorithm is derived and discussed in Section 2.

1. NOTATION AND PRELIMINARIES

Consider a single item whose inventory may be replenished by placing orders of unlimited size. Orders arrive after a given lead time. Stockouts are backlogged. In this section, we briefly review the main inventory models in which (r, Q) policies are optimal and their average cost is of the form given by (1) because these results are scattered throughout or are not available in the open literature.

For any $t \geq 0$, let

$D(t)$ = the total demand in $[0, t)$;
$IP(t)$ = the inventory position at time t;
$IL(t)$ = the inventory level at time t.

Consider first the simplest of all models for which optimal (r, Q) policies exist, namely the case of Poisson demands and constant lead times. It is well known (see, e.g., Hadley and Whitin, and Zipkin 1986a) that the inventory position process $IP(t)$ and the inventory level process $IL(t)$ have limiting distributions. Indeed, with $IP(\infty)$ and $IL(\infty)$ denoting random variables with these distributions we have

$$IL(\infty) = IP(\infty) - LD(\infty), \tag{2}$$

where

$IP(\infty)$ is uniformly distributed on

$$\{r+1, r+2, \ldots, r+Q\}, \tag{3}$$

and

$LD(\infty)$ is a nonnegative integer valued random variable which is *independent* of $LP(\infty)$. (4)

In fact, $LD(\infty)$ represents the total demand in a lead time, which in the above model has a Poisson distribution with a mean λL, where λ denotes the per unit demand rate. Now assume that the cost structure consists of a fixed cost K per order, an inventory carrying cost of h per unit carried in stock per unit of time, and a backlogging cost of p per unit of backlogged demand, per unit of time. In view of (2)–(4) we obtain, as in Hadley and Whitin,

$$\begin{aligned} C(r, Q) &= \frac{K\lambda}{Q} + h \sum_{j=0}^{\infty} \text{Prob}[IL(\infty) = j] \\ &\quad + p \sum_{j=-\infty}^{-1} (-j)\text{Prob}[IL(\infty) = j] \\ &= \frac{K\lambda}{Q} + \frac{1}{Q} \\ &\quad \cdot \sum_{y=r+1}^{r+Q} \left\{ h \sum_{i=0}^{y} (y-i)p_i + p \sum_{i=y+1}^{\infty} (i-y)p_i \right\}, \end{aligned}$$

where $p_j = \text{Prob}[LD(\infty) = j]$. (In the Poisson demand model, we have $p_j = e^{\lambda L}(\lambda L)^j/j!$, $j = 0, 1, \ldots$.) Thus, $C(r, Q)$ may be written in the form of (1) with $\kappa = K\lambda$, and

$$\begin{aligned} G(y) &= h \sum_{j=0}^{y} (y-j)p_j + p \sum_{j=y+1}^{\infty} (j-y)p_j \\ &= (h+p) \sum_{j=0}^{y-1} P_j + p(\lambda L - y), \end{aligned} \tag{5}$$

where $P_j = \text{Prob}[LD(\infty) \leq j]$. (In our basic model, (1) may also be derived using the renewal reward theorem; the chosen derivation allows for the extensions discussed below.)

Several more general cost structures result in $C(r, Q)$ being of the form of (1). For example, the inventory carrying cost and/or the backlogging cost per unit of time may be specified as a general convex function $h(\cdot)$ or $p(\cdot)$ of the inventory or backlog sizes. In some systems, a one-time penalty π is incurred for each stockout, i.e., for each unit which cannot be delivered upon demand. (The parameters p and π may either be given explicitly or they may arise in Lagrangian relaxations of service level constraints, i.e., upper bounds with respect to the average backlog or the *fill rate*, the fraction of sales which can be satisfied upon demand.) The stockout penalty cost results in an additional term in the average cost expression:

$$\begin{aligned} \bar{C}(r, Q) &= C(r, Q) + \lambda\pi \,\text{Prob}[IL(\infty) < 0] \\ &= C(r, Q) + \frac{\lambda\pi}{Q} \sum_{y=r+1}^{r+Q} \sum_{j=y}^{\infty} p_j, \end{aligned}$$

where the first equality follows from the Poisson Arrivals See Time Averages (PASTA) property. Once again it is easy to verify that (1) continues to apply with $G(\cdot)$ replaced by

$$\bar{G}(y) = G(y) + \lambda\pi(1 - P_{y-1}).$$

It is also easy to verify that for all of the above cost structures $-G(\cdot)$ is unimodal. Our algorithm, described in the next section, merely uses the fact that $C(r, Q)$ is of the form of (1) as well as the unimodality of $-G(\cdot)$.

Characterization of the distribution of $IL(\infty)$ via (2)–(4) applies to considerably more general continuous-time models than the basic model with Poisson demands and fixed lead times. We refer to Zipkin (1988) for an elegant and comprehensive treatment of these generalized models. See also Federgruen and Zheng (1988).

The Poisson demand assumption **D1** may, for example, be relaxed to **D2**: $\{D(t)\}$ is a renewal process; or **D3**: The demand process is a counting process with a Markovian representation, i.e., there exists an auxiliary process such that the joint process $\{A(t), D(t)\}$ is Markovian. Moreover, $(A(t), IP(t)) \xrightarrow{D} (A(\infty), IP(\infty))$ as $t \to \infty$ with $A(\infty)$ and $IP(\infty)$ independent.

Conditions (2)–(4) continue to hold when $\{D(t)\}$ is a general, nondecreasing process with stationary increments and continuous sample paths (see Zipkin 1986b, remark after Proposition 1, Zheng 1989, and Browne and Zipkin 1991). Here r and Q are to be treated as continuous variables and (1) continues to hold after replacing the summation by an integral.

The characterization of the limiting inventory level distribution by (2)–(4) is valid under fixed lead times as well as certain types of *random* lead times (see Zipkin 1986a, 1988).

In all of the above models, $LD(\infty)$ may be interpreted as the "lead time demand," appropriately derived. For example, under **D2** and fixed lead times

L, $LD(\infty)$ represents the number of demands (renewals) that occur in the time interval $[0, L]$ under the equilibrium renewal process. In some settings, a specific distribution for the lead time demand $LD(\infty)$ is directly estimated or surmised, rather than being derived from an underlying demand and lead time process.

2. THE OPTIMIZATION ALGORITHM

Consider an inventory model in which the long-run average cost of an (r, Q) policy is of the form of (1) with $-G(\cdot)$ a unimodal function. In this section, we describe our optimization procedure.

As we point out in the Introduction, the unimodality of $-G(\cdot)$ implies that for fixed Q, the Q smallest $G(\cdot)$ values can be achieved in contiguous points and $C^*(Q) = \min_r C(r, Q)$ is achieved if the sum in (1) consists of these values. These values and, hence, the optimal reorder levels $r^*(1), \ldots, r^*(Q)$ (for given order quantities $1, \ldots, Q$) are easy to identify by the following procedure. Let y_1 be an integer that minimizes $G(y)$ over all integers. We generate a sequence $\{y_1, y_2, \ldots,\}$ inductively. Assuming that $\{y_1, \ldots, y_Q\}$ have been generated, let $L(Q) = \min\{y_1, \ldots, y_Q\}$, $R(Q) = \max\{y_1, \ldots, y_Q\}$. Then let

$$y_{Q+1} = \begin{cases} L(Q) - 1 & \text{if } G(L(Q) - 1) \leqslant G(R(Q) + 1) \\ R(Q) + 1 & \text{otherwise.} \end{cases}$$

Clearly, for any given Q, $\{y_1, \ldots, y_Q\}$ are contiguous, and their $G(\cdot)$ values constitute the Q smallest ones due to the unimodality of $-G(\cdot)$. This leads directly to the following lemma.

Lemma 1. *For any given integer* $Q \geqslant 1$, $r^*(Q) = L(Q) - 1$.

Proof. The proof is by (1).

Clearly, $L(Q) = L(Q - 1)$ or $L(Q) = L(Q - 1) - 1$. This implies the following corollary.

Corollary 1. $r^*(Q) - 1 \leqslant r^*(Q + 1) \leqslant r^*(Q)$ *for all integers* $Q \geqslant 1$.

This corollary may be derived (with considerably more effort) from the results in Sahin (1982) (see Zipkin 1988).

We conclude that

$$C^*(Q) = \left[\kappa + \sum_{i=1}^{Q} G(y_i)\right] \Big/ Q,$$

so that

$$C^*(Q + 1) = [QC^*(Q) + G(y_{Q+1})]/(Q + 1). \qquad (6)$$

Note that $C^*(Q + 1) < C^*(Q)$ if and only if $G(y_{Q+1}) < C^*(Q)$. This suggests the following exceedingly simple characterization of Q^*, the optimal order size.

Lemma 2. Q^* *is the smallest integer* q *with the property* $C^*(q) \leqslant G(y_{q+1})$.

Proof. It follows from (6) that $C^*(Q)$ is decreasing for $Q < Q^*$. Moreover, for $Q > Q^*$,

$$\begin{aligned} &C(Q) - C(Q^*) \\ &= \left[\sum_{i=Q^*+1}^{Q} G(y_i) - (Q - Q^*)C^*(Q)\right] \Big/ Q \\ &\geqslant (Q - Q^*)[G(y_{Q^*+1}) - C^*(Q^*)]/Q \geqslant 0. \end{aligned}$$

Lemmas 1 and 2 clearly suggest an efficient algorithm for finding an optimal reorder level and order size (r^*, Q^*). Below we give a detailed algorithm for the case where $\Delta G(y) = G(y + 1) - G(y)$ is easy to compute (this is, for example, the case in the Poisson demand model, see (5)). For notational convenience only, we restrict ourselves to the most common case, where $y_1 > 0$.

Algorithm OPT

```
Step 0. Calculate G(0) and ΔG(0); L := 0;
while ΔG(L) < 0 do
  begin L := L + 1, evaluate ΔG(L), G(L + 1) :=
        G(L) + ΔG(L)
  end;
S := κ + G(L), Q := 1, C* := S, r := L − 1,
R := L + 1;

Step 1. Repeat
begin if G(r) ≤ G(R)
  then if C* ≤ G(r)
    then stop.
    else begin S := S + G(r), r := r − 1,
      if r < 0, evaluate ΔG(r) and G(r) :=
      G(r + 1) − ΔG(r),
    end;
else if C* ≤ G(R)
  then stop.
  else begin S := S + G(R), evaluate ΔG(R),
          G(R + 1) := G(R) + ΔG(R), R := R + 1
  end;
  Q := Q + 1, C* := S/Q
end.
```

Complexity of the Algorithm

Assume that the probability density function of $LD(\infty)$, i.e., the numbers $\{p_j : j = 0, 1, \ldots\}$, are given as input to the problem. We first analyze the complexity of the algorithm for the (most prevalent) case where the cost structure is *linear*, i.e., where it consists of linear holding, backlogging and stockout penalty costs (in addition to the fixed order costs). We first analyze the case where $r^* > 0$. The entire algorithm consists of:

i. evaluation of the $r^* + Q^*$ first values of the difference function $\Delta G(y) = (h + p)P_y - \lambda\pi p_y - p$: each such evaluation requires five elementary operations (three additions and two multiplications; in the absence of stockout penalty costs only three operations are required);
ii. $r^* + Q^*$ additions to compute the $\{G(1), \ldots, G(r^* + Q^*)\}$ values;
iii. at most, $7Q^*$ elementary operations ($3Q^*$ comparisons; $3Q^*$ additions and Q^* divisions) to execute the remaining work in Step 1.

The algorithm thus requires no more than $6r^* + 13Q^*$ elementary operations ($4r^* + 11Q^*$ operations in the absence of stockout penalty costs). In case $r^* \leqslant 0$, the total number of operations is reduced to $13Q^*$ because only Q^* values of the $G(\cdot)$ function are computed. It is useful to compare this complexity bound with a lower bound for the amount of work that is required to do a single evaluation of the cost function $C(\cdot, \cdot)$ at the optimal point (r^*, Q^*). Even when the cost structure is linear this evaluation requires the computation of $P_{r^*+Q^*}$ and the latter requires $r^* + Q^*$ additions. We conclude as follows.

Theorem 1. *Algorithm* **OPT** *determines the optimal* (r^*, Q^*) *policy. Under linear cost structures and when* $r > 0$, *the amount of work* (*measured as the number of elementary operations required*) *is at worst* 13 *times the amount of work required to evaluate* $P_{r^*+Q^*}$ (*and at worst* 11 *times in the absence of stockout penalty costs*).

We conclude this paper with a few comments on the case where the cost structure consists of general nonlinear holding, backlogging or stockout penalty cost functions (the *nonlinear* case). In the linear case, the function $G(\cdot)$ is most easy to evaluate via $\Delta G(\cdot)$ (see (5)). In the nonlinear case, it is more efficient to evaluate the $G(\cdot)$ values directly. The algorithm needs to be modified accordingly. We may replace Step 0 by a bisection search to locate the minimum point y_1. To gain efficiency, $\{G(r^*), \ldots, G(r^* + Q^*)\}$ should be evaluated in Step 1 when needed. The resulting algorithm requires (see Federgruen and Zheng (1988): at most $\lceil \log_2 U \rceil + Q^* + 1$ evaluations of the $G(\cdot)$ function with U an priori upper bound for y_1; and at most $8Q^*$ elementary operations to execute the remaining work. Note that a single evaluation of the average cost function $C(\cdot, \cdot)$ requires at least Q^* evaluations of the function $G\cdot$).

Note Added in Proof

Shortly prior to receiving the galley proofs of this paper, we became aware of G. Rubal'sky's "Calculations of Optimum Parameters in an Inventory Control Problem" (*Eng. Cybern.* **10,** 182–187; translated from Russian), which describes a similar method for the determination of optimal r- and q-values.

REFERENCES

ATKINS, D., AND P. IYOGUN 1988. Priodic Versus Can-Order Policies for Coordinated Multi-item Inventory Systems. *Mgmt. Sci.* **34,** 791–795.

BROWNE, S., AND P. ZIPKIN. 1991. Inventory Models With Continuous, Stochastic Demands. *Anns. Appl. Prob.* **1** (3), 419–435.

FEDERGRUEN, A., AND Y. S. ZHENG. 1988. A Simple and Efficient Algorithm for Computing Optimal (*r, Q*) Policies in Continuous-Review Stochastic Inventory Systems. Working Paper (Unabridged Version). Decision Sciences Department, The Wharton School, University of Pennsylvania, Philadelphia.

GALLIHER, H., P. MORSE AND M. SIMMOND. 1959. Dynamics of Two Classes of Continuous-Review Inventory Systems. *Opns. Res.* **7,** 362–384.

HADLEY, G., AND T. M. WHITIN. 1963. *Analysis of Inventory Systems.* Prentice-Hall, Englewood Cliffs, N.J.

LEE, H., AND S. NAHMIAS. 1989. Single Product, Single-Location, Models, Chap. 2. In *Handbook in Operations Research and Management Science,* Vol. 4: *Logistics of Production and Inventory,* S. Graves, A. Rinnooy Kan and P. Zipkin (eds). North Holland, Amsterdam.

RICHARD, F. 1975. Comments of the Distribution of Inventory Position in a Continuous Review (*s, S*) Inventory System. *Opns. Res.* **23,** 366–371.

SAHIN, I. 1982. On the Objective Function Behavior in (*s, S*) Inventory Models. *Opns. Res.* **30,** 709–725.

SAHIN, I. 1990. *Regenerative Inventory Systems,* Springer-Verlag, New York.

STIDHAM, S. 1977. Cost Models for Stochastic Clearing Systems. *Opns. Res.* **25,** 100–127.

ZHENG, Y. S. 1989. Properties of Stochastic Inventory Systems. *Mgmt. Sci.* (to appear).

ZHENG, Y. S., AND F. CHEN. 1990. Inventory Policies With Quantized Ordering. Working Paper, Decision Sciences Department. The Wharton School, University of Pennsylvania, Philadelphia.

ZHENG, Y. S., AND A. FEDERGRUEN. 1991. Finding Optimal (*s*, *S*) Policies is About as Simple as Evaluating a Single Policy. *Opns. Res.* **39,** 654–665.

ZIPKIN, P. 1986a. Stochastic Leadtimes in Continuous-Time Inventory Models. *Naval Res. Logist. Quart.* **33,** 763–774.

ZIPKIN, P. 1986b. Inventory Service-Level Measures: Convexity and Approximation. *Mgmt. Sci.* **32,** 975–981.

ZIPKIN, P. 1988. Lecture Notes in Inventory Theory. Graduate School of Business, Columbia University, New York.

ON CONSTRAINED BOTTLENECK EXTREMA

JON LEE

Yale University, New Haven, Connecticut

(Received September 1990; revision received April 1991, accepted June 1991)

This note treats optimization problems that involve two criteria on set systems. One is a bottleneck criterion, and the other is general. In particular, recent algorithms for such problems on the edge-sets of networks are related to the fundamental work of J. Edmonds and D. R. Fulkerson involving more general set systems.

In a recent note by Berman, Einav and Handler (1990), algorithms are given for graph problems that involve two criteria with respect to the edge-set. The first problem is that of minimizing a bottleneck objective, subject to a single generic constraint. The second problem is that of minimizing a generic objective function, subject to a single bottleneck constraint. The purpose of this note is to point out that the first problem can be studied in the broader setting of clutters. The analysis of the second problem involves even less structure.

Let $\mathscr{C}$ be an arbitrary set of subsets of the finite set $\mathscr{E}$. Let w and f be arbitrary real-valued functions on $\mathscr{E}$ and $\mathscr{C}$, respectively. Let F and W be real numbers. We define two problems on $\mathscr{C}$.

Problem $\mathscr{P}_1$

$$\min_{\mathscr{X}\in\mathscr{C}}\{\max_{e\in\mathscr{X}} w(e): f(\mathscr{X}) \leq F\}.$$

Problem $\mathscr{P}_2$

$$\min_{\mathscr{X}\in\mathscr{C}}\{f(\mathscr{X}): \max_{e\in\mathscr{X}} w(e) \leq W\}.$$

For the purposes of solving $\mathscr{P}_1$, if f is monotone nondecreasing, we may as well assume that $\mathscr{C}$ is a clutter (i.e., the elements of $\mathscr{C}$ are pairwise incomparable with respect to set inclusion), because we can replace $\mathscr{C}$ by the clutter of its minimal (nonempty) elements without altering the optimal value of $\mathscr{P}_1$. The set $\mathscr{C}_F \equiv \{\mathscr{X} \in \mathscr{C}: f(\mathscr{X}) \leq F\}$ is also a clutter on E. Hence, the following "threshold algorithm" of Edmonds and Fulkerson (1970) can be used to solve $\mathscr{P}_1$. Accumulate elements of $\mathscr{E}$, in order of nondecreasing weight, collecting them in a set $\mathscr{S}$, until $\mathscr{S}$ contains an element of $\mathscr{C}_F$. We can determine whether $\mathscr{S}$ contains an element of $\mathscr{C}_F$ by solving the following problem.

Problem $\mathscr{P}_{\mathscr{S}}$

$$\min_{\mathscr{X}\in\mathscr{C}}\{f(\mathscr{X}): \mathscr{X} \subseteq \mathscr{S}\}.$$

If the procedure is implemented using bisection search on the elements of $\mathscr{E}$, ordered by w, an $\mathscr{O}(t(\mathscr{C})\log|\mathscr{E}|)$ algorithm results, where $t(\mathscr{C})$ is the maximum time to solve $\mathscr{P}_{\mathscr{S}}$ (over $\mathscr{S} \subseteq \mathscr{E}$).

If f is not monotone, we may still treat the case in which $\mathscr{C}$ is not a clutter, by replacing $\mathscr{C}_F$ by the clutter of its minimal (nonempty) elements.

Since $\mathscr{C}_F$ is a clutter, the following duality equation holds (Edmonds and Fulkerson):

$$\min_{\mathscr{X}\in\mathscr{C}_F}\max_{e\in\mathscr{X}} w(e) = \max_{y\in\mathscr{C}_F^*}\min_{e\in y} w(e),$$

where $\mathscr{C}_F^*$ is the "blocker" of $\mathscr{C}_F$. That is, $\mathscr{C}_F^*$ is the set of minimal subsets of $\mathscr{E}$ that have nonempty intersection with every element of $\mathscr{C}_F$. In general, it is not easy

Subject classifications: Mathematics, combinatorics: clutters. Networks/graphs: algorithms. Programming: bottleneck objective with a side-constraint. *Area of review:* OPTIMIZATION.