Metrical Service Systems: Deterministic Strategies

Marek Chrobak [1] Lawrence L. Larmore [1]

UCR-CS-93-1

[1]Department of Computer Science, University of California, Riverside, CA 92521. Research supported by NSF grant CCR9112067.

Abstract

In this paper we introduce a new model of on-line problems called *metrical service systems*. An instance of a metrical service system is specified by a metric space M and a set $\mathcal{R}$ of *requests*, where each $r \in \mathcal{R}$ is a subset of M. There is one *server* that can move in M. At every time step a request $r \in \mathcal{R}$ is specified, and the server must move to a point $x \in r$ to serve the request r. We also have the following *on-line* restriction: the choice of x must be made before future requests are given. The sequence of points occupied by the server during this process is called a *service*. The goal is to minimize the service cost, which is the total distance moved by the server. Even for simple metrical service systems computing the optimal service on-line is impossible. Therefore, we concentrate on so-called *competitive* strategies that compute a service that only approximates an optimal one. More precisely, a strategy is called *c-competitive* if, for each request sequence, it computes a service whose cost is bounded by c times the optimal service cost of this sequence plus a constant.

First, we develop a theory of on-line strategies for metrical service systems. In particular, we show that the existence of an on-line c-competitive strategy is equivalent to the existence of a *c-potential function*. The concept of the *minimum c-potential* is particularly useful in lower-bound proofs.

We illustrate our method using two examples. We consider the *k-point request problem*, a metrical service system where each request set has at most k points. We give optimally competitive deterministic strategies for two cases: (1) for arbitrary k and uniform M, and (2) for $k = 2$ and arbitrary M. These solutions, and especially the lower bound proofs, illustrate the usefulness of the techniques described in the first part of the paper.

0.1 Introduction

On-line optimization problems. Most computational problems dealt with in the classic theory of algorithms are *off-line*, in the sense that they involve the assumption, explicit or not, that all data are present on input before the computation begins. This paradigm was justified by the nature of computational problems found in early computer applications, but it does not apply to algorithmic problems, arising in some contemporary applications, in which only partial information about the data is available. Such problems are not given in the form: "given x, compute $f(x)$."

An important class of such problems is that of *on-line optimization problems*, in which input data are given incrementally at certain time steps, and the algorithm must output partial results after each step, without knowledge of data that will arrive later. The task is to compute a solution that is close to optimal, in spite of the fact that only partial information about the input data is available.

For example, many scheduling problems are on-line: a scheduler must assign a processor (or, say, a memory block) to a given task without any knowledge of the tasks that will arrive later. The goal is to minimize the number of used processors (memory blocks, *etc.*). The problem is that optimality of a schedule is a global property, depending on the whole sequence of task specifications. This means that the allocation of a given task in an optimal schedule may depend on the data that will arrive in the future. For this reason, on-line algorithms normally cannot compute an optimal schedule.

An on-line problem can be viewed as a game, in which the input data are generated by an adversary, who attempts to make the algorithm perform as badly as possible. Because of this formulation, we use the term *on-line strategy* instead of *on-line algorithm*. Another, perhaps more important reason for using this term, is that some problems we consider are not finite, nor even discrete, and strategies theoretically need to remember an uncountable amount of information.

Competitiveness. Motivated by scheduling and other on-line optimization problems, researchers have recently begun to develop a theory of on-line strategies. Since for such problems computing an optimal solution is usually impossible, the notion of *competitiveness* has been introduced. Informally, a strategy is called *c-competitive* if it computes a solution whose cost is asymptotically within c times the optimal cost.

The server problem. One famous example of an on-line problem is the *k-server problem* [13], defined as follows. Suppose that there are k mobile *servers*, each of which must, at any time, be located at some point in a given metric space M. At every time step, a *request* $r \in M$ is given. An on-line strategy needs to choose one of the servers and move it to r, to *serve* the request. The goal is the minimize the total *cost*, defined to be the total distance traveled by the servers in M.

There is a known lower bound of k for the competitiveness of the k-server problem, and the *k-server conjecture* by Manasse, McGeoch, and Sleator [13] states that this bound can be achieved, *i.e.*, that there is a k-competitive strategy for the k-server problem, for each integer k. The conjecture is known to be true for $k = 2$ [4, 13]. Among other results, Fiat *et al.* [10] and Grove [12] proved that, for each k, a c_k-competitive strategy exists for the k-server problem, where c_k is an exponential function of k. Chrobak and Larmore [5] give an 11-competitive strategy for the 3-server problem.

Metrical service systems. Our task has been to develop a formalism for investigation of on-line problems. Such a formalism should encompass a wide variety of on-line problems, and at the same time it should produce some general techniques for constructing on-line competitive strategies, and for proving lower bounds for the competitive constant.

In this paper we present a model which we call *metrical service systems* (MSS). A metrical service system $\mathcal{S}$ is specified by a metric space M and a set $\mathcal{R}$ of *requests*, where each request $r \in \mathcal{R}$ is a subset of M. The input data consist of an initial server position x^0 and a sequence of requests $\varrho = r^1 r^2 \ldots r^n$. There is one server that resides initially at x^0. At time t, request r^t must be *served* by moving the server to some point $x^t \in r^t$. This is an on-line problem, because the choice of $x^t \in r^t$ must be made before future requests $r^{t+1} \ldots r^n$ are given. A strategy $\mathcal{A}$ for $\mathcal{S}$ is called *c-competitive* if, for each x^0 and ϱ, it finds a service whose cost is at most c times the optimal service, plus possibly an additive constant independent of ϱ. We say that a strategy $\mathcal{A}$ is *optimally competitive* if its competitiveness constant is the smallest possible.

By specifying different request sets, one can obtain different on-line problems. In particular, the k-server problem in a metric space M can be expressed as a metrical service system. Let M' be the set of all server configurations, that is all k-element multisets of points in M. This is a metric space under the *minimum matching* metric, defined as follows: The distance between two k-multisets $u, v \in M'$, for $u = \{u_1, \ldots, u_k\}$ and $v = \{v_1, \ldots, v_k\}$, is $\min_\pi \sum u_i v_{\pi(i)}$, where the minimum is over all permutations π of $\{1, 2, \ldots, n\}$. (For $x, y \in M$, xy denotes the distance between x and y.) Associate each $x \in M$ with the set r_x consisting of all k-multisets that contain x. Then the k-server problem for M can be viewed as an MSS in M' where the requests are the sets r_x.

Our results. First, we present a number of general results about metrical service systems. We prove that the existence of an on-line c-competitive strategy is equivalent to the existence of a so-called *c-potential* function. Various potential function methods have been used in the past to prove competitiveness of on-line strategies. Our result gives a solid theoretical foundation for this technique by showing that, in order to prove c-competitiveness, one *must*, explicitly or not, provide a c-potential function. Additionally, we show that if there is a c-competitive strategy $\mathcal{A}$ then, under some very general topological conditions, there exists a so-called *work-function based* c-competitive strategy $\mathcal{A}'$, that is, a strategy whose complete information about the past is described by a *work function* which, for each $x \in M$, gives the optimal cost of serving the past request sequence and ending at x. In the general case, without any assumption on the topological structure of M, a slightly weaker result holds, since $\mathcal{A}'$, besides the current work function, needs to keep track of the number of previous requests.

We also show that if there is any c-potential function, then there is also a unique *minimum* c-potential function, and that this function is a fixed point of a certain *boost* operator. The minimum c-potential possesses certain symmetry properties that give rise to a new lower-bound technique, which can be used in place of the well-known adversary method. We also show a method to compute such a minimum c-potential function for a given metrical service system $\mathcal{S}$. If $\mathcal{S}$ is finite and all distances in M are integers, then this method converges in a finite number of steps. We use this algorithm in our experimental work to compute optimally competitive strategies for some finite metrical service systems.

In the second part of the paper we consider the *k-point request problem* ($k-$PRP), in which the set $\mathcal{R}$ consists of all non-empty subsets of a given metric space M with cardinality at most k. We investigate two instances of this problem: arbitrary k in uniform spaces (all distances in M equal), and $k = 2$ in arbitrary metric spaces.

For the $k-$PRP in uniform spaces, we give a k-competitive strategy and we prove that it is optimally competitive.

For the $2-$PRP in arbitrary metric spaces, we give a 9-competitive strategy, and we prove it is optimally competitive for metric spaces in general by showing that no smaller competitiveness can be achieved if M is the real line.

Both algorithms are special cases of a so-called λ-*Cheap-and-Lazy* strategy, denoted CL_λ, that balances two different greedy approaches to the problem: one that always moves the server to the nearest point in the request, and another that always moves the server to the point which minimizes the new work function (this point is the server location in the optimal service for the request sequence consisting of all past requests). CL_λ minimizes the weighted sum: (distance) + $\lambda\cdot$(value of the new work function). The smaller λ, the "lazier" the strategy is. Our strategy for $k-$PRP in uniform spaces is identical to CL_λ for any $\lambda > 1$. For $2-$PRP we use strategy CL_3.

The lower bound proofs in both examples illustrate the usefulness of the approach based on the minimum c-potential. Previously, lower bounds on the competitive constant were usually proven using an adversary argument. The idea of such an argument is to show that an adversary, by choosing an appropriate request sequence, can force any on-line strategy to return to its initial configuration, incurring positive cost which is also at least c times the optimal cost. Since this "cycle" can be repeated an arbitrary number of times, this implies a lower bound of c on the competitive constant. In some problems, this adversary strategy ends when the algorithm is not in the initial configuration, but in one that is very similar to it, and thus we cannot simply repeat the same strategy, since this time the algorithm may behave differently. This phenomenon occurs, for example, in the case of $2-$PRP, where the constant 9 cannot be forced in a cycle. This problem is difficult to handle with the adversary argument. However, the properties of the minimal c-potential imply that, without loss of generality, a c-competitive algorithm makes the same moves in such similar situations (as formalized in Section 0.3). In Section 0.5 we show how to use this approach to obtain a simple proof of the lower bound of 9 for the competitiveness of $2-$PRP.

The cow problem and layered graph traversal. Baeza-Yates, Culberson and Rawlins [1] considered the so-called "cow problem," in which a cow needs to find a gate in an infinitely long fence. They also use competitive analysis; their goal was to minimize the worst case ratio between the distance traveled by the cow and the distance to the gate. They present a 9-competitive algorithm for this problem and prove that the constant 9 is optimal. It is possible to exploit the relationship between the cow problem and the $2-$PRP on the line to get a different lower bound proof of 9 on the competitive constant. A more general problem, called *layered graph traversal*, was introduced by Papadimitriou and Yannakakis [14], who gave a 9-competitive optimal algorithm for width-2 graphs. Fiat *et al.* [9] proved that $k-$PRP and the width-k graph traversal problem have the same competitive constants, which gives another proof of the upper bound of 9 on $2-$PRP. They also prove that there exists an $O(9^k)$-competitive algorithm for width-k graphs, and that 2^{k-2} is a lower bound for the competitiveness of this problem.

Other formalisms. A number of other models of on-line problems have been proposed in the literature. Friedman and Linial [11] introduced a model equivalent to metrical service systems which they call *set chasing*. In our terminology, they consider the MSS $\mathrm{CR}_2 = \left\langle \mathbb{R}^2, \mathcal{C} \right\rangle$, where $\mathbb{R}^2$ is the real plane with the Euclidean metric, and $\mathcal{C}$ is the set of all convex subsets of $\mathbb{R}^2$, and they prove that this problem is c-competitive, for some constant c not specified explicitly in their paper.

A more general model, called *task systems*, was given by Borodin *et al.* [3]. A task system is specified by a 4-tuple $\langle M, \mu, T, f \rangle$, where M is a set of states, μ is a state-to-state distance function (not necessarily symmetric), T is a set of tasks (requests), and f is a cost function that gives the cost of serving a given task from a given state. If μ is symmetric, then the system is called *metrical*. The main result from [3] is that each n-state metrical task system has a $(2n-1)$-competitive strategy, and that this constant is optimal, in the sense that no better constant is possible for some systems with n states. Each finite MSS is clearly also a metrical task system, so their upper bound extends to finite metrical service systems.

Besides the models discussed above, there are also other, even more general, models of on-line problems, including *on-line games*, introduced by the authors in [7], and *request-answer games*, introduced by Ben-David *et al.* in [2]. This work is based, in part, on some ideas and results announced in [7].

Comments on infinite spaces. There is a difference in emphasis between our work and that of [3]. We are interested in investigating problems that have a c-competitive strategy, independently of the size of the underlying metrical space. In fact, in this paper, we do not assume that our spaces are finite, nor even that they are discrete. Thus one can say that we concentrate on problems where competitiveness is determined by the structure of the request set, not by the number of states. Our results can be applied even if the whole space is not known in advance, and at every step only the points in the current request are revealed.

Investigating infinite spaces is important, both for theoretical and practical reasons. Theoretical motivations are obvious. Some problems are competitive in arbitrary spaces, and thus there is no reason to restrict the research to finite cases. Another reason is that certain phenomena can occur only in infinite spaces.

But there are also practical aspects. First, infinite spaces do occur in practice. A robot in a room has infinitely many possible locations. In some applications the actual space may be unknown, in which case we need to use the infinite space of all "imaginable" positions. An important source of problems where configuration spaces are infinite are randomized problems, in which the space of probability distributions is always infinite. One could argue that, instead of dealing with an infinite space, we can prove competitiveness for all finite "subproblems," where a subproblem is obtained by restricting the space and the requests to a finite subset of a metric space. This argument is invalid. Results about competitiveness for subspaces do not carry over to the whole space. It is quite easy to construct an example of an MSS that does not have a competitive strategy, but in which every finite subproblem has a 1-competitive strategy. There is also an infinite MSS with a 1-competitive strategy, in which there are finite subproblems of arbitrarily high competitiveness. (Recall that, by [3], all finite metrical service systems have competitive strategies.) Finally, competitiveness proofs for the infinite case may be actually simpler than those for finite subproblems. Restriction to a subproblem introduces additional structure that may sometimes obscure the essence of a problem and complicate the proofs.

0.2 Preliminaries

In this section we give a formal definition of metrical service systems, and introduce the concept of a *work function*.

0.2.1 Definition of Metrical Service Systems

Recall that a metric space is specified by a set M and a distance function that to each pair of points $x, y \in M$ assigns the distance between x and y. For simplicity, we will denote the metric space by M, and the distance between x and y will be denoted by xy. The distance function satisfies the following restrictions: for all $x, y, z \in M$, (1) $xx = 0$, (2) $xy > 0$ if $x \neq y$, and (3) $xy + yz \geq xz$. The last condition is called the *triangle inequality*. We assume the reader is familiar with the basic definitions and properties which relate to metric spaces, such as *limit*, *closed set*, *compact set*, and *continuous function*, as well as standard theorems relating to these definitions. For example, we will frequently use the theorem that a continuous real-valued function on a compact metric space achieves a minimum and a maximum.

A *metrical service system* is a pair $\mathcal{S} = \langle M, \mathcal{R}\rangle$, where M is a metric space and $\mathcal{R}$ is a family of non-empty subsets of M. We assume[1] that $\bigcup \mathcal{R} = M$. Each $r \in \mathcal{R}$ is called a *request*. Let $x^0 \in M$, and $\varrho = r^1 \dots r^n \in \mathcal{R}^*$ be a sequence of requests. We say that a sequence $\sigma = x^1 \dots x^n$ is a *service for* ϱ if $x^t \in r^t$ for each $t = 1, \dots, n$. The *cost* of service σ, starting from x^0, is the total distance traveled by the server. Denote by $|\pi|$ the length of the path $\pi = x^0 \dots x^n$ in M, that is $|\pi| = \sum_{t=1}^n x^{t-1}x^t$. In particular, the cost of a service σ starting from x^0 equals $|x^0\sigma|$.

Let $opt(x^0, \varrho)$ denote the optimal cost of serving ϱ starting from x^0, that is

$$opt(x^0, \varrho) \ = \ \inf_{\sigma} |x^0\sigma|,$$

where the infimum is over all services σ for ϱ. If there exists a service sequence σ for ϱ such that $|x^0\sigma| = opt(x^0, \varrho)$, then σ is called an *optimal service* for ϱ starting from x^0.

With $\mathcal{S}$ we associate the following optimization problem:

Problem OPTIMAL_SERVICE($\mathcal{S}$):

Instance: $x^0 \in M$, $\varrho \in \mathcal{R}^*$.

Task: Find an optimal service for ϱ starting from x^0.

Note that this problem may not have a solution. The optimal cost $opt(x^0, \varrho)$ is always well-defined, but there may not be an optimal service that realizes the infimum in the definition of $opt(x^0, \varrho)$, unless the given MSS satisfies certain topological conditions. (See Theorem 0.2.2.1.) In the *on-line* version of Problem OPTIMAL_SERVICE($\mathcal{S}$), we additionally require that the point x^t be written before requests $r^{t+1}, \dots, r^n$ are read. See Section 0.3 for detailed treatment of on-line strategies.

0.2.2 Compactly Shielded Systems

As we shall see later, in some metrical service systems certain "pathological" phenomena may occur that do not seem to appear in systems that arise naturally from applications. In this section we will define a class of "nice" systems in which such pathological phenomena do not occur, and which possess certain properties that are desirable both from the theoretical and practical point of view. The definition below seems rather complicated; our intention is to give a condition that would cover most known instances of MSS problems, including all finite systems, the $k-$PRP, and the metrical service systems arising in the k-server problem.

[1] This is only a technical assumption. If a point does not belong to any $r \in \mathcal{R}$ then it will never be used, so we can ignore it. Having such useless points introduces some technical difficulties and unnecessarily complicates other definitions.

Given sets $U, V, Y \subseteq M$, we say that a set Y is a *shield between* U and V if for all $u \in U$ and $v \in V$ there is a $y \in Y$ such that $uv = uy + yv$. We say that an MSS $\mathcal{S} = \langle M, \mathcal{R} \rangle$ is *compactly shielded* if the following conditions hold:

Closed requests: Every $r \in \mathcal{R}$ is closed, *i.e.*, contains all of its own limit points.
Compact shielding: For any compact $X \subseteq M$, any $r \in \mathcal{R}$, and any bounded set $Z \subseteq r$, there exists a compact shield $Y \subseteq r$ between X and Z.

Intuitively, compactly shielded systems have the following property: Suppose that the server is somewhere on X, and that the new request r is bounded, for simplicity. Then, by taking $Z = r$, no sensible strategy would move the server anywhere except to a point of Y. If r is not bounded, a similar property holds, by showing that we can restrict attention to bounded subsets of r.

If M has the property that bounded and closed sets are compact, then the first condition implies the second – simply let Y be the closure of Z. Euclidean metric spaces and many others have this property.

The $k-$PRP is compactly shielded for any metric space, because every request is finite, hence compact. The MSS for the k-server problem is also compactly shielded, independently of the underlying metric space M, as follows. For any $x = \{x_1 \dots x_k\} \in \Lambda^k M$, and for each $q \in M$, let $Y_{x,q} = \{x - \{x_i\} \cup \{q\} \mid i = 1, \dots, k\}$. Then r_q is shielded from X by $Y = \cup_{x \in X} Y_{x,q}$. If X is compact, Y is compact. We leave the details to the reader.

Theorem 0.2.2.1 *If $\mathcal{S}$ is a compactly shielded MSS, then, for each $x^0 \in M$ and $\varrho \in \mathcal{R}^*$, there exists an optimum service for ϱ starting from x^0.*

Proof: Fix x^0 and $\varrho = r^1 \dots r^n$. Let B be the ball of radius $opt(x^0, \varrho) + 1$ centered at x^0, and let $Z^t = r^t \cap B$. Let $Y^0 = \{x^0\}$, and recursively choose Y^t to be a compact shield between Z^t and Y^{t-1}.

We show that only service sequences drawn from $Y^1 \times \dots \times Y^n$ need be considered. Pick a service $\sigma = x^1 \dots x^n$ for ϱ whose cost is within 1 of optimal, *i.e.*, such that

$$|x^0\sigma| = \sum_{i=1}^{n} x^{t-1}x^t < opt(x^0, \varrho) + 1.$$

Then $x^t \in Z^t$ for all t. We recursively define a service $\sigma' = y^1 \dots y^n$, where $y^t \in Y^t$ for $t = 1, \dots, n$, whose cost does not exceed the cost of σ. Let $y^0 = x^0$, and choose $y^t \in Y^t$ such that $y^{t-1}x^t = y^{t-1}y^t + y^t x^t$. Then

$$|y^0\sigma'| = \sum_{t=1}^{n} y^{t-1}y^t = \sum_{t=1}^{n} \left(y^{t-1}x^t - y^t x^t\right) \leq \sum_{t=1}^{n} \left(y^{t-1}x^t - y^{t-1}x^{t-1}\right) \leq \sum_{t=1}^{n} x^{t-1}x^t = |x^0\sigma|.$$

Finally, continuity of cost on the compact set $Y^1 \times \dots \times Y^n$ implies existence of an optimal service.
□

We will need the following property of compact spaces.

Lemma 0.2.2.2 *Let X be a compact metric space. Let $\{f^n\}$ be a monotone increasing sequence of real functions defined on X, such that $\sup_n f^n(x)$ exists for each $x \in X$, and such that, for some constant $\lambda > 0$, $f^n(x) - f^n(y) \leq \lambda xy$ for all $x, y \in M$ and all n. Then*

$$\sup_n \inf_{x \in X} f^n(x) \;=\; \inf_{x \in X} \sup_n f^n(x).$$

Proof: ($\leq$) Write $f^* = \sup_n f^n$. Since $f^n \leq f^*$, we have $\inf_{x \in X} f^n(x) \leq \inf_{x \in X} f^*(x)$, and consequently $\sup_n \inf_{x \in X} f^n(x) \leq \inf_{x \in X} f^*(x)$.

($\geq$) Without loss of generality, $\lambda = 1$. Fix an arbitrary $\epsilon > 0$. Let $a = \inf_{x \in X} f^*(x)$. By compactness, there is a finite set $Y \subseteq X$ such that $\min_{y \in Y} xy \leq \epsilon$ for each $x \in X$. Since Y is finite and $\min_{y \in Y} f^*(y) \geq a$, there exists an integer m_ϵ such that $\min_{y \in Y} f^n(y) \geq a - \epsilon$ for all $n \geq m_\epsilon$. Now, by the choice of Y, we have $\inf_{x \in X} f^n(x) \geq a - 2\epsilon$ for $n \geq m_\epsilon$, and consequently $\sup_n \inf_{x \in X} f^n(x) \geq a - 2\epsilon$. This completes the proof, because ϵ can be made arbitrarily small. □

0.2.3 Work Functions

Let $\mathbb{R}$ denote the set of real numbers. Throughout this section, fix a metrical service system $\mathcal{S} = \langle M, \mathcal{R} \rangle$.

Work functions. We say that $\omega : M \to \mathbb{R}$ is a *work function* if the following conditions hold:

(wf1) $\omega(x) - \omega(y) \leq xy$ for all $x, y \in M$.
(wf2) For some point $x \in M$, $\sup_y \{xy - \omega(y)\} < \infty$.

In (wf2) we can require that the inequality holds for all $x \in M$, instead of for just one. Let $\mathcal{W}_M$ denote the set of all work functions on M. We will write simply $\mathcal{W}$ instead of $\mathcal{W}_M$ if M is understood from context.

For any $x \in M$, the function $\chi_x(y) = xy$ will be called the *characteristic function* of x, or the *cone* at x. Note that χ_x is a work function.

If M is a bounded metric space, (wf1) implies (wf2). Generally, condition (wf2) means that, in the sup-norm metric, ω is finitely far from any cone.

The update operator. For any $\omega \in \mathcal{W}$ and any non-empty set $r \subseteq M$, let $\omega \wedge r : M \to \mathbb{R}$ be defined by

$$\omega \wedge r(x) = \inf_{y \in r} \{\omega(y) + xy\}.$$

The operator $\wedge$ is called *update*. We say that a work function ω is *supported* by $r \subseteq M$ if $\omega \wedge r = \omega$.

Observation 0.2.3.1 *Let $\omega \in \mathcal{W}$ be a work function, and $r \subseteq M$ a nonempty set. Then*

(i) $\omega_{|r} = (\omega \wedge r)_{|r}$.
(ii) *If ω is supported by r, then* $\inf(\omega) = \inf(\omega \wedge r) = \inf(\omega)_{|r}$.

(iii) $\omega \wedge r$ *is supported by* r.

(iv) $\omega \wedge r \geq \omega$.

Lemma 0.2.3.2 *If ω is a work function, and $r \subseteq M$, then $\omega \wedge r$ is a work function.*

Proof: We verify conditions (wf1) and (wf2). Note that $\inf f - \inf g \leq \sup(f - g)$ for all real-valued functions f and g. Therefore

$$\begin{aligned} \omega \wedge r(x) - \omega \wedge r(y) &= \inf_{z \in r}\{\omega(z) + zx\} - \inf_{z \in r}\{\omega(z) + zy\} \\ &\leq \sup_{z \in r}\{\omega(z) + zx - \omega(z) - zy\} \;\leq\; xy, \end{aligned}$$

proving (wf1). Since $\omega \wedge r \geq \omega$, we have (wf2). □

The update operator extends naturally to arbitrary sequences of requests. Let $\omega \wedge \varepsilon = \omega$ (where ε = empty string) and recursively, for any $\varrho \in \mathcal{R}^*$ and $r \in \mathcal{R}$, let $\omega \wedge (\varrho r) = (\omega \wedge \varrho) \wedge r$.

Given $x \in M$ and $\varrho \in \mathcal{R}^*$, the optimal cost of serving ϱ starting from x can be expressed using work functions and the update operator, as stated in the following lemma.

Lemma 0.2.3.3 *Let $x^0 \in M$, $\varrho = r^1 \dots r^n \in \mathcal{R}^*$, and $\omega = \chi_{x^0} \wedge \varrho$. Then, for each $x \in r^n$,*

$$\omega(x) = \inf_{\sigma} |x^0 \sigma x|,$$

where the infimum is over all services σ for $r^1 \dots r^{n-1}$. Thus the optimal service cost is $opt(x^0, \varrho) = \inf(\omega)$.

Proof: For the sake of uniformity, define $r^0 = \{x^0\}$. Let $\omega^t = \chi_{x^0} \wedge r^1 \dots r^t$, for $t = 1, ..., n$. Thus $\omega = \omega^n$. Fix $x \in M$.

The proof is by induction on n. For $n = 0$, the result is trivial. Let $n > 0$. We will prove the "$\leq$" and "$\geq$" inequalities separately.

($\leq$) If $\sigma x = x^1 \dots x^n$ is any service for ϱ, where $x^n = x$, then, by the inductive hypothesis,

$$|x^0 \sigma x| \;=\; \sum_{t=1}^{n} x^{t-1} x^t \;=\; \sum_{t=1}^{n-1} x^{t-1} x^t + x^{n-1} x \;\geq\; \omega^{n-1}(x^{n-1}) + x^{n-1} x \;\geq\; \omega^n(x) = \omega(x).$$

($\geq$) For any fixed $\epsilon > 0$, we will construct a service $\sigma_\epsilon x = x^1 \dots x^n$, such that $|x^0 \sigma_\epsilon x| \leq \omega(x) + 2\epsilon$. Pick $x^{n-1} \in r^{n-1}$ such that

$$\omega(x) \;=\; \omega^n(x) \;>\; \omega^{n-1}(x^{n-1}) + x^{n-1} x - \epsilon.$$

By the inductive hypothesis, we can pick a service $x^1 \dots x^{n-1}$ for $r^1 \dots r^{n-1}$ such that

$$\sum_{t=1}^{n-1} x^{t-1} x^t \;<\; \omega^{n-1}(x^{n-1}) + \epsilon.$$

Thus $|x^0 \sigma_\epsilon x| = \sum_{t=1}^{n} x^{t-1} x^t < \omega(x) + 2\,\epsilon$. Let $\epsilon \to 0$, and we are done. □

0.3 Competitive Strategies

0.3.1 Potential Functions

Throughout the remainder of this section, let $\mathcal{S} = \langle M, \mathcal{R}\rangle$ be a metrical service system, and let $c > 0$ be a fixed real constant.

Solvent functions. Let $\mathcal{W} = \mathcal{W}_M$ be the set of work functions on $\mathcal{S}$. By $\textsc{Part}_{\mathcal{S}}$ we denote the set of partial functions $\phi : \mathcal{W} \times M \to \mathbb{R}$. For a given $\phi \in \textsc{Part}_{\mathcal{S}}$, let $\mathsf{Dom}\,\phi$ be the *domain* of ϕ, that is, the set of all $(\omega, x) \in \mathcal{W} \times M$, such that $\phi(\omega, x)$ is defined. If $\phi, \psi \in \textsc{Part}_{\mathcal{S}}$, we write $\phi \le \psi$ if $\mathsf{Dom}\,\psi \subseteq \mathsf{Dom}\,\phi$ and $\phi(\omega, x) \le \psi(\omega, x)$ for all $(\omega, x) \in \mathsf{Dom}\,\psi$. It is convenient to assume that $\phi(\omega, x) = \infty$ for $(\omega, x) \notin \mathsf{Dom}\,\phi$: note that this interpretation is consistent with the definition of the partial order "$\le$" on $\textsc{Part}_{\mathcal{S}}$. We can also write $\inf\phi$ to denote the infimum of ϕ on $\mathsf{Dom}\,\phi$, since it is the same as the global infimum of ϕ.

Let $\perp^c \in \mathbb{R}^{\mathcal{W}\times M}$ be the *bottom function*, defined by $\perp^c(\omega, x) = -c\,\inf(\omega)$. A function $\phi \in \textsc{Part}_{\mathcal{S}}$ is called *c-solvent* if $\phi \ge \perp^c$, or merely *solvent* if the value of c is understood from context. By $\textsc{Solv}^c_{\mathcal{S}}$ we denote the set of c-solvent functions on $\mathcal{S}$.

We use the word "solvent" for this property because we think of ϕ as representing a savings account for the server. The condition means that, should there be no more requests, the server has enough in this account to verify c-competitiveness.

The boost operator. The *boost operator* is a function $\mathcal{B} : \textsc{Part}_{\mathcal{S}} \to \textsc{Part}_{\mathcal{S}}$, defined as follows. For all $(\omega, x) \in \mathcal{W} \times M$,

$$\mathcal{B}\phi(\omega, x) \quad = \quad \sup_{r\in\mathcal{R}} \inf_{y\in r} \{xy + \phi(\omega\wedge r, y)\}.$$

We use notation $\mathcal{B}\phi(\omega, x)$ instead of more explicit $\mathcal{B}(\phi)(\omega, x)$, in order to simplify the notation. Note that the domain of $\mathcal{B}\phi$ may be different than the domain of ϕ.

Fact 0.3.1.1 *The boost operator has the following properties:*

(i) *If $\phi \in \textsc{Solv}^c_{\mathcal{S}}$ then $\mathcal{B}\phi \in \textsc{Solv}^c_{\mathcal{S}}$, for each $\phi \in \textsc{Part}_{\mathcal{S}}$.*

(ii) *$\mathcal{B}$ is monotone, that is, $\phi \le \psi$ implies that $\mathcal{B}\phi \le \mathcal{B}\psi$, for all $\phi, \psi \in \textsc{Part}_{\mathcal{S}}$.*

(iii) *$\mathcal{B}$ commutes with adding a constant, that is $\mathcal{B}(\phi + \lambda) = \mathcal{B}\phi + \lambda$ for all $\lambda \in \mathbb{R}$ and $\phi \in \textsc{Part}_{\mathcal{S}}$.*

Proof: Parts (ii) and(iii) follow directly from the properties of sup and inf operations, since both of them are monotone and commute with constants. To prove (i), note that for $(\omega, x) \in \mathcal{W} \times M$, by the solvency of ϕ,

$$\mathcal{B}\phi(\omega, x) \;\ge\; \sup_r \{-c\,\inf(\omega\wedge r)\} \;=\; \sup_r \left\{c\,\sup(-\omega)_{|r}\right\} \;=\; -c\,\inf(\omega).$$

In the last equality it is necessary to make use of the fact that $\bigcup_{r\in\mathcal{R}} r = M$. $\square$

The Lipschitz operator. The *c-Lipschitz operator* is a function $\mathcal{L}^c : \text{PART}_{\mathcal{S}} \to \mathbb{R}^{\mathcal{W}\times M}$, defined as follows. For any given $(\omega, x) \in \mathcal{W} \times M$,

$$\mathcal{L}^c\phi(\omega, x) \;=\; \inf_{(\theta,u)\in\mathcal{W}\times M} \{xu + c\,\sup(\theta - \omega) + \phi(\theta, u)\}.$$

Note that $\mathcal{L}^c\phi$ is always a total function. If the value of c is understood from context, we write $\mathcal{L}$ instead of $\mathcal{L}^c$, and *Lipschitz* instead of *c-Lipschitz.*

Fact 0.3.1.2 *The Lipschitz operator has the following properties:*

(i) *If* $\phi \in \text{SOLV}^c_{\mathcal{S}}$ *then* $\mathcal{L}\phi \in \text{SOLV}^c_{\mathcal{S}}$*, for each* $\phi \in \text{PART}_{\mathcal{S}}$*.*

(ii) $\mathcal{L}$ *is monotone, that is,* $\phi \leq \psi$ *implies that* $\mathcal{L}\phi \leq \mathcal{L}\psi$*, for all* $\phi, \psi \in \text{PART}_{\mathcal{S}}$*.*

(iii) $\mathcal{L}\phi \leq \phi$ *for each* $\phi \in \text{PART}_{\mathcal{S}}$*.*

Proof: Part (ii) follows from the monotonicity of sup and inf. Part (iii) follows from the definition of $\mathcal{L}$, by instantiating $\theta = \omega$ and $u = x$. In order to prove (i), note that for $(\omega, x) \in \mathcal{W} \times M$, by the solvency of ϕ,

$$\mathcal{L}\phi(\omega, x) \;\geq\; \inf_{\theta\in\mathcal{W}} \{c\,\sup(\theta - \omega) - c\,\sup(\theta)\} \;\geq\; -c\,\inf(\omega).$$

This completes the proof. $\square$

Lemma 0.3.1.3 $\mathcal{B}\mathcal{L}\phi \leq \mathcal{L}\mathcal{B}\phi$ *for each* $\phi \in \text{PART}_{\mathcal{S}}$*.*

Proof: For each $(\omega, x) \in \mathcal{W} \times M$,

$$\begin{aligned}
\mathcal{B}\mathcal{L}\phi(\omega, x) &= \sup_{r\in\mathcal{R}} \inf_{y\in r} \left\{ xy + \inf_{(\theta,u)} \{yu + c\,\sup(\theta - \omega\wedge r) + \phi(\theta, u)\} \right\} \\
&= \sup_{r\in\mathcal{R}} \inf_{y\in r} \inf_{(\theta,u)} \left\{ xy + yu + c\,\sup(\theta - \omega)_{|r} + \phi(\theta, u) \right\} \\
&= \inf_{(\eta,z)} \left\{ xz + c\,\sup(\eta - \omega) + \sup_{r\in\mathcal{R}} \inf_{y\in r} \inf_{(\theta,u)} \left\{ zy + yu + c\,\sup(\theta - \eta)_{|r} + \phi(\theta, u) \right\} \right\} \\
&\leq \inf_{(\eta,z)} \left\{ xz + c\,\sup(\eta - \omega) + \sup_{r\in\mathcal{R}} \inf_{y\in r} \{zy + \phi(\eta\wedge r, y)\} \right\} \\
&= \mathcal{L}\mathcal{B}\phi(\omega, x).
\end{aligned}$$

The last inequality is obtained by instantiating $\theta = \eta\wedge r$ and $u = y$. $\square$

We define $\phi \in \mathbb{R}^{\mathcal{W}\times M}$ to be *c-Lipschitz* if $\mathcal{L}^c\phi = \phi$. Thus, by Fact 0.3.1.2, (iii), ϕ is c-Lipschitz if and only if $\mathcal{L}\phi \geq \phi$, that is

$$\phi(\omega, x) - \phi(\theta, u) \;\leq\; xu + c\,\sup(\theta - \omega), \tag{0.1}$$

for all $(\omega, x),\ (\theta, u) \in \mathcal{W} \times M$. We will refer to (0.1) as the *c-Lipschitz property* for ϕ, or simply the as the *Lipschitz property* if c is understood.

Fact 0.3.1.4 *Let $\phi \in \mathrm{PART}_{\mathcal{S}}$. If ϕ is c-Lipschitz then*

(i) *$\mathcal{B}\phi$ is c-Lipschitz.*

(ii) *$\phi(\omega + a, x) = \phi(\omega, x) - c\,a$ for each $(\omega, x) \in \mathcal{W} \times M$ and $a \in \mathbb{R}$.*

Proof: (i) By Fact 0.3.1.2, (iii), it is sufficient to show $\mathcal{L}\mathcal{B}\phi \geq \mathcal{B}\phi$. By Lemma 0.3.1.3, we have $\mathcal{L}\mathcal{B}\phi \geq \mathcal{B}\mathcal{L}\phi = \mathcal{B}\phi$.

(ii) We only need to prove "$\geq$", since "$\leq$" follows by substituting $-a'$ for a and $\omega' + a'$ for ω. In inequality (0.1), let $(\theta, u) = (\omega + a, x)$. $\square$

Potential functions. Let $\phi \in \mathrm{PART}_{\mathcal{S}}$. We say that ϕ is a *c-potential* for $\mathcal{S}$, or simply a *c-potential* if $\mathcal{S}$ is understood, if the following conditions are satisfied:

(pot1) $\mathsf{Dom}\,\phi \neq \emptyset$,

(pot2) $\phi \in \mathrm{SOLV}^c_{\mathcal{S}}$,

(pot3) $\mathcal{B}\phi \leq \phi$.

Note that (pot3) implies that $\mathsf{Dom}\,\phi$ is closed under update in the following sense: if $(\omega, x) \in \mathsf{Dom}\,\phi$, then for each $r \in \mathcal{R}$ there is $y \in r$ such that $(\omega \wedge r, y) \in \mathsf{Dom}\,\phi$.

Fact 0.3.1.5 *Let $\phi \in \mathrm{PART}_{\mathcal{S}}$. If ϕ is a c-potential then $\mathcal{B}\phi$ and $\mathcal{L}\phi$ are c-potentials.*

Proof: Assume ϕ is a c-potential. Then $\mathcal{B}\phi \leq \phi$. Thus $\mathcal{B}\mathcal{B}\phi \leq \mathcal{B}\phi$ by the monotonicity of $\mathcal{B}$, proving that $\mathcal{B}\phi$ is a c-potential.

By the monotonicity of $\mathcal{L}$ and Lemma 0.3.1.3, $\mathcal{B}\mathcal{L}\phi \leq \mathcal{L}\mathcal{B}\phi \leq \mathcal{L}\phi$. $\square$

The minimum potential function. The minimum c-potential, if it exists, has properties useful in proving lower bounds on competitive constants. We prove these properties below.

Theorem 0.3.1.6 *Suppose that there exists a c-potential for $\mathcal{S}$. Then there exists a unique minimum c-potential $\phi^* = \phi^*_c$ for $\mathcal{S}$, and ϕ^* is equal to the minimum fixed point of $\mathcal{B}$ on $\mathrm{SOLV}^c_{\mathcal{S}}$. Furthermore, ϕ^* is Lipschitz.*

Proof: For every $(\omega, x) \in \mathcal{W} \times M$ define $\phi^*(\omega, x) = \inf_\phi \phi(\omega, x)$, where the infimum is over all c-potentials for $\mathcal{S}$. Clearly $\phi^* \geq \perp^c$. For any c-potential ϕ, by the monotonicity of the the boost operator, $\mathcal{B}\phi^* \leq \mathcal{B}\phi \leq \phi$. Thus $\mathcal{B}\phi^* \leq \phi^*$, proving that ϕ^* is a c-potential.

By Fact 0.3.1.5, $\mathcal{B}\phi^*$ is a c-potential. Thus, from the minimality of ϕ^*, we obtain $\mathcal{B}\phi^* \geq \phi^*$, proving that $\mathcal{B}\phi^* = \phi^*$. If $\psi \in \mathrm{SOLV}^c_{\mathcal{S}}$ is any fixed point of $\mathcal{B}$, that is, $\mathcal{B}\psi = \psi$, and if $\mathsf{Dom}\,\psi \neq \emptyset$, then ψ is also a c-potential and thus $\psi \geq \phi^*$. Therefore ϕ^* is the minimum fixed point of $\mathcal{B}$ on $\mathrm{SOLV}^c_{\mathcal{S}}$.

By Fact 0.3.1.2 (iii), $\mathcal{L}\phi^* \leq \phi^*$. By Fact 0.3.1.5, and the minimality of ϕ^*, $\mathcal{L}\phi^* \geq \phi^*$. Thus $\mathcal{L}\phi^* = \phi^*$, *i.e.*, ϕ^* is Lipschitz. $\square$

From the Lipschitz property of ϕ^* and Fact 0.3.1.4, (ii), we have:

Fact 0.3.1.7 *Let ϕ^* denote the minimum c-potential. For each $(\omega, x) \in \mathcal{W} \times M$ and $a \in \mathbb{R}$ we have $\phi^*(\omega + a, x) = \phi^*(\omega, x) - c\,a$.*

Similar metrical service systems. Suppose that we are given two metrical service systems $\mathcal{S} = \langle M, \mathcal{R} \rangle$ and $\mathcal{S}' = \langle M', \mathcal{R}' \rangle$. Let $F : M \to M'$ be a function, and let $\lambda > 0$. We say that F is a λ-*similarity* from $\mathcal{S}$ to $\mathcal{S}'$ if F is one-to-one and onto, $F(x)F(y) = \lambda\, xy$ for all $x, y \in M$, and $\mathcal{R}' = \{F(r) : r \in \mathcal{R}\}$. Note that then F^{-1} is a λ^{-1}-similarity from $\mathcal{S}'$ to $\mathcal{S}$.

If $\lambda = 1$, we say that F is an *isomorphism*. If $\mathcal{S} = \mathcal{S}'$, we say that F is a λ-*symmetry*. Define a *symmetry* to be a 1-symmetry.

Given a λ-similarity F from $\mathcal{S}$ to $\mathcal{S}'$, we overload notation by letting F denote the function from any structure associated with $\mathcal{S}$ to the corresponding structure for $\mathcal{S}'$. Write $\mathcal{W}' = \mathcal{W}_{M'}$. If $\omega \in \mathcal{W}$, let $F(\omega) \in \mathcal{W}'$ be defined by $F(\omega)(F(x)) = \lambda\, \omega(x)$. If $\phi \in \text{PART}_{\mathcal{S}}$, let $F(\phi)$ be defined by $F(\phi)(F(\omega), F(x)) = \lambda\, \phi(\omega, x)$.

Similarity preserves all relevant structure, *i.e.*, any statement about $\mathcal{S}$ is equivalent to the corresponding statement about $\mathcal{S}'$.

Fact 0.3.1.8 *Let F be a λ-similarity from $\mathcal{S}$ to $\mathcal{S}'$, for some $\lambda > 0$. Then*

(i) *If $\phi \in \text{PART}_{\mathcal{S}}$, then $\mathcal{B}F(\phi) = F(\mathcal{B}\phi)$. Thus, ϕ is a c-potential for $\mathcal{S}$ if and only if $F(\phi)$ is a c-potential for $\mathcal{S}'$.*
(ii) *$\mathcal{L}^c F(\phi) = F(\mathcal{L}^c \phi)$. Thus, ϕ is c-Lipschitz if and only if $F(\phi)$ is c-Lipschitz.*
(iii) *If $\phi \in \text{PART}_{\mathcal{S}}$, then ϕ is a c-potential iff $F(\phi)$ is a c-potential. Furthermore, $F(\phi_c^*) = \psi_c^*$, where ψ_c^* is the minimum c-potential for $\mathcal{S}'$. In particular, if $\mathcal{S} = \mathcal{S}'$, then ϕ_c^* is a fixed point of F.*

0.3.2 The Minimum Potential in a Compactly Shielded System

Throughout this subsection, let $\mathcal{S} = \langle M, \mathcal{R} \rangle$ be a fixed MSS, and let $c > 0$. Fix a function $\varphi^0 \in \mathbb{R}^{\mathcal{W} \times M}$ and for each integer $n \geq 0$ let $\varphi^n = \mathcal{B}^n \varphi$. We will study properties of the sequence $\{\varphi^n\}$.

Fact 0.3.2.1 *Assume that $\mathcal{B}\varphi^0 \geq \varphi^0$. Then the sequence $\{\varphi^n\}$ has the following properties:*

(i) *$\varphi^{n+1} \geq \varphi^n$, for each n.*
(ii) *If φ^0 is solvent then each φ^n is solvent.*
(iii) *If φ^0 is c-Lipschitz then each φ^n is c-Lipschitz.*
(iv) *If $\{\varphi^n\}$ converges (pointwise), and $\varphi^* = \lim_n \varphi^n$, then $\mathcal{B}\varphi^* \geq \varphi^*$.*

Proof: Part (i) follows from the monotonicity of $\mathcal{B}$: Since $\mathcal{B}$ is monotone, $\mathcal{B}^n$ is also monotone. Thus $\varphi^{n+1} = \mathcal{B}^n \mathcal{B}\varphi^0 \geq \mathcal{B}^n \varphi^0 = \varphi^n$. Part (ii) follows directly from (i). Part (iii) follows from

Lemma 0.3.1.3: $\mathcal{L}\varphi^n = \mathcal{L}\mathcal{B}^n\varphi^0 \geq \mathcal{B}^n\mathcal{L}\varphi^0 = \mathcal{B}^n\varphi^0 = \varphi^n$, so φ^n is c-Lipschitz. In order to prove (iv), notice that, by monotonicity, $\mathcal{B}\varphi^* \geq \mathcal{B}\varphi^{n-1} = \varphi^n$ for each n. Passing to the limit, we obtain $\mathcal{B}\varphi^* \geq \varphi^*$. □

Lemma 0.3.2.2 *Assume that $\mathcal{S}$ is compactly shielded. Suppose φ is c-Lipschitz, $\mathcal{B}\varphi^0 \geq \varphi^0$, and that the pointwise limit $\varphi^* = \lim_n \varphi^n$ exists. Then $\mathcