Controlled Flooding Search with Delay Constraints

Nicholas B. Chang and Mingyan Liu
Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI 48109-2122
Email: {changn,mingyan}@eecs.umich.edu

Abstract

In this paper we consider the problem of query and search in a network, e.g., searching for a specific node or a piece of data. We limit our attention to the class of TTL (time-to-live) based controlled flooding search strategies where query/search packets are broadcast and relayed in the network until a preset TTL value carried in the packet expires. Every unsuccessful search attempt results in an increased TTL value (i.e., larger search area) and the same process is repeated. Every search attempt also incurs a cost (in terms of packet transmissions and receptions) and a delay (time till timeout or till the target is found). The primary goal is to derive search strategies (i.e., sequences of TTL values) that minimize a worst-case cost measure subject to a worst-case delay constraint. We present a constrained optimization framework and derive a class of optimal strategies, shown to be randomized strategies, and obtain their performance as a function of the delay constraint. We also use these results to discuss the trade-off between search cost and delay within the context of flooding search.

Index Terms

data query and search, TTL, controlled flooding search, wireless sensor and ad hoc networks, constrained optimization, randomized strategy, competitive analysis

I. Introduction

Query and search form an important functionality for many network applications. Searching for a destination node whose location is unknown is a prime example frequently encountered by ad hoc network routing protocols and services, e.g., [1], [2], [3]. Other examples include the search for certain data of interest in an environmental monitoring sensor network [4], and more broadly, the search for a shared file in a peer-to-peer (P2P) network. A good search mechanism should have a short response time, i.e. the time it takes to find the object of interest, and should do so with minimal cost. Such cost refers to the amount of processing and transmissions incurred by the search. This is particularly important in a wireless context where cost is associated with energy consumption, measured by the amount of packet transmissions and receptions.

There are a variety of mechanisms one may use to conduct search. These include maintaining a centralized directory service, or by sending out a query packet that traverses the network in a certain way, e.g., the rumor routing proposed in [4] and the random walk method in [5].

In this paper we focus on the class of TTL-based controlled flooding search. This decentralized search mechanism is widely used, e.g., in ad hoc routing protocols [6] as well as in P2P networks [7]. Under this scheme the node originating the search (also referred to as the em source node) sends out a query packet that carries an integer TTL (time-to-live) value. If the underlying network is wired, this query will be transmitted once along each outgoing link of the source. In a wireless network, this packet reaches all immediate neighbors in a single broadcast transmission. If the search target is found at a neighboring node, it will reply to the source. Otherwise it decrements the TTL value by one and retransmits the query packet containing the new TTL value. This continues until TTL reaches zero. Thus how much of the network has been queried is controlled by the TTL value. If the target is not found in this search area, the source node will eventually time out and initiate another *round* of search covering a bigger area using a larger TTL value, and presumably setting a larger timeout value for that round of search. This process continues until either the object is found or the source gives up. Hence the performance of a search strategy both in terms of cost and delay is determined by the sequence of TTL values used.

Controlled flooding search has previously been studied in [7], [8], [6], [9]. In [8] it was shown that if the target location distribution is known, then a dynamic programming formulation can be used to find TTL sequences that minimize the average search cost. When this distribution is not known, [7] derived the optimal deterministic TTL sequence that minimizes a worst-case cost measure. For the same cost measure [8] showed that the best strategies are *randomized* TTL sequences. [9] further

This work is supported by NSF award ANI-0238035, ONR grant N00014-03-1-0232, and through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011.

derived the optimal randomized strategy. All these studies focused only on the cost of search and did not consider the delay of search.

The primary goal of this paper is to derive TTL controlled flooding search strategies that perform well both in terms of the search cost and the search delay. While there are a variety of ways to handle multiple (potentially conflicting) objectives, in this study we will approach this by formulating a constrained optimization problem. Specifically, we will attempt to minimize a cost measure subject to a delay constraint. The solution to this problem results in search strategies that minimize the search cost and locate the target within a specified time constraint.

Imposing a delay constraint also allows us to study the trade-off between search cost and search delay. This trade-off can be seen by considering the strategy of flooding the entire network (e.g., setting the TTL to be the maximum allowed value, assumed sufficient to cover the whole network). Such a strategy would most likely result in a short search delay, as the target node is likely to be found during the first round. On the other hand, this strategy is not the most cost effective [7], [8], as it results in a large amount of packet transmissions and retransmissions. Conversely, it may be the case that some strategy incurs extremely low cost but suffers from high delay. Under the constrained optimization framework studied in this paper, we will be able to conveniently address the above performance trade-off.

The main contributions of this paper are summarized as follows:

1) We provide an analytical framework within which the delay of search strategies can be studied along with their cost. The previously studied unconstrained problem (simply minimizing a cost measure) [7], [9] becomes a special case under this framework when the delay constraint is not binding. In this sense, this method is a key generalization of prior work and presents a much more powerful analytical tool. To the best of our knowledge, delay has not been studied in this context before.
2) When a worst-case delay constraint is imposed, we derive a class of optimal strategies that minimize worst-case cost measure among all strategies that satisfy the delay constraint. In [8] we showed that randomized strategies outperform deterministic strategies in the unconstrained problem. With this study we show that the same holds when a worst-case delay constraint is imposed.
3) We establish an understanding of the trade-off between delay constraints and corresponding optimal achievable cost, and show specifically how the two conflicting objectives can affect each other.

In addition to the above, our problem formulation and abstraction generalize to a much larger class of optimization problems involving constrained resource allocation. This is discussed at the end of this paper.

The rest of the paper is organized as follows. In Section II we present the network model and assumptions used in this study. We then introduce our performance measures and objectives, along with the main results of this paper in Section III. In Section IV we derive the optimal worst-case strategies satisfying a delay constraint. These results are discussed and examined in Section V. Section VI concludes the paper.

II. Network Model

A. Model and Assumptions

We will limit our analysis to the case of searching for a single target, which is assumed to exist in the network. For the rest of our discussion we will use the term *object* to indicate the target of a search, be it a node, a piece of data or a file. Within the context of controlled flooding search, the distance between two nodes is measured in number of *hops*, assuming that the network is connected. Two nodes being one hop away means they can reach each other in one transmission.

We measure the position of an object by its distance to the source initiating the searching. We will use the term *object location* to indicate the minimum TTL value needed to locate the object, denoted by X. The term *network dimension* refers to the minimum TTL required to reach every node in the network, denoted by L. $\bar{F}_X(u) = P(X > u)$ denotes the tail distribution of the random variable X.

We will assume that a TTL value of u will reach all nodes within u hops of the source and will find the object with probability 1 if it is located within u hops, when the timer expires. This assumption implies that a timeout event is equivalent to not finding the object in the u-hop neighborhood, and that flooding the entire network will for sure locate the object. This is reasonable in a wired network, as long as packet loss is low and timeout values are properly set to sufficiently account for delay in the network. On the other hand, this assumption is a simplification in a wireless network because packet collisions and corruption losses can cause the query propagation process to be much more random and less reliable. This assumption nevertheless allows us to reveal some very interesting fundamental features of the problem and obtain valuable insights.

A search strategy $\mathbf{u}$ is a TTL sequence of certain length N, $\mathbf{u} = [u_1, u_2, \cdots, u_N]$. It can be either fixed/deterministic where $u_i, i = 1, \cdots, N$, are deterministic values, or random where u_i are drawn from probability distributions. For a fixed strategy we assume that $\mathbf{u}$ is an increasing sequence. For randomized strategies, we assume all realizations are increasing sequences.

The requirement for the sequence to be increasing is a natural one under the assumption that search with TTL u will always find the object if it is indeed within u of the source. Note that in a specific search experiment we may not need to use the entire sequence; the search stops whenever the object is found.

In practice, it is natural to only consider integer-valued (*discrete*) policies. However, considering real-valued sequences proves to be helpful in deriving optimal integer-valued strategies. For this reason we will also consider *continuous* (real-valued) strategies, denoted by $\mathbf{v}$, where $\mathbf{v} = [v_1, v_2, \cdots, v_N]$, and v_i is either a fixed or continuous random variable that takes real values. When considering discrete strategies, TTL values are integers and the object location X is assumed to be a positive integer taking values between 1 and L. In analyzing continuous strategies, X is assumed to be a real number in the interval $[1, L]$.

A strategy is *admissible* if it locates any object of finite location with probability 1. For a fixed strategy this implies $u_N = L$. For a random strategy, this implies $Pr(u_i = L) = 1$ for some $1 \leq i \leq N$. In the asymptotic case as $L \to \infty$, a strategy $\mathbf{u}$ is admissible if $\forall\, x \geq 1,\ \exists\, n \in \mathbb{Z}^+$ s.t. $Pr\,(u_n \geq x) = 1$. This implies that in the asymptotic case, $\mathbf{u}$ is an infinite-length vector. We let V denote the set of all real-valued admissible strategies (random and fixed). U denotes the set of all integer-valued admissible strategies (random and fixed).

B. Search Cost and Delay

We will associate a cost $C(u)$ with a single round of search using TTL value u. The functional form of this cost depends on the properties of the network as well as the underlying broadcast techniques used. In our analysis we will ignore these details and simply assume that such a function is obtainable, i.e., by estimating the number of transmissions and receptions, etc. $C(u)$ is therefore an abstraction of the lower layer properties, and for the rest of our discussion we will no longer regard network as wired or wireless, but only discuss in terms of the search cost $C(u)$.

Note that in general, a node receiving the search query will be unaware whether the object is found at another node in the same round (except perhaps when the object is found at one of its neighbors, or some other more sophisticated schemes are employed). Thus this node will continue the process by decrementing the TTL value and passing on the search query. We can therefore regard the search cost as being *paid in advance*, i.e., the search cost for each round is determined by the TTL value and not by whether the object is located in that round.

We next introduce the search delay function. We denote by $D_t(u)$ the timeout value used when searching with TTL u. This is the delay incurred when the object is not found using u, i.e., when $u < X$. On the other hand, if $u > X$, then the object will be found within this round of search. The delay incurred in this case is the amount of time it takes for the query to propagate X hops and for the reply to reach back to the source. We will denote this delay by $D_r(X)$ for object location X. Therefore mathematically the search delay of using TTL value u can be written as:

$$I(u < X) D_t\,(u) + I(X \leq u) D_r\,(X)\ \ ,$$

where I is the indicator function: $I(A) = 1$ if A is true and 0 otherwise.

For real-valued sequences, we require that the function $C(v)$ and $D(v)$ be defined for all $v \in [1, \infty)$, while for integer-valued sequences we only require that the cost and delay functions be defined for positive integers. When the cost function is invertible, we write $C^{-1}(\cdot)$ to denote its inverse. We will adopt the natural assumption that $C(v_1) > C(v_2)$ and $D(v_1) > D(v_2)$ if $v_1 > v_2$. We define the following class of cost functions for real-valued sequences.

Definition 1: The function $C : [1, \infty) \to [C(1), \infty)$ belongs to the class $\mathbb{C}$ if $0 < C(1) < \infty$, $C(v)$ is increasing and differentiable (hence continuous), and $\lim_{v \to \infty} C(v) = \infty$. Note that for every $y \in [C(1), \infty)$, there exists exactly one $v \in [1, \infty)$ such that $C(v) = y$.

When considering discrete strategies, we will restrict our results to the following subclass of $\mathbb{C}$:

Definition 2: A function $C(\cdot) \in \mathbb{C}$ belongs to the class $\mathbb{C}_q$ for some $q \geq 1$ if: (1) $\lim_{x \to \infty} \frac{C(x+1)}{C(x)} = q$ and (2) $\frac{C(x+1)}{C(x)} \geq q$ for all $x \in [1, \infty)$.

Note that since $C(\cdot)$ is strictly increasing, for $q = 1$ condition (2) is automatically satisfied. The case of $q = 1$ also contains all polynomial cost functions. The case of $q > 1$ includes for example exponential cost functions of the form q^x. Therefore even though this is a subclass of $\mathbb{C}$, it remains very general.

III. Problem Formulation and Main Results

A. Problem Formulation

We will consider the search performance in the asymptotic regime as $L \to \infty$. This is because it is difficult if at all possible to obtain a general strategy that is optimal for all finite-dimension networks as the optimal TTL sequence often depends on the specific value of L. In this sense, an asymptotically optimal strategy may provide much more insight into the intrinsic structure of the problem. It will become evident that asymptotically optimal TTL sequences also perform very well in a network of arbitrary finite dimension.

Let $J_X^{\mathbf{u}}$ denote the expected search cost of using strategy $\mathbf{u}$ when the object location is X. This quantity can be calculated as follows:

$$J_X^{\mathbf{u}} = E_{\mathbf{u}} E_X \left[\sum_{k=1}^{\infty} I\left(X > u_{k-1}\right) C(u_k) \right] = E_{\mathbf{u}} \left[\sum_{k=1}^{\infty} \bar{F}_X(u_{k-1}) C(u_k) \right] \ , \tag{1}$$

where $u_0 = 0$, $E_{\mathbf{u}}$ and E_X denote expectations with respect to $\mathbf{u}$ and X, respectively. The expectation and summation can be interchanged due to the Monotone Convergence Theorem [10]. We will drop the variable from the subscript when it is clear which variable the expectation is taken with respect to.

Similarly, let $D_X^{\mathbf{u}}$ denote the expected search delay induced by strategy $\mathbf{u}$ for X. This quantity can be calculated as follows:

$$D_X^{\mathbf{u}} = E_{\mathbf{u}} E_X \left[\sum_{k=1}^{\infty} \{ I(X > u_k) D_t(u_k) + I(u_k \geq X > u_{k-1}) D_r(X) \} \right] = E_{\mathbf{u}} \left[\sum_{k=1}^{\infty} \bar{F}_X(u_k) D_t(u_k) \right] + E_X \left[D_r(X) \right] \ . \tag{2}$$

When the distribution of X is known in advance, a natural objective is to determine strategies that minimize $J_X^{\mathbf{u}}$ subject to some constraint on $D_X^{\mathbf{u}}$. In general, such computations are numerical and the optimal solutions can be determined by standard constrained optimization techniques [11] [12]. In Section IV-B, we will derive the optimal strategy for a particular distribution of X and delay constraint under which the optimal strategy has a very interesting structure.

On the other hand, when the distribution of X is not known in advance, as is often the case, then we need a different approach. In this study we adopt a worst-case performance measure. Consider an omniscient observer who knows the object location in advance and will use a TTL of X, incurring an expected cost of $E[C(X)]$. We can then measure the performance of a strategy $\mathbf{u}$ by the following:

$$\rho^{\mathbf{u}} = \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \ , \tag{3}$$

where $\{p_X(x)\}$ denotes the set of all probability distributions for X such that $E[C(X)] < \infty$. The term $\rho^{\mathbf{u}}$ is an upper-bound, or worst-case measure, on the ratio between the cost of strategy $\mathbf{u}$ and the omniscient observer, over all X. We will refer to $\rho^{\mathbf{u}}$ as the competitive ratio, or *worst-case cost ratio*, of $\mathbf{u}$. This type of worst-case measure is commonly used in many online decision and computation problems [13]. It was introduced in [7] as a method of analyzing flooding strategies, and generalized in [8] to study randomized strategies.

We apply a similar worst-case analysis to delay. The minimum expected delay is $E[D_r(X)]$, obtainable by either an omniscient observer or a strategy that uses the highest TTL ($\mathbf{u} = [L]$ as $L \to \infty$). Hence the *worst-case delay ratio* is defined as:

$$\tau^{\mathbf{u}} = \sup_{\{p_X(x)\}} \frac{D_X^{\mathbf{u}}}{E[D_r(X)]} \ , \tag{4}$$

where we note in this case $\{p_X(x)\}$ is the set of all location distributions such that $E[D_r(X)] < \infty$. Note that the worst-case cost and delay ratios are always strictly greater than 1 for any admissible strategy as it is impossible to equal or do better than the omniscient observer.

We define the following set:

$$U_d = \left\{ \mathbf{u} \in U : \sup_{\{p_X(x)\}} \frac{D_X^{\mathbf{u}}}{E[D_r(X)]} \leq d \right\} \ , \tag{5}$$

for some constant $d > 1$. This is the set of all strategies whose delay is always within a factor d of the delay of the omniscient observer, regardless of X. We will call d the *delay constraint*. Note that as $d \to \infty$, the delay constraint becomes less restrictive and the set U_d approaches U.

We seek a strategy that satisfies this delay constraint d and has the smallest worst-case cost ratio, i.e. achieves the minimum worst-case cost ratio among all $\mathbf{u} \in U_d$:

$$\rho_d^* = \inf_{\mathbf{u} \in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \ . \tag{6}$$

This essentially constitutes our constrained optimization problem (P), rewritten as follows:

$$\inf_{\mathbf{u}} \quad \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \tag{7}$$

$$\text{s.t.} \quad \sup_{\{p_X(x)\}} \frac{D_X^{\mathbf{u}}}{E[D_r(X)]} \leq d \tag{8}$$

Note that the two supremums in (P), one in the objective function and the other in the constraint, are in general *not* achieved under the same distribution $p_X(x)$. The intention for adopting such a worst-case formulation, which may be viewed as somewhat conservative, is to place an upper bound on both the delay and the cost over all possible locations.

The above definitions also hold analogously for continuous strategies, by simply replacing U with V, and replacing the set $\{p_X(x)\}$ with $\{f_X(x)\}$, which is the set of density functions such that $E[C(X)] < \infty$, or $E[D_r(X)] < \infty$ depending on whether we consider worst-case cost or delay. We will thus denote $\rho^{\mathbf{v}}$, $\tau^{\mathbf{v}}$ and V_d as the continuous versions of (3), (4), (5), respectively. We will use the same notation ρ_d^* to denote the minimum worst-case cost ratio achieved by continuous strategies satisfying a delay constraint d; the distinction should be clear from the context. V_d is defined as follows for any $d > 1$:

$$V_d = \left\{ \mathbf{v} \in V : \sup_{\{f_X(x)\}} \frac{D_X^{\mathbf{v}}}{E[D_r(X)]} \leq d \right\} \ . \tag{9}$$

We now show that there is no loss in generality in assuming that $D_t(\cdot) = D_r(\cdot)$. Let $\tilde{D}_X^{\mathbf{u}}$ denote expected delay of strategy $\mathbf{u}$ for object location X *when these two functions are equal*. Then note the following:

$$\frac{D_X^{\mathbf{u}}}{E[D_r(X)]} = \frac{E[D_t(X)]}{E[D_r(X)]} \frac{E_{\mathbf{u}}\left[\sum_{k=1}^{\infty} \bar{F}_X(u_k) D_t(u_k)\right]}{E[D_t(X)]} + 1 = \frac{E[D_t(X)]}{E[D_r(X)]} \left(\frac{\tilde{D}_X^{\mathbf{u}}}{E[D_t(X)]} - 1 \right) + 1 \tag{10}$$

Hence, the delay ratio when $D_t \neq D_r$ is simply a rescaling of the ratio when D_t and D_r are the same functions. Specifically, a strategy $\mathbf{u}$ satisfies $D_X^{\mathbf{u}}/E[D_r(X)] \leq d$ if and only if:

$$\frac{\tilde{D}_X^{\mathbf{u}}}{E[D_t(X)]} \leq \frac{E[D_r(X)]}{E[D_t(X)]} (d-1) + 1 \tag{11}$$

Therefore, the set U_d that we defined for the case of $D_t \neq D_r$ can easily be redefined if $D_t = D_r$, by simply rescaling the delay constraint d. Note that this result holds in both the discrete and continuous cases. Therefore for the rest of the analysis, we will assume these two functions are equal while noting that the results apply to the unequal case by scaling the constant d. We let $D(u) = D_t(u) = D_r(u)$ for all u. It follows that using a TTL value u for object location X will incur a delay of $D(\min\{X, u\})$.

B. Main Results

In this section we present our main results to be proven and discussed later in this paper. We begin by examining optimal continuous strategies, i.e., finding the strategy in V_d that achieves minimum worst-case cost ratio. We define the following class of continuous strategies:

Definition 3: Assume that the cost function $C(\cdot) \in \mathbb{C}$. Let $\mathbf{v}[r, F_{v_1}(\cdot)]$ denote a jointly defined sequence $\mathbf{v} = [v_1, v_2, ...]$ with a configurable parameter r, generated as follows:

(J.1) The first TTL value v_1 is a continuous random variable taking values in the interval $\left[1, C^{-1}(r \cdot C(1))\right)$, with its cdf given by some nondecreasing, right-continuous function $F_{v_1}(x) = Pr(v_1 \leq x)$. Note that this means $F_{v_1}(1) = 0$ and $F_{v_1}\left(C^{-1}(r \cdot C(1))\right) = 1$.

(J.2) The k-th TTL value v_k is defined by $v_k = C^{-1}\left(r^{k-1} C(v_1)\right)$ for all positive integers k.

From (J.1) and (J.2), it can be seen that r and $F_{v_1}(\cdot)$ uniquely define the TTL strategy, and that given the selection of v_1, the cost of successive TTL values essentially form a geometric sequence of base r, i.e., $C(v_k) = r^{k-1} C(v_1)$. More discussion on this structure is given in Section V.

Our main theorem regarding the class of continuous strategies V is as follows.

Theorem 1: When $C(\cdot) \in \mathbb{C}$ and $C(\cdot) = \beta D(\cdot)^m$ for some $m, \beta > 0$, we have:

(1) For any fixed $1 < d < m+1$,

$$\inf_{\mathbf{v} \in V_d} \sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} = \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ . \tag{12}$$

Moreover, this minimum worst-case ratio is achieved by using the strategy $\mathbf{v}[r, \frac{1}{\ln r} \ln \frac{C(\cdot)}{C(1)}]$ with $r = e^{\frac{m}{d-1}}$.

(2) For $d \geq m+1$, we have:

$$\inf_{\mathbf{v} \in V_d} \sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} = e \ . \tag{13}$$

Moreover, this minimum worst-case ratio is achieved by using the strategy $\mathbf{v}[r, \frac{1}{\ln r} \ln \frac{C(\cdot)}{C(1)}]$ with $r = e$.

Note that the optimal strategy of Theorem 1 can be adjusted for different delay constraints by varying the parameter r. These optimal continuous strategies will be used to derive discrete strategies which perform well in the worst-case and are optimal for a subset of $\mathbb{C}$. In particular, we have the following.

Theorem 2: When $C(\cdot) \in \mathbb{C}$ and $C(\cdot) = \beta D(\cdot)^m$ for some $m, \beta > 0$, we have:

(1) For $1 < d < m+1$,

$$\inf_{\mathbf{u} \in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \leq \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ . \tag{14}$$

(2) For $d \ \geq \ m+1$,

$$\inf_{\mathbf{u} \in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \leq e \ . \tag{15}$$

Whether the upper bounds in Theorem 2 become equalities appears to depend on the specific cost function $C(\cdot)$. By restricting our attention to cost functions $C(\cdot) \in \mathbb{C}_1$ we have the following result.

Theorem 3: Consider $C(\cdot) \in \mathbb{C}_1$ and $C(\cdot) = \beta D(\cdot)^m$ for some $m, \beta > 0$.

(1) For $1 < d < m+1$,

$$\inf_{\mathbf{u} \in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} = \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ , \tag{16}$$

where this minimum worst-case ratio can be achieved by the discrete strategy $\mathbf{u}^*$ constructed as follows. Take the strategy $\mathbf{v}^*[e^{\frac{m}{d-1}}, \frac{d-1}{m} \ln \frac{C(\cdot)}{C(1)}]$ given by Definition 3, and set $u_k^* = \lfloor v_k^* \rfloor$ for all k to obtain the discrete strategy $\mathbf{u}^* = [u_1^*, u_2^*, ...]$.

(2) For $d \geq m+1$, we have

$$\inf_{\mathbf{u} \in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} = e \ . \tag{17}$$

Moreover, this minimum worst-case cost ratio is achieved by the strategy $\mathbf{u}^* = \lfloor \mathbf{v}^* \rfloor$, where $\mathbf{v}^*$ denotes the strategy $\mathbf{v}^*[e, \ln \frac{C(\cdot)}{C(1)}]$.

This result shows that we can take the floor of the optimal continuous strategy to obtain a discrete strategy which is optimal when the cost is a subclass of $\mathbb{C}$.

C. Discussion of Main Results

The main results described in the previous subsection are derived under a worst-case performance measure. This implies that for any object location, the optimal (for $1 < d < m+1$) strategy $\mathbf{v}$ of Theorem 1 has an expected search cost within $\frac{d-1}{m} e^{\frac{m}{d-1}}$ times the expected cost of the omniscient observer. Similarly, its expected delay is always within factor d of the delay incurred by an omniscient observer.

The differentiation between the two cases, $1 < d < m+1$ vs. $d \geq m+1$, in all three theorems is due to the fact that the optimization problem (P) has an active/binding constraint in the former, and an inactive/non-binding constraint in the latter, as we show in the next section.

The main results rely on the relationship $C(\cdot) = \beta D(\cdot)^m$ for some $m, \beta > 0$, where the factor m essentially describes the relative rate at which the cost and delay functions grow with respect to TTL. The first thing to note is that the constant positive factor β simply cancels out in the cost or delay ratio calculated in (3) and (4). Hence we can assume that $\beta = 1$ without loss of generality.

Secondly, the relationship $C(\cdot) = D(\cdot)^m$ holds, for example, in a very representative case of a searching in a 2-dimensional network with search cost proportional to the number of transmissions incurred. In this case $C(v)$ is well approximated by a quadratic function (see e.g., [7], [8]) and $D(v)$ can be chosen to be a linear function of v (implying $m = 2$), or quadratic (implying $m = 1$). Another scenario described by this relationship is when $m = 1$, where the cost and delay scale in the same fashion. This could be a good model in a linear network with constant node density where both cost and delay increase proportionally (linearly) to the number of transmissions.

IV. Optimal Strategies with Delay Constraints

In this section we prove the results shown in the previous section, i.e., the solution to problem (P). The solution approach we take is outlined as follows. We first (in Section IV-B) consider the continuous version of problem (P) and derive a tight lower-bound to the minimum worst-case cost under the delay constraint. This is accomplished by interchanging the $\inf$ and $\sup$ in Eqn. (6), and introducing a constrained optimization problem whose objective is to minimize the average search cost subject to a delay constraint. Then in Section IV-C we derive a class of randomized continuous strategies whose worst-case cost ratio matches this lower bound for all d, proving that they are optimal. These continuous strategies are then used in Section IV-D to derive good discrete strategies whose performance is at least as good in the worst case. We will also prove that they are optimal for a subclass of $\mathbb{C}$.

Unless otherwise stated, all proofs can be found in the Appendix.

A. Preliminaries

The following lemmas are critical in our subsequent analysis. We will let $J_x^{\mathbf{u}}$ and $D_x^{\mathbf{u}}$ denote the expected search cost and delay, respectively, of using strategy $\mathbf{u}$ when $P(X = x) = 1$.

Lemma 1: For any search strategy $\mathbf{v} \in V$ and $\mathbf{u} \in U$, we have

$$\sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} = \sup_{x \in [1,\infty)} \frac{J_x^{\mathbf{v}}}{C(x)} , \tag{18}$$

$$\sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} = \sup_{x \in \mathbb{Z}^+} \frac{J_x^{\mathbf{u}}}{C(x)} , \tag{19}$$

where $\mathbb{Z}^+$ denotes the set of natural numbers.

Proof of this lemma can be found in [9]. In words, this lemma states that the cost ratio is maximized when the object location is a single point. We also have an analogous lemma for search delay.

Lemma 2: For any search strategy $\mathbf{v} \in V$ and $\mathbf{u} \in U$,

$$\sup_{\{f_X(x)\}} \frac{D_X^{\mathbf{v}}}{E[D(X)]} = \sup_{x \in [1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x)} , \tag{20}$$

$$\sup_{\{p_X(x)\}} \frac{D_X^{\mathbf{u}}}{E[D(X)]} = \sup_{x \in \mathbb{Z}^+} \frac{D_x^{\mathbf{u}}}{D(x)} . \tag{21}$$

The proof of Lemma 2 can be found in [14]. These two lemmas reduce the space over which the worst-case cost or delay can occur, and thus are very useful in subsequent analysis.

B. A Tight Lower Bound

Consider any $d > 1$. To establish a tight lower bound to the minimum worst-case cost ratio, we can interchange infimum and supremum [11] to obtain the following:

$$\sup_{\{f_X(x)\}} \inf_{\mathbf{v} \in V_d} \frac{J_X^{\mathbf{v}}}{E[C(X)]} \leq \inf_{\mathbf{v} \in V_d} \sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} . \tag{22}$$

Any lower bound of the left-hand side of (22) can be found by fixing some object location distribution f_X and finding the strategy within V_d that minimizes the expected cost. Note that the strategy in V_d that minimizes the cost may be randomized, which makes the minimization very difficult.

Therefore, we further lower-bound the left hand side by considering a larger set of strategies than V_d. In particular, let $V_d(Y)$ denote the following set of strategies for some object location Y such that $E[D(Y)] < \infty$:

$$V_d(Y) = \left\{ \mathbf{v} \in V : \frac{D_Y^{\mathbf{v}}}{E[D(Y)]} \leq d \right\} . \tag{23}$$

Clearly, $V_d(Y) \supseteq V_d$ for any Y because any strategy $\mathbf{v} \in V_d$ has a delay ratio upper bounded by d for all object locations. Therefore,

$$\sup_{\{f_X(x)\}} \inf_{\mathbf{v} \in V_d(Y)} \frac{J_X^{\mathbf{v}}}{E[C(X)]} \leq \sup_{\{f_X(x)\}} \inf_{\mathbf{v} \in V_d} \frac{J_X^{\mathbf{v}}}{E[C(X)]} , \tag{24}$$

because for any object location X, the infimum on the right hand side is over a smaller set.

A valid lower bound of the left hand side of (24) can be obtained by choosing particular distributions for X and Y, and finding the strategy within $V_d(Y)$ that minimizes the expected cost. To obtain a tight lower bound, we need to find a combination of X and Y such that the optimal average-cost strategy under X satisfying the delay constraint induced by Y has a high expected cost ratio. It is important to note that it is *not* necessary that Y and X have the same distribution; this property allows us to obtain tight lower-bounds.

We consider the following problem (P1), whose solution not only provides a tight lower bound to (22) but also serves as an example for deriving optimal average-cost strategies subject to a delay constraint.

Problem 1: Suppose $C(\cdot) = \beta D(\cdot)^m$. Let $\bar{F}_X(x) = P(X > x) = \left(\frac{C(x)}{C(1)}\right)^{-\alpha}$, and $\bar{F}_Y(x) = P(Y > x) = \left(\frac{C(x)}{C(1)}\right)^{-\alpha+1-\frac{1}{m}}$, for some $\alpha > 1$ and for all $x \geq 1$. Consider the following constrained optimization problem:

$$\inf_{\mathbf{v}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} \quad \text{s. t. } \frac{D_Y^{\mathbf{v}}}{E[D(Y)]} \leq d \tag{25}$$

We solve the above problem for the following choice of α:
(1) If $1 < d < m+1$, choose α to be such that

$$1 < \alpha < 1 + \frac{m+1-d}{m(d-1)} \ . \tag{26}$$

(2) If $d \geq m+1$, choose any $\alpha > 1$.

The distinction between the two cases is that Problem 1 under the former ($1 < d < m+1$) has a binding constraint, while it has a non-binding constraint under the latter ($d \geq m+1$), which also means in this case Problem 1 reduces to an unconstrained optimization problem.

Solution: The optimal strategy $\mathbf{v}$ for this problem satisfies $C(v_j)/C(1) = \gamma^j$ for all j. The value of γ depends on d as follows (details can be found in the Appendix).

If $1 < d < m+1$, then γ is:

$$\gamma = \left(1 + \frac{(\alpha-1)m}{[(\alpha-1)m+1](d-1)}\right)^{\frac{1}{\alpha-1}} \ . \tag{27}$$

The optimal cost ratio for this case is given by:

$$\frac{J_X^{\mathbf{v}}}{E[C(X)]} = \left[\frac{d-1}{m\alpha} + d\left(\frac{\alpha-1}{\alpha}\right)\right]\gamma \ . \tag{28}$$

If $d \geq m+1$, then $\gamma = \alpha^{\frac{1}{\alpha-1}}$ and the optimal cost ratio is $\alpha^{\frac{1}{\alpha-1}}$.

Using this solution, we see that as α approaches 1 from above, the optimal cost ratio for the case $1 < d < m+1$ has the following limit:

$$\lim_{\alpha \to 1^+} \frac{J_X^{\mathbf{v}}}{E[C(X)]} = \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ , \tag{29}$$

where the limit is reached from below. When $d \geq m+1$, then the optimal cost ratio satisfies: $\lim_{\alpha \to 1^+} \alpha^{\frac{1}{\alpha-1}} = e$. Hence, the highest minimum cost ratio is lower bounded as follows:

Theorem 4: When $C(\cdot) = \beta D(\cdot)^m$ for $\beta, m > 0$, for any $1 < d < m+1$ the best worst-case cost ratio is lower bounded by the following:

$$\inf_{\mathbf{v} \in V_d} \sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} \geq \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ .$$

Therefore, any strategy in V_d which achieves a worst-case cost ratio of $\frac{(d-1)}{m} e^{\frac{m}{d-1}}$ must be optimal.

Similarly, when $d \geq m+1$, we have:

$$\inf_{\mathbf{v} \in V_d} \sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E[C(X)]} \geq e \ . \tag{30}$$

Therefore any strategy in V_d which achieves a worst-case cost ratio of e must be optimal.

We next derive strategies achieving the lower bounds established above.

C. Optimal Delay-Constrained Strategies

For convenience, we summarize the main results given in Section III in Figure 1. In this and the next subsection we will prove these results. We proceed to find strategies that match the lower bounds established in the previous subsection. To do so, we will consider strategies of the form $\mathbf{v}[r, F_{v_1}(x)]$ given by Definition 3.

Lemma 3: Assume $C(\cdot) \in \mathbb{C}$ and $C(\cdot) = \beta D(\cdot)^m$ for some $\beta, m > 0$. Then for any strategy $\mathbf{v}[r, F_{v_1}(\cdot)]$, its worst-case delay ratio is given by:

$$\sup_{x \in [1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x)} = \sup_{1 \leq z < r^{\frac{1}{m}}} \left\{\frac{1}{r^{\frac{1}{m}}-1} \frac{g(r^{\frac{1}{m}}) + (r^{\frac{1}{m}}-1)g(z)}{zD(1)} \frac{g'(z)}{D(1)}\right\} + 1 \ ,$$

	Continuous	**Discrete**
$1 < d < m+1$	Thm 1 (1) $\rho_d^* = \frac{d-1}{m} e^{\frac{m}{d-1}}$ optimal strategy: $\mathbf{v}^*[e^{\frac{m}{d-1}}, \frac{d-1}{m} \ln \frac{C(\cdot)}{C(1)}]$	Thm 2 (1) $\rho_d^* \leq \frac{d-1}{m} e^{\frac{m}{d-1}}$ Thm 3 (1), $C(\cdot) \in \mathbb{C}_1$ $\rho_d^* = \frac{d-1}{m} e^{\frac{m}{d-1}}$ optimal: $\mathbf{u}^* = \lfloor \mathbf{v}^* \rfloor$
$d \geq m+1$	Thm 1 (2) $\rho_d^* = e$ optimal strategy: $\mathbf{v}^*[e, \ln \frac{C(\cdot)}{C(1)}]$	Thm 2 (2) $\rho_d^* \leq e$ Thm 3 (2), $C(\cdot) \in \mathbb{C}_1$ $\rho_d^* = e$ optimal: $\mathbf{u}^* = \lfloor \mathbf{v}^* \rfloor$

Fig. 1. Summary of results on optimal worst-case strategies under (P).

where $g'(z)$ denotes the derivative of g with respect to z, and $g(z)$ is defined as follows for $1 \leq z < r^{\frac{1}{m}}$:

$$g(z) = D(1) + \int_{D(1)}^{z \cdot D(1)} \bar{F}_{v_1}(D^{-1}(y)) \, dy \ . \tag{31}$$

Consider the family of strategies of the form $\mathbf{v}[r, \frac{1}{\ln r} \ln \frac{C(\cdot)}{C(1)}]$, we have:

$$g(z) = D(1) \left[z - \frac{mz \ln z}{\ln r} + \frac{zm}{\ln r} - \frac{m}{\ln r} \right] \ , \tag{32}$$

and $g'(z) = D(1)(1 - m\frac{\ln z}{\ln r})$ for all z. We have the following results regarding this family of strategies:

(1) The worst-case delay ratio of these strategies is $\frac{m}{\ln r} + 1$. This is easily verified by using Lemma 3.
(2) The worst-case cost ratio of these strategies is $\frac{r}{\ln r}$. This result was proven in [9].

We consider two special cases of this family of strategies. The first case is when $r = e^{\frac{m}{d-1}}$ for some $1 < d < m+1$. With the above results, the worst-case delay ratio of this strategy is exactly d. Hence this specific strategy belongs to V_d. On the other hand its worst-case cost ratio is $\frac{(d-1)}{m} e^{\frac{m}{d-1}}$ (plugging $r = e^{\frac{m}{d-1}}$ into $\frac{r}{\ln r}$), achieving the lower bound established in Theorem 4.

The second case is when $r = e$. In this case we achieve a worst-case delay ratio of $m+1$, and the worst-case cost ratio of exactly e. Hence when $d \geq m+1$, this strategy belongs to V_d and is optimal since it matches the lower-bound established in Theorem 4. If $d > m+1$ then the delay constraint becomes inactive/non-binding under this strategy. Thus for $d > m+1$ this is also the solution to the unconstrained problem. This result was proven separately in [9] within the context of an unconstrained optimization problem, which we have now shown to be a special case of the more general result in this paper.

Combining these two cases together, we obtain Theorem 1. Therefore, we have obtained the optimal worst-case continuous strategies for any delay constraint $d > 1$.

D. Optimal Discrete Strategies

We are interested in deriving robust integer-valued strategies for the TTL-based controlled flooding search, i.e. finding $\mathbf{u} \in U_d$ achieving the minimum worst-case cost ratio. We will use our optimal continuous strategies of the previous subsection to derive discrete strategies that perform well in the worst-case. For notation, we will denote $\lfloor \mathbf{v} \rfloor$ to be the strategy $[\lfloor v_1 \rfloor, \lfloor v_2 \rfloor, \cdots]$.

We begin with the following lemma:

Lemma 4: For all $x \in \mathbb{Z}^+$, we have $D_x^{\lfloor \mathbf{v} \rfloor} \leq D_x^{\mathbf{v}}$ and $J_x^{\lfloor \mathbf{v} \rfloor} \leq J_x^{\mathbf{v}}$. That is, we can take the floor of any continuous strategy to find a discrete strategy that performs just as well if the object location is restricted to integers.

Using this result, we can prove Theorem 2. The proof is given in the Appendix.

This theorem gives an upper bound on the best worst-case discrete strategy, for all $C(\cdot) \in \mathbb{C}$. It appears that the actual value of the minimum worst-case cost will depend on the specific function $C(\cdot)$. A general result is currently not available, but if

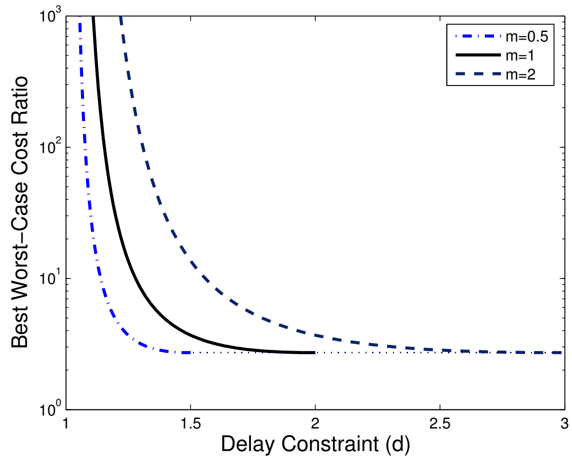

Fig. 2. When $C(\cdot) = \beta D(\cdot)^m$, a logarithmic plot of the minimum worst-case cost ratio as a function of the delay constraint d. Dotted portions indicate when the delay constraint is not binding and hence the unconstrained strategy of Theorem 1, part (2) is optimal. For $d \geq 3$, the best worst-case cost ratio is e for all three curves.

we restrict ourselves to cost functions of the simple polynomial form $C(\cdot) \in \mathbb{C}_1$, then we can obtain Theorem 3 presented earlier. This proof is provided in the Appendix.

V. Applications, Examples and Discussion

A. *Cost-Delay Tradeoff*

Having derived optimal strategies for any delay constraint, it is worth examining how the delay constraint affects the minimum achievable worst-case cost ratio. Figure 2 depicts the tradeoff between optimal worst-case cost ratio as given by Theorem 1 and the delay constraint d when $C(\cdot) = \beta D(\cdot)^m$. The dotted portion of each curve indicates when the delay constraint is not binding, i.e., for $d \geq m + 1 = 1.5, 2, 3$, respectively. In these cases the optimal unconstrained strategy (using $r = e$) has a minimum worst-case cost ratio of e. Note that the plot is logarithmic. As d approaches 1 from above, the best worst-case cost ratio approaches ∞ for all m. Hence, as the constraint on delay becomes tighter, the minimum worst-case cost increases unboundedly.

For any fixed d, as m increases the minimum worst-case cost also increases. This can be understood by fixing some delay function $D(\cdot)$. As m increases, the cost function $C(\cdot) = \beta D(\cdot)^m$ increases faster. For any given delay constraint, it then becomes more difficult to achieve a low cost ratio.

B. *Examples*

We present an example scenario where the search delay grows linearly in the TTL value used, while the search cost grows quadratically. Specifically, consider $D(x) = \beta x$ for all x and $C(x) = \xi x^2$ so $m = 2$. As mentioned earlier, this could be a good representation of a two-dimensional network, where transmissions are on the order of x^2, and the delay is proportional to number of hops.

From Theorem 1, the optimal strategy is $\mathbf{v}[e^{\frac{2}{d-1}}, (d-1)\ln x]$ whenever $1 < d < 3$. When $d \geq 3$, the optimal strategy is $v[e, 2\ \ln x]$. Figure 3 depicts the cost and delay ratio curves, with respect to object location, of the corresponding optimal strategies when $d = 1.5,\ 2$, and 3.

Note that both the delay and cost ratio curves approach their maximum values very rapidly. Hence, the worst-case value of cost and delay under asymptotic network size (as $L \to \infty$) can approximate the performance when the network size is finite. At the same time, the worst-case is approached asymptotically. Hence the cost ratio at any finite object location is less than the worst-case cost ratio.Also note that the cost and delay ratio curves are smooth and nearly flat with respect to object location. Thus the actual object location does not significantly change the performance of these strategies. One can view this as a built-in robustness for both the cost and delay criteria.

Similar results hold for other values of d and m, and other functional forms of $C(\cdot)$ and $D(\cdot)$. They are not repeated here.

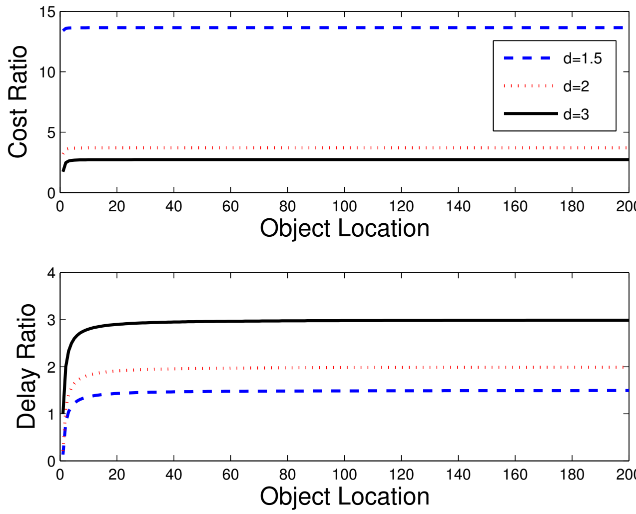

Fig. 3. Plot of cost and delay ratios of optimal strategies under different delay constraints, when cost is quadratic and delay is linear, i.e. $C(\cdot) = D(\cdot)^m$ for $m = 2$. Note that the delay ratio and cost ratio curves approach their maximum values very rapidly.

C. Comparison

It was shown in [8] and [9] that when adopting a worst-case cost measure, randomized strategies outperform deterministic ones. The results of the previous sections show that randomized strategies also perform better when delay constraints are added. Here we illustrate this in more detail.

Note that both the optimal deterministic strategy for Problem 1 and the optimal randomized strategies of Section IV-C share the property that the costs of the TTL values grow geometrically. That is, for any realization, $C(v_k) = r^{k-1}C(v_1)$ for all k. It was shown in [7] that the unconstrained optimal deterministic strategy under linear cost $C(\cdot)$ is also a geometric sequence: $u_k = 2^{k-1}$ for all k.

Below we compare deterministic and randomized geometric strategies to examine the effect of randomization. For deterministic geometric strategies with parameter r, $C(v_k) = r^{k-1}C(v_1)$ for all $k \geq 1$. Consider when both cost and delay are linear, so $D(v) = C(v) = v$ for all v. Then for any $k \geq 0$ and $x = r^k + \epsilon$, where $0 < \epsilon \leq r^{k+1} - r^k$, we have:

$$\frac{D_x^{\mathbf{v}}}{D(x)} = \frac{\sum_{j=0}^{k} r^j + r^k + \epsilon}{r^k + \epsilon} = \frac{r^{k+1} - 1}{(r-1)(r^k + \epsilon)} + 1 \ .$$

For each k, this ratio is maximized by taking the limit as ϵ approaches 0 from above. The maximum value of this ratio over all k is derived by letting $k \to \infty$, giving:

$$\sup_{x \in [1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x)} = \lim_{k \to \infty} \frac{r^{k+1} - 1}{(r-1)r^k} + 1 = \frac{2r-1}{r-1} \ , \tag{33}$$

which is strictly greater than 2 for all values of r. At the same time, similar calculations show that the worst-case cost ratio for such strategies is $\frac{r^2}{r-1}$.

Now consider the randomized strategies $\mathbf{v}[r, \frac{1}{\ln r} \ln \frac{C(\cdot)}{C(1)}]$, shown to be optimal in Theorem 1. Every realization of $\mathbf{v}$ is a geometric deterministic strategy with growth rate r. For any $r > 1$ and $m = 1$, it was shown that the worst-case cost ratio of $\mathbf{v}$ is $\frac{r}{\ln r}$ and the worst-case delay ratio is $\frac{1}{\ln r} + 1$.

In Figure 4 we plot the worst-case cost and delay ratios, as functions of r, for the aforementioned geometric deterministic and randomized strategies. Note that for any r, the randomized strategy achieves a lower worst-case cost *and* a lower worst-case delay than its deterministic counterpart. Hence, randomization has the effect of decreasing worst-case cost and delay at the same time.

In addition, note that the worst-case delay ratio of the randomized strategies approaches 1 as $r \to \infty$, but for the fixed strategies this limit is 2. In fact, for randomized geometric strategies using $r \geq e$, the worst-case delay ratio is always below 2. The class of optimal randomized strategies in Theorem 1 used $r \geq e$ for all values of d. Therefore even by arbitrarily

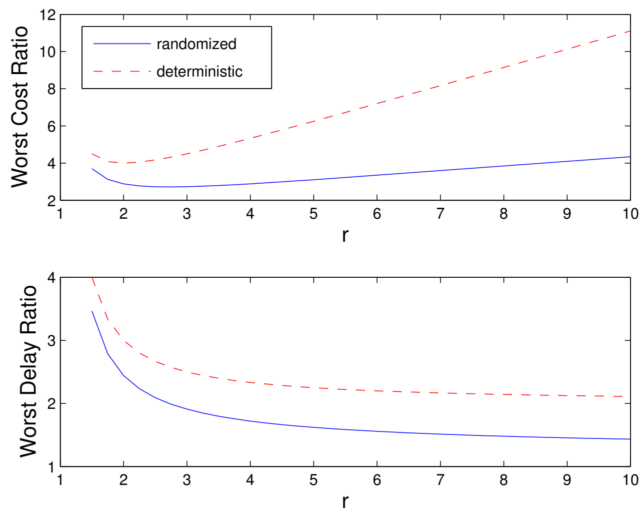

Fig. 4. For the strategies discussed in Section V-C, comparison of deterministic and randomized strategies as a function of r. Note that for any r, the randomization achieves lower worst-case cost ratio and delay ratio.

increasing the value of r for deterministic geometric strategies, it is not possible to match the worst-case delay ratio of the optimal randomized geometric strategies that we have derived in this study.

By varying the cost/delay functions and m, the curves in Figure 4 may change but the general relationship between randomized and deterministic strategies will still hold.

D. Generalization of the Problem Abstraction

We end our discussion by presenting the main problem abstraction adopted in this study in a more general fashion.

The problem introduced in this paper is primarily motivated by the flooding search application in networks. However, further consideration reveals some quite general and appealing features about the abstraction of this problem that can potentially be applied to a variety of problems involving constrained resource allocation. Here we restate the same problem in a more general context.

Consider an individual who seeks to complete a task (e.g., a computing job). There is a minimum level of resources/effort X required to accomplish the task (e.g., an updating step size in the computing job). X is a random variable whose distribution may be unknown to the individual. Its realization is not known in advance. The individual may choose from a range of resources/effort levels she is willing to put in the job, and the outcome (e.g., the precision of the computing result obtained) depends on the effort level. If she chooses a level $v \geq X$, then the task returns successfully and the process terminates. Otherwise the task returns failure and the individual increases her resources/effort level and try again. When a level v is chosen, the individual commits to paying a cost of $C(v)$ (e.g., memory and processing needed in the computing job), regardless of whether she succeeds or not. At the same time, with a level v the job takes a certain amount of time to return (either with a success or a failure), and this delay is given by $D\left(\min\left\{v, X\right\}\right)$.

The successive resource levels $\mathbf{v} = [v_1, v_2, \cdots]$ chosen by the individual form a *strategy*, which determines the total cost paid and time expended by the individual in accomplishing the job. As the cost is committed when a level is chosen, the individual must balance between selecting too low a level (more likely to be unsuccessful) and too high a level (more costly or wasteful). As can be seen, when one wishes to find a low cost strategy subject to a delay constraint, a constrained optimization problem is obtained. If furthermore the objective and the constraint are in the form of worst-case cost/delay measure, then a formulation akin to the one presented in this paper arises.

VI. Conclusion

In this paper we studied the class of TTL-based controlled flooding search and presented a constrained optimization framework in order to derive strategies that minimize a worst-case search cost measure subject to a worst-case search delay constraint. Optimal strategies were obtained in the continuous as well as discrete cases and their performance was studied. These results were used to discuss the trade-off between cost and delay using this type of search method. We also showed there the abstraction underlying the search application has a broad generalization that can be applied to solve a range of constrained resource allocation problems.

References

[1] D. Johnson and D. Maltz, "Dynamic source routing in ad hoc wireless networks," *Mobile Computing*, pp. 153–181, 1999.
[2] J. Xie, R. Talpade, T. McAuley, and M. Liu, "AMRoute: Ad Hoc Multicast Routing Protocol," *ACM Mobile Networks and Applications (MONET) Special Issue on Mobility of Systems, Users, Data and Computing*, vol. 7, no. 6, pp. 429–439, December 2002.
[3] C. Carter, S. Yi, P. Ratanchandani, and R. Kravets., "Manycast: Exploring the space between anycast and multicast in ad hoc networks," *Proceedings of the Ninth Annual International Conference on Mobile Computing and Networks (MobiCOM'03)*, September 2003, San Diego, California.
[4] D. Braginsky and D.Estrin, "Rumor routing algorithm for sensor networks," *Proc. International Conference on Distributed Computing Systems (ICDCS-22)*, 2002.
[5] S. Shakkottai, "Asymptotics of query strategies over a sensor network," *Proceedings of IEEE Infocom*, March 2004, Hong Kong.
[6] Z. Cheng and W. Heinzelman, "Flooding strategy for target discovery in wireless networks," *Proceedings of the Sixth ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2003)*, Sept 2003.
[7] Y. Baryshinikov, E. Coffman, P. Jelenkovic, P. Momcilovic, and D. Rubenstein, "Flood search under the california split rule," *Operations Research Letters*, vol. 32, no. 3, pp. 199–206, May 2004.
[8] N. Chang and M. Liu, "Revisiting the TTL-based controlled flooding search: Optimality and randomization," *Proceedings of the Tenth Annual International Conference on Mobile Computing and Networks (MobiCom'04)*, pp. 85–99, September 2004, Philadelphia, PA.
[9] N.B.Chang and M. Liu, "Optimal controlled flooding search in a large wireless network," *Third International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'05)*, pp. 229–237, April 2005, Riva Del Garda, Italy.
[10] P. Billingsley, *Probability and Measure*, John Wiley & Sons, 1995, New York, USA.
[11] S. Nash and A. Sofer, *Linear and Nonlinear Programming*, McGraw-Hill, 1996, New York, USA.
[12] E. Altman, *Constrained Markov Decision Processes*, Chapman & Hall/CRC, 1999, Boca Raton, Florida.
[13] A. Borodin and R. El-Yaniv, *Online Comptutation and Competetive Analysis*, Cambridge University Press, 1998, Cambridge, UK.
[14] N. Chang and M. Liu, "Controlled flooding search with delay constraints," *EECS Technical Report CGR 05-08*, 2005, University of Michigan, Ann Arbor.

VII. Appendix

A. Solution to Problem 1

To begin, we calculate the mean of object location cost and delay, noting that X takes values on $[1,\infty)$:

$$E\left[C(X)\right] = \int_0^\infty Pr\left(C\left(X\right) > x\right)\ dx = C(1) + \int_{C(1)}^\infty \left[\frac{C(C^{-1}(x))}{C(1)}\right]^{-\alpha}\ dx = \frac{\alpha}{\alpha-1}C(1)\ . \tag{34}$$

$$E\left[D(Y)\right] = E\left[C(Y)^{1/m}\right] = \int_0^\infty Pr\left(C\left(Y\right)^{1/m} > x\right)\ dx = \frac{1+(\alpha-1)m}{(\alpha-1)m}C(1)^{1/m}\ . \tag{35}$$

Using (2) to evaluate the delay ratio for (deterministic) $\mathbf{v}$:

$$\frac{D_Y^{\mathbf{v}}}{E[D(Y)]} = \frac{\sum_{k=1}^\infty \bar{F}_Y(v_k)D(v_k)}{E[D(Y)]} + 1 \tag{36}$$

By rearranging (36) and observing that only the numerator of the cost ratio depends on $\mathbf{v}$, Problem 1 is equivalent to:

$$\inf_{\mathbf{v}} J_X^{\mathbf{v}} \quad \text{s. t.} \ \sum_{k=1}^\infty \bar{F}_Y(v_k)D(v_k) \leq (d-1)E[D(Y)].$$

Therefore, we define the Lagrangian for $\lambda \geq 0$:

$$\begin{aligned} G(\mathbf{v},\lambda) &= \sum_{k=0}^\infty \bar{F}_X(v_k)C(v_{k+1}) - \lambda\left((d-1)E[D(Y)] - \sum_{k=1}^\infty \bar{F}_Y(v_k)D(v_k)\right) \\ &= \sum_{k=0}^\infty \left(\frac{C(v_k)}{C(1)}\right)^{-\alpha} C(v_{k+1}) - \lambda(d-1)E[D(Y)] + \lambda\sum_{k=1}^\infty \left(\frac{C(v_k)}{C(1)}\right)^{-\alpha+1-1/m} C(v_k)^{1/m} \end{aligned}$$

A necessary condition [11] for optimality of $\mathbf{v}$ is that the partial derivative of G with respect to v_j is 0, for all j. In other words,

$$\frac{\partial G}{\partial v_j} = \frac{\partial C(v_j)}{\partial v_j}C(1)^\alpha\left[C(v_{j-1})^{-\alpha} - \alpha C(v_{j+1})C(v_j)^{-\alpha-1} + \lambda(1-\alpha)C(v_j)^{-\alpha}C(1)^{\frac{1}{m}-1}\right] = 0 \tag{37}$$

Because the derivative of the cost function is strictly positive and $C(1) > 0$, then equation (37) is satisfied if and only if the term inside the brackets is equal to 0. Setting this term equal to 0, letting $\bar{\lambda} = \lambda C(1)^{\frac{1}{m}-1}$ for notational convenience, and rearranging yields the following recursion for $j \geq 1$:

$$C(v_{j+1}) = \frac{C(v_j)}{\alpha}\left[\left(\frac{C(v_j)}{C(v_{j-1})}\right)^{\alpha} + \bar{\lambda}(1-\alpha)\right] , \tag{38}$$

Hence, any optimal strategy must satisfy the recursion given by (38). Let $\gamma_j = \frac{C(v_j)}{C(v_{j-1})}$ for all $j \geq 1$. This quantity indicates the relative amount of cost increase after every unsuccessful search. Then (38) reduces to:

$$\gamma_{j+1} = \frac{1}{\alpha}\left(\gamma_j^{\alpha} - \bar{\lambda}(\alpha-1)\right) \tag{39}$$

Note that the value of γ_1 uniquely defines the remaining values γ_j for all $j \geq 2$. At the same time, the entire sequence $\{\gamma_j\}$ uniquely defines the strategy. Hence it remains to determine values of γ_1 that define optimal strategies.

Lemma 5: Fix $\bar{\lambda} \geq 0$. A necessary condition for optimality is that for all $j \geq 1$, $\gamma_j = \gamma$ where γ is the unique solution to the following equation:

$$\frac{1}{\alpha}\gamma^{\alpha} - \gamma = \bar{\lambda}\frac{(\alpha-1)}{\alpha} . \tag{40}$$

Hence, any optimal strategy for Problem 1 will have costs increasing geometrically by factor γ.

Proof: It should be noted for completeness that equation (40) has a unique solution because the function $\frac{1}{\alpha}x^{\alpha} - x$ is strictly increasing in x (this can be seen by differentiating with respect to x), is equal to 0 when $x = \alpha^{\frac{1}{\alpha-1}}$, is continuous, and increases to ∞ as $x \to \infty$. At the same time, $\bar{\lambda}\frac{(\alpha-1)}{\alpha}$ is a nonnegative finite quantity.

Note that if $\gamma_1 = \gamma$, then $\gamma_j = \gamma$ for all $j \geq 2$. Hence, it suffices to prove that $\gamma_1 = \gamma$ is necessary for optimality. We proceed by contradiction.

Case 1: $\gamma_1 > \gamma$.

Note that if $\gamma_j > \gamma$ for some j, then we have the following:

$$\gamma_{j+1} = \frac{1}{\alpha}\left(\gamma_j^{\alpha} - \bar{\lambda}(\alpha-1)\right) = \frac{1}{\alpha}\gamma_j^{\alpha} - \gamma_j + \gamma_j - \frac{\bar{\lambda}(\alpha-1)}{\alpha} > \frac{1}{\alpha}\gamma^{\alpha} - \gamma - \frac{\bar{\lambda}(\alpha-1)}{\alpha} + \gamma_j = \frac{\bar{\lambda}(\alpha-1)}{\alpha} - \frac{\bar{\lambda}(\alpha-1)}{\alpha} + \gamma_j = \gamma_j$$

where the last inequality follows from the fact that $\frac{1}{\alpha}x^{\alpha} - x$ is strictly increasing in x, as noted earlier. Hence we have the following: if $\gamma_j > \gamma$ for some j, then $\gamma_{j+1} > \gamma_j$. This means that because $\gamma_1 > \gamma$, then $\gamma_2 > \gamma_1 > \gamma$, and so on. Hence by induction, the $\{\gamma_j\}$ form a strictly increasing sequence, where $\gamma_j > \gamma$ for each j. So for each $j \geq 1$ by rearranging the recursion (39):

$$\frac{\gamma_{j+1}}{{\gamma_j}^{\alpha}} = \frac{1}{\alpha} - \bar{\lambda}\frac{(\alpha-1)}{\gamma_j^{\alpha}\alpha} \geq \frac{1}{\alpha} - \bar{\lambda}\frac{(\alpha-1)}{\gamma^{\alpha}\alpha} = \gamma^{1-\alpha} , \tag{41}$$

where the inequality holds because $\gamma_j^{\alpha} > \gamma^{\alpha}$, and the last equality holds from the definition of γ. The inequality becomes strict when $\bar{\lambda} > 0$.

Note that for any $j \geq 1$, we have by the definition of γ_j that $C(v_j) = C(1)\prod_{k=1}^{j}\gamma_k$. Hence, the expected cost of any such strategy is given by:

$$J_X^{\mathbf{v}} = \sum_{j=0}^{\infty}\bar{F}_X(v_j)C(v_{j+1}) = \sum_{j=0}^{\infty}\left(\frac{C(v_j)}{C(1)}\right)^{-\alpha}C(v_{j+1}) = \sum_{j=0}^{\infty}C(1)\gamma_1\prod_{k=1}^{j}\frac{\gamma_{k+1}}{\gamma_k^{\alpha}} , \tag{42}$$

where the product is defined to be equal to 1 if $j = 0$.

We have shown that if $\gamma_1 > \gamma$, then $\frac{\gamma_{k+1}}{{\gamma_k}^{\alpha}} > \gamma^{1-\alpha}$ for all k. Hence for any such strategy where $\gamma_1 > \gamma$, the expected search cost is lower-bounded by:

$$\sum_{j=0}^{\infty}C(1)\gamma_1\prod_{k=1}^{j}\frac{\gamma_{k+1}}{\gamma_k^{\alpha}} > \sum_{j=0}^{\infty}C(1)\gamma\prod_{k=1}^{j}\gamma^{1-\alpha} \tag{43}$$

However, note that the right-hand side of the above equation is simply the expected search cost for a strategy such that $\gamma_j = \gamma$ for all j (plug $\gamma_j = \gamma$ into 42). Hence from (43), any strategy where $\gamma_j > \gamma$ for all j has an expected search cost strictly greater than using $\gamma_j = \gamma$, and these strategies cannot be optimal.

Case 2: $\gamma_1 < \gamma$.

Note that for any optimal strategy, $\gamma_j > 1$ for all j, because only strictly increasing TTL sequences can be optimal. Hence the sequence $\{\gamma_j\}$ is always lower-bounded by 1. Note that if $\gamma_j < \gamma$ for some j, then we have the following:

$$\gamma_{j+1} = \frac{1}{\alpha}\left(\gamma_j^\alpha - \bar{\lambda}(\alpha-1)\right) = \frac{1}{\alpha}\gamma_j^\alpha - \gamma_j + \gamma_j - \frac{\bar{\lambda}(\alpha-1)}{\alpha} < \frac{1}{\alpha}\gamma^\alpha - \gamma - \frac{\bar{\lambda}(\alpha-1)}{\alpha} + \gamma_j = \frac{\bar{\lambda}(\alpha-1)}{\alpha} - \frac{\bar{\lambda}(\alpha-1)}{\alpha} + \gamma_j = \gamma_j$$

Hence, if $\gamma_1 < \gamma$, then $\gamma_2 < \gamma_1 < \gamma$, and so on.

Because the $\{\gamma_j\}$ are bounded, then the sequence converges (since all monotonic bounded sequences converge). Let $\gamma_\infty = \lim_{j\to\infty}\gamma_j$. Because the γ_j are strictly less than γ and form a decreasing sequence, then $\gamma_\infty < \gamma$. On the other hand, we have:

$$\lim_{j\to\infty}\gamma_{j+1} = \lim_{j\to\infty}\left\{\frac{1}{\alpha}\left[\gamma_j^\alpha - \bar{\lambda}(\alpha-1)\right]\right\} \Longrightarrow \gamma_\infty = \frac{1}{\alpha}\left[\gamma_\infty^\alpha - \bar{\lambda}(\alpha-1)\right] \Longrightarrow \frac{1}{\alpha}\gamma_\infty^\alpha - \gamma_\infty = \bar{\lambda}\frac{(\alpha-1)}{\alpha}$$

We defined γ as the unique number satisfying (40). Because we have just shown that γ_∞ is bounded and also satisfies the same equation we have that $\gamma_\infty = \gamma$. However, this contradicts the fact that $\gamma_\infty < \gamma$, which we showed earlier. Hence, it is not possible to have $\gamma_1 < \gamma$ if the $\{\gamma_j\}$ are bounded.

Therefore, combining Case 1 and Case 2 proves that $\gamma_j = \gamma$ for all j is the only possibly optimal strategy for fixed $\bar{\lambda}$. ■

For any strategy $\mathbf{v}$ where $\gamma_j = \gamma$, we have the following: $C(v_j) = C(1)\prod_{k=1}^{j}\gamma_k = C(1)\gamma^j$. Therefore we have the following geometric sum, which converges since $\gamma > 1$ (necessary for increasing sequence) implies $0 < \gamma^{1-\alpha} < 1$:

$$\sum_{k=1}^{\infty}\bar{F}_Y(v_k)D(v_k) = C(1)^{1/m}\sum_{k=1}^{\infty}\left(\frac{C(v_k)}{C(1)}\right)^{1-\alpha} = C(1)^{1/m}\sum_{k=1}^{\infty}(\gamma^{1-\alpha})^k = \frac{C(1)^{1/m}}{\gamma^{\alpha-1}-1} \quad (44)$$

Using (44) into (36), we can see that it is possible to achieve a delay ratio arbitrarily close to 1 by choosing a sufficiently high enough value of γ. Therefore, for every $d > 1$, there exists a strategy achieving delay ratio below d. Hence the optimal strategy for Problem 1 must satisfy the Kuhn-Tucker condition (see [11] and [12]):

$$\bar{\lambda}\left(\sum_{k=1}^{\infty}\bar{F}_Y(v_k)D(v_k) - (d-1)E[D(Y)]\right) = 0 \quad (45)$$

Therefore either $\bar{\lambda} = 0$ or the delay constraint is satisfied with equality. We use this to prove solutions for two cases, depending on whether $1 < d < m+1$ or $d \geq m+1$.

Case 1: $1 < d < m+1$

When $\bar{\lambda} = 0$, then $\lambda = 0$ and we have an unconstrained optimization problem. In this case, $\gamma = \alpha^{\frac{1}{\alpha-1}}$ from Lemma 5. The summation in (44) is then equal to $C(1)^{1/m}/(\alpha-1)$ for this value of γ. From (36), we know that dividing this summation by (35) and then adding 1 gives the delay ratio:

$$\frac{D_Y^{\mathbf{v}}}{E[D(Y)]} = \frac{\sum_{k=1}^{\infty}\bar{F}_Y(v_k)D(v_k)}{E[D(Y)]} + 1 = \frac{m}{1+(\alpha-1)m} + 1 \; .$$

From inequality (26) on α, this delay ratio is thus strictly greater than d. Hence, this strategy does not meet the delay inequality requirement.

Therefore, we seek solutions for which the delay constraint is met with equality, i.e. the term inside the brackets of (45) is equal to 0. In this case, (44) needs to be equal to $(d-1)C(1)^{\frac{1}{m}}\frac{(\alpha-1)m+1}{(\alpha-1)m}$, and solving for γ gives:

$$\gamma = \left(1 + \frac{(\alpha-1)m}{[(\alpha-1)m+1](d-1)}\right)^{\frac{1}{\alpha-1}} \quad (46)$$

From the earlier equation (40) relating γ and $\bar{\lambda}$, we have that $\bar{\lambda}$ can be calculated as:

$$\bar{\lambda} = \gamma\left(\frac{\gamma^{\alpha-1}-\alpha}{\alpha-1}\right) = \gamma\left(\frac{m}{[(\alpha-1)m+1](d-1)} - 1\right)$$

The cost ratio can be calculated by multiplying both sides of (38) by $\bar{F}_X(v_j)$ and then summing over $j \geq 1$ to give:

$$\sum_{j=1}^{\infty}\bar{F}_X(v_j)C(v_{j+1}) = \frac{1}{\alpha}\sum_{j=0}^{\infty}\bar{F}_X(v_j)C(v_{j+1}) - \bar{\lambda}\frac{\alpha-1}{\alpha}\sum_{j=1}^{\infty}C(v_j)\bar{F}_X(v_j)$$

The left-hand sum is simply $J_X^{\mathbf{v}} - C(v_1)$, so rearranging and solving for $J_X^{\mathbf{v}}$ gives:

$$J_X^{\mathbf{v}} \frac{\alpha-1}{C(1)\alpha} = \frac{C(v_1)}{C(1)} - \bar{\lambda}\frac{\alpha-1}{C(1)\alpha}\sum_{j=1}^{\infty} C(v_j)\bar{F}_X(v_j) \Longrightarrow \frac{J_X^{\mathbf{v}}}{E[C(X)]} = \frac{C(v_1)}{C(1)} - \bar{\lambda}\frac{\sum_{j=1}^{\infty} C(v_j)\bar{F}(v_j)}{E[C(X)]} \tag{47}$$

To evaluate this ratio, note that:

$$\sum_{j=1}^{\infty} C(v_j)\bar{F}(v_j) = C(1)\sum_{k=1}^{\infty}\left(\frac{C(v_k)}{C(1)}\right)^{1-\alpha} = \frac{C(1)}{\gamma^{\alpha-1}-1} = C(1)(d-1)\left(1+\frac{1}{(\alpha-1)m}\right)$$

Dividing by $E[C(X)]$ from equation (34) gives:

$$\frac{\sum_{j=1}^{\infty} C(v_j)\bar{F}(v_j)}{E[C(X)]} = (d-1)\left(\frac{(\alpha-1)m+1}{m\alpha}\right) \tag{48}$$

Using (46), the corresponding $\bar{\lambda}$, and (48) into (47) gives:

$$\begin{aligned}\frac{J_X^{\mathbf{v}}}{E[C(X)]} &= \frac{C(v_1)}{C(1)} - \bar{\lambda}(d-1)\left(\frac{(\alpha-1)m+1}{m\alpha}\right) = \gamma - \gamma\left(\frac{m}{[(\alpha-1)m+1](d-1)} - 1\right)(d-1)\left(\frac{(\alpha-1)m+1}{m\alpha}\right)\\ &= \left[\frac{d-1}{m\alpha} + d\left(\frac{\alpha-1}{\alpha}\right)\right]\gamma \end{aligned} \tag{49}$$

Hence (49) gives the optimal cost ratio when $1 < d < m+1$.

Case 2: $d \geq m+1$.

As explained in Case 1, it follows from Lemma 5 that when $\bar{\lambda} = 0$ (the unconstrained case), then using $\gamma = \alpha^{\frac{1}{\alpha-1}}$ is optimal. Because this strategy is the optimal unconstrained strategy, it achieves minimum average-cost when it satisfies the delay constraint. It was shown that the delay ratio for this strategy is $1 + \frac{m}{1+(\alpha-1)m}$, which is always strictly less than $m+1$ for all $\alpha > 1$. Hence for $d \geq m+1$, the delay constraint is not binding and the optimal strategy uses $\gamma = \alpha^{\frac{1}{\alpha-1}}$. From (47), the optimal cost ratio is $\alpha^{\frac{1}{\alpha-1}}$ because $\bar{\lambda} = 0$.

B. Proof of Lemma 3

Proof: Because $C(v_k) = r^{k-1}C(v_1)$ for all positive integers k, then $D(v_k) = r^{\frac{k-1}{m}}D(v_1)$. Let $S_k = D(v_1) + D(v_2) + \cdots + D(v_k)$. Then for any k,

$$\begin{aligned} E\left[S_k\right] &= E\left[\sum_{j=1}^{k} D(v_j)\right] = E\left[\sum_{j=1}^{k} r^{\frac{j-1}{m}} D(v_1)\right] = \sum_{j=0}^{k-1} r^{\frac{j}{m}} E\left[D(v_1)\right]\\ &= E\left[D(v_1)\right]\sum_{j=0}^{k-1} r^{\frac{j}{m}} = E\left[D(v_1)\right]\frac{r^{\frac{k}{m}}-1}{r^{\frac{1}{m}}-1} \end{aligned} \tag{50}$$

For any $1 \leq l < C^{-1}(rC(1))$,

$$\begin{aligned} E\left[D(v_1)|v_1 \leq l\right] &= \int_0^{\infty} Pr\left(D(v_1) > y | v_1 \leq l\right)\, dy = D(1) + \int_{D(1)}^{\infty} \frac{Pr\left(D^{-1}(y) < v_1 \leq l\right)}{Pr(v_1 \leq l)}\, dy\\ &= D(1) + \frac{1}{F_{v_1}(l)}\left[\int_{D(1)}^{D(l)}\left[\bar{F}_{v_1}(D^{-1}(y)) - \bar{F}_{v_1}(l)\right]\, dy\right] = \frac{1}{F_{v_1}(l)}\left[D(1) + \int_{D(1)}^{D(l)} \bar{F}_{v_1}(D^{-1}(y))\, dy - D(l)\cdot\bar{F}_{v_1}(l)\right] \end{aligned}$$

Consider any real number $x \geq 1$. There must exist a positive integer n such that $r^{n-1}C(1) \leq C(x) < r^n C(1)$. Equivalently, $r^{\frac{n-1}{m}}D(1) \leq D(x) < r^{\frac{n}{m}}D(1)$. Then we have:

$$\begin{aligned} D_x^{\mathbf{v}} &= D_x^{\mathbf{v}}|_{v_n > x} Pr\left(v_n > x\right) + D_x^{\mathbf{v}}|_{v_n \leq x} Pr\left(v_n \leq x\right)\\ &= E\left[S_{n-1} + D(x)|v_n > x\right] Pr\left(v_n > x\right) + E\left[S_n + D(x)|v_n \leq x\right] Pr\left(v_n \leq x\right) \end{aligned} \tag{51}$$

$$\begin{aligned} &= E[S_{n-1}] + D(x) + E[D(v_n)|v_n \leq x] Pr(v_n \leq x)\\ &= \frac{r^{\frac{n-1}{m}}-1}{r^{\frac{1}{m}}-1} E\left[D(v_1)\right] + D(x) + r^{\frac{n-1}{m}} E\left[D(v_1)\middle| v_1 \leq D^{-1}\left(\frac{D(x)}{r^{\frac{n-1}{m}}}\right)\right] F_{v_1}\left(D^{-1}\left(\frac{D(x)}{r^{\frac{n-1}{m}}}\right)\right) \end{aligned}$$

Using (51), we obtain the following:

$$
\begin{aligned}
D_x^{\mathbf{v}} &= \frac{r^{\frac{n}{m}}-1}{r^{\frac{1}{m}}-1}E\left[D(v_1)\right] + D(x) + r^{\frac{n-1}{m}}\cdot\left[D(1) + \int_{D(1)}^{\frac{D(x)}{r^{\frac{n-1}{m}}}} \bar{F}_{v_1}(D^{-1}(y))\ dy - \frac{D(x)}{r^{\frac{n-1}{m}}}\bar{F}_{v_1}\left(D^{-1}\left(\frac{D(x)}{r^{\frac{n-1}{m}}}\right)\right)\right] \\
&= \frac{r^{\frac{n-1}{m}}}{r^{\frac{1}{m}}-1}\left[E\left[D(v_1)\right] + (r^{\frac{1}{m}}-1)\left\{D(1) + \int_{D(1)}^{\frac{D(x)}{r^{\frac{n-1}{m}}}} \bar{F}_{v_1}\left(D^{-1}(y)\right)\ dy\right\}\right] \\
&\quad + D(x) - D(x)\bar{F}_{v_1}\left(D^{-1}\left(\frac{D(x)}{r^{\frac{n-1}{m}}}\right)\right) - \frac{E\left[D(v_1)\right]}{r^{\frac{1}{m}}-1}
\end{aligned}
\tag{52}
$$

Letting $z = \frac{D(x)}{r^{\frac{n-1}{m}}D(1)}$, we obtain the following:

$$
\begin{aligned}
\frac{D_x^{\mathbf{v}}}{D(x)} &= \frac{1}{r^{\frac{1}{m}}-1}\left[\frac{E[D(v_1)] + (r^{\frac{1}{m}}-1)\left\{D(1) + \int_{D(1)}^{zD(1)} \bar{F}_{v_1}(D^{-1}(y))\ dy\right\}}{zD(1)}\right] + 1 - \bar{F}_{v_1}(D^{-1}(z)) - \frac{E[D(v_1)]}{(r^{\frac{1}{m}}-1)zr^{\frac{n-1}{m}}D(1)} \\
&= \frac{1}{r^{\frac{1}{m}}-1}\left[\frac{g(r^{\frac{1}{m}}) + (r^{\frac{1}{m}}-1)g(z)}{zD(1)}\right] + 1 - \frac{g'(z)}{D(1)} - \frac{g(r^{\frac{1}{m}})}{(r^{\frac{1}{m}}-1)zr^{\frac{n-1}{m}}D(1)}
\end{aligned}
\tag{53}
$$

Only the last (righthandmost) term depends on n. As x goes to ∞, this term goes to 0. Hence, the supremum of the delay ratio becomes:

$$
\sup_{x\in[1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x)} = \sup_{1\le z<r^{\frac{1}{m}}} \left\{\frac{1}{r^{\frac{1}{m}}-1}\frac{g(r^{\frac{1}{m}}) + (r^{\frac{1}{m}}-1)g(z)}{zD(1)} - \frac{g'(z)}{D(1)}\right\} + 1\ , \tag{54}
$$

■

C. Proof of Theorem 2

Proof: For $1 < d < m+1$, consider the strategy $\mathbf{v}^*[e^{\frac{m}{d-1}}, \frac{d-1}{m}\ln\frac{C(\cdot)}{C(1)}]$ of Theorem 1. Let $\mathbf{u}^* = \lfloor \mathbf{v}^* \rfloor$. From Lemma 4, we have that $\mathbf{v}^* \in V_d$ implies $\mathbf{u}^* \in U_d$. Also from Lemma 4, $J_x^{\mathbf{u}^*} \le J_x^{\mathbf{v}^*}$ for all integers x. From Theorem 1, the worst case ratio of $\mathbf{v}^*$ is $\frac{d-1}{m}e^{\frac{m}{d-1}}$. Hence, the worst-case cost ratio of $\mathbf{u}^*$ over all integers is less than or equal to $\frac{d-1}{m}e^{\frac{m}{d-1}}$, which establishes the theorem for $1 < d < m+1$.

For $d \ge m+1$, similar steps can be applied to the floor of $\mathbf{v}^*[e, \ln\frac{C(\cdot)}{C(1)}]$ to establish the theorem. ■

D. Proof of Lemma 4

Proof: Fix any $x \in \mathbb{Z}^+$. Note that $\lfloor v_k \rfloor > x$ if and only if $v_k > x$, since x is an integer. Therefore for all k: $I(\lfloor v_k \rfloor > x) = I(v_k > x)$ $w.p.1$, (with probability 1). In addition, $D(\lfloor (v_k) \rfloor) \le D(v_k)$ $w.p.1$, since the delay function is increasing. This gives us:

$$
D_x^{\lfloor \mathbf{v} \rfloor} = D(x) + E\left[\sum_{k=1}^{\infty} I(\lfloor v_k \rfloor > x)D(\lfloor v_k \rfloor)\right] \le D(x) + E\left[\sum_{k=1}^{\infty} I(v_k > x)D(v_k)\right] = D_x^{\mathbf{v}}\ .
$$

This proves the delay part of the lemma. Similarly, $C(\cdot)$ being increasing implies $C(\lfloor v_k \rfloor) \le C(v_k)$ $w.p.1$. Therefore,

$$
J_x^{\lfloor \mathbf{v} \rfloor} = E\left[\sum_{k=1}^{\infty} I(\lfloor v_{k-1} \rfloor > x)C(\lfloor v_k \rfloor)\right] \le E\left[\sum_{k=1}^{\infty} I(v_{k-1} > x)C(v_k)\right] = J_x^{\mathbf{v}}\ ,
$$

which establishes the inequality on the expected cost. ■

E. Proof of Theorem 3

Because $\mathbb{C}_1 \subset \mathbb{C}$, then from Theorem 2 the best worst-case cost ratio is upper-bounded as follows:

(1) For $1 < d < m+1$,

$$
\inf_{\mathbf{u}\in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \le \frac{(d-1)}{m}e^{\frac{m}{d-1}}\ . \tag{55}
$$

(2) For $d \geq m+1$,

$$\inf_{\mathbf{u}\in U_d} \sup_{\{p_X(x)\}} \frac{J_X^{\mathbf{u}}}{E[C(X)]} \leq e \ . \tag{56}$$

Hence, to prove the Theorem it suffices to prove that the minimum worst cost ratio is upper bounded by the above quantities.

Consider any $\mathbf{u} \in U$. For any integer $x \geq 2$:

$$\frac{J_x^{\mathbf{u}}}{C(x)} = \lim_{\epsilon\to 0} \frac{J_{x-1+\epsilon}^{\mathbf{u}}}{C(x+\epsilon)} = \sup_{y\in[x-1,x)} \frac{J_y^{\mathbf{u}}}{C(y+1)} \ , \tag{57}$$

since $J_{x-1+\epsilon}^{\mathbf{u}} = J_x^{\mathbf{u}}$ for all $0 < \epsilon \leq 1$, and $C(\cdot)$ is strictly increasing and continuous. Hence we have:

$$\sup_{x\in\mathbb{Z}^+} \frac{J_x^{\mathbf{u}}}{C(x)} = \sup\left\{\frac{J_1^{\mathbf{u}}}{C(1)}, \sup_{x\in\mathbb{Z}^+,\ x\geq 2}\ \sup_{y\in[x-1,x)} \frac{J_y^{\mathbf{u}}}{C(y+1)}\right\} = \sup\left\{\frac{J_1^{\mathbf{u}}}{C(1)}, \sup_{x\in[1,\infty)} \frac{J_x^{\mathbf{u}}}{C(x+1)}\right\} \tag{58}$$

We can repeat the above steps for delay to obtain for any $\mathbf{u} \in U$:

$$\sup_{x\in\mathbb{Z}^+} \frac{D_x^{\mathbf{u}}}{D(x)} = \sup\left\{\frac{D_1^{\mathbf{u}}}{D(1)}, \sup_{x\in[1,\infty)} \frac{D_x^{\mathbf{u}}}{D(x+1)}\right\} = \sup_{x\in[1,\infty)} \frac{D_x^{\mathbf{u}}}{D(x+1)} \ , \tag{59}$$

because $\frac{D_1^{\mathbf{u}}}{D(1)} = 1$.

Next, it can be shown similarly to Lemmas 1 and 2:

$$\sup_{\{f_X(x)\}} \frac{J_X^{\mathbf{v}}}{E\left[C(X+1)\right]} = \sup_{x\in[1,\infty)} \frac{J_x^{\mathbf{v}}}{C(x+1)} \ , \tag{60}$$

$$\sup_{\{f_X(x)\}} \frac{D_X^{\mathbf{v}}}{E\left[D(X+1)\right]} = \sup_{x\in[1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x+1)} \ . \tag{61}$$

Define $\tilde{V}_d$ similarly to V_d in (9), but replace $D(X)$ with $D(X+1)$. Also, define $\tilde{V}_d(Y)$ similarly to $V_d(Y)$ in (23) but replace $D(X)$ with $D(X+1)$. In other words:

$$\tilde{V}_d = \left\{\mathbf{v}\in V :\ \sup_{x\in[1,\infty)} \frac{D_x^{\mathbf{v}}}{D(x+1)} \leq d\right\} \tag{62}$$

$$\tilde{V}_d(Y) = \left\{\mathbf{v}\in V :\ \frac{D_Y^{\mathbf{v}}}{E\left[D(Y+1)\right]} \leq d\right\} \tag{63}$$

Note that from (59) and $U \subseteq V$, we have that $U_d \subseteq \tilde{V}_d$. Therefore, if we can prove a lower-bound on the best cost ratio in $\tilde{V}_d$, the same bound will also apply to U_d. To proceed, we use the following:

$$\sup_{\{f_X(x)\}} \inf_{\mathbf{v}\in\tilde{V}_d(Y)} \frac{J_X^{\mathbf{v}}}{E\left[C(X+1)\right]} \leq \sup_{\{f_X(x)\}} \inf_{\mathbf{v}\in\tilde{V}_d} \frac{J_X^{\mathbf{v}}}{E\left[C(X+1)\right]} \leq \inf_{\mathbf{v}\in\tilde{V}_d} \sup_{x\in[1,\infty)} \frac{J_x^{\mathbf{v}}}{E\left[C(x+1)\right]} \tag{64}$$

We will use this inequality to obtain a tight upper-bound. In order to do this, we use the following:

Lemma 6: Let X_α be the random variable with tail distribution $P(X_\alpha > x) = \left(\frac{C(x)}{C(1)}\right)^{-\alpha}$, and let Y_α be the random variable with $P(Y_\alpha > x) = \left(\frac{C(x)}{C(1)}\right)^{-\alpha+1-\frac{1}{m}}$, for some $\alpha > 1$ and for all $x \geq 1$. Then we have for $C(\cdot) \in \mathbb{C}_1$:

$$\lim_{\alpha\to 1^+} \frac{E[C(X_\alpha+1)]}{E[C(X_\alpha)]} = q \ . \tag{65}$$

$$\lim_{\alpha\to 1^+} \frac{E[D(Y_\alpha+1)]}{E[D(Y_\alpha)]} = q^{\frac{1}{m}} \ . \tag{66}$$

$$\tag{67}$$

Proof: Because $\frac{C(x+1)}{C(x)} \geq q$ for all x, we have that $\frac{E[C(X_\alpha+1)]}{E[C(X_\alpha)]} \geq q$ for all α. Hence to complete the proof, we need to show that for any $\epsilon > 0$, there exists $\bar{\alpha}$ such that $\frac{E[C(X_\alpha+1)]}{E[C(X_\alpha)]} < q + \epsilon$ for all $1 < \alpha < \bar{\alpha}$.

Fix $\epsilon > 0$. Since $\lim_{x\to\infty} \frac{C(x+1)}{C(x)} = q$, there exists a x^* such that $\frac{C(x+1)}{C(x)} < q + \frac{\epsilon}{2}$ for all $x > x^*$. Let $1(\cdot)$ denote the indicator function; so $1(A) = 1$ if A is true, otherwise it equals 0. Thus we have:

$$\begin{aligned} E[C(X_\alpha+1)1(X_\alpha > x^*)] <& \left(q+\frac{\epsilon}{2}\right) E[C(X_\alpha)1(X_\alpha > x^*)] \\ \leq& \left(q+\frac{\epsilon}{2}\right) E[C(X_\alpha)] \end{aligned} \tag{68}$$

At the same time, we have:

$$\lim_{\alpha\to 1^+} \frac{E[C(X_\alpha+1)1(X_\alpha \leq x^*)]}{E[C(X_\alpha)]} \leq \lim_{\alpha\to 1^+} \frac{C(x^*+1)}{E[C(X_\alpha)]} = 0 \ ,$$

because $C(x^*+1) < \infty$ and $E[C(X_\alpha)] = \frac{\alpha}{\alpha-1}$, which approaches ∞ as α goes to 1. Hence, there exists an $\bar{\alpha}$ such that for all$1 < \alpha < \bar{\alpha}$:

$$\frac{E[C(X_\alpha+1)1(X_\alpha \leq x^*)]}{E[C(X_\alpha)]} < \frac{\epsilon}{2} \tag{69}$$

Therefore, combining (68) and (69) gives for all $1 < \alpha < \bar{\alpha}$

$$\frac{E[C(X_\alpha+1)]}{E[C(X_\alpha)]} = \frac{E[C(X_\alpha+1)1(X_\alpha > x^*)]}{E[C(X_\alpha)]} + \frac{E[C(X_\alpha+1)1(X_\alpha \leq x^*)]}{E[C(X_\alpha)]} < q + \frac{\epsilon}{2} + \frac{\epsilon}{2} = q + \epsilon \ ,$$

which completes the proof for the cost function.

Note that because $C(\cdot) \in \mathbb{C}_q$ and $D(\cdot) = C(\cdot)^{1/m}$, then $D(\cdot) \in \mathbb{C}_{q^{1/m}}$. Hence we can repeat the above steps, replacing X_α with Y_α and $C(\cdot)$ with $D(\cdot)$ to prove the delay portion of the lemma. ■

By modifying Problem 1 and using the above lemma, we obtain the following tight upper-bound:

Problem 2: Define X_α and Y_α as in Lemma 6.

$$\inf_{\mathbf{v}} \frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha+1)]} \quad \text{s.t.} \ \frac{D^{\mathbf{v}}_{Y_\alpha}}{E[D(Y_\alpha+1)]} \leq d$$

Solution: Similar to Problem 1, the optimal cost ratio for this problem satisfies the following:

(1) For $1 < d < m+1$:

$$\lim_{\alpha\to 1^+} \frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha+1)]} = \frac{(d-1)}{m} e^{\frac{m}{d-1}} \ , \tag{70}$$

where the limit is reached from below.

(2) For $d \geq m+1$, then:

$$\lim_{\alpha\to 1^+} \frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha+1)]} = e \ , \tag{71}$$

Proof: Because only the numerator of the cost ratio depends on $\mathbf{v}$, then Problem 2 is equivalent to the following:

$$\inf_{\mathbf{v}} J^{\mathbf{v}}_{X_\alpha} \quad \text{s. t.} \ \sum_{k=1}^{\infty} \bar{F}_{Y_\alpha}(v_k) D(v_k) \leq \left(d \ \frac{E[D(Y_\alpha+1)]}{E[D(Y_\alpha)]} - 1\right) E[D(Y_\alpha)].$$

We note that for α sufficiently close to 1, the above problem can be solved by using the solution to Problem 1. To see this, define the following:

$$\tilde{d} = d \ \frac{E[D(Y_\alpha+1)]}{E[D(Y_\alpha)]} \tag{72}$$

Note that $\lim_{\alpha\to 1^+} \tilde{d} = d$ from Lemma 6.

If $1 < d < m+1$, then this implies:

$$\lim_{\alpha\to 1^+} (\tilde{d}-1)[1+(\alpha-1)m] = d-1 < m \ . \tag{73}$$

Hence there must exist an $\bar{\alpha}$ such that for all $1 < \alpha < \bar{\alpha}$, we have: $(\tilde{d}-1)[1+(\alpha-1)m] < m$. Rearranging this gives for all $1 < \alpha < \bar{\alpha}$:

$$1 < \alpha < 1 + \frac{m+1-\tilde{d}}{m(\tilde{d}-1)} \ , \tag{74}$$

which is equivalent to (26), with $\tilde{d}$ replacing d. Note also that $(\tilde{d}-1)[1+(\alpha-1)m] < m$ implies $\tilde{d} < m+1$. Hence, the above problem is equivalent to part (1) of Problem 1, with $\tilde{d}$ replacing d. So using (49), we have the following for the optimal strategy:

$$\frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha)]} = \left[\frac{\tilde{d}-1}{m\alpha} + \tilde{d}\left(\frac{\alpha-1}{\alpha}\right)\right]\gamma \tag{75}$$

Therefore we have the following limit:

$$\lim_{\alpha\to1^+}\frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha+1)]} = \lim_{\alpha\to1^+}\frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha)]}\frac{E[C(X_\alpha)]}{E[C(X_\alpha+1)]} = \lim_{\alpha\to1^+}\left[\frac{\tilde{d}-1}{m\alpha} + \tilde{d}\left(\frac{\alpha-1}{\alpha}\right)\right]\gamma\cdot\lim_{\alpha\to1^+}\frac{E[C(X_\alpha)]}{E[C(X_\alpha+1)]} \tag{76}$$

$$= \frac{d-1}{m}e^{\frac{m}{d-1}} \ . \tag{77}$$

This establishes (70).

If $d \geq m+1$, then $\tilde{d} > d > m+1$. Hence we can use the solution to part (2) of Problem 1, to obtain:

$$\lim_{\alpha\to1^+}\frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha+1)]} = \lim_{\alpha\to1^+}\frac{J^{\mathbf{v}}_{X_\alpha}}{E[C(X_\alpha)]}\frac{E[C(X_\alpha)]}{E[C(X_\alpha+1)]} = \lim_{\alpha\to1^+}\alpha^{\frac{1}{\alpha-1}}\cdot\lim_{\alpha\to1^+}\frac{E[C(X_\alpha)]}{E[C(X_\alpha+1)]} = e \ , \tag{78}$$

which establishes equation (71). ■

Plugging this result into (64) gives that for $1 < d < m+1$:

$$\inf_{\mathbf{v}\in\tilde{V}_d}\sup_{x\in[1,\infty)}\frac{J^{\mathbf{v}}_x}{C(x+1)} \geq \frac{d-1}{m}e^{\frac{m}{d-1}} \ . \tag{79}$$

Finally, using the fact that $U_d \subseteq \tilde{V}_d$ gives:

$$\inf_{\mathbf{u}\in U_d}\sup_{x\in\mathbb{Z}^+}\frac{J^{\mathbf{v}}_x}{C(x)} = \inf_{\mathbf{u}\in U_d}\sup_{x\in[1,\infty)}\frac{J^{\mathbf{v}}_x}{C(x+1)} \geq \frac{d-1}{m}e^{\frac{m}{d-1}} \ .$$

Combining this with Lemma 1 and Theorem 2 proves the theorem for $1 < d < m+1$.

Similar steps can be applied for the case $d \geq m+1$.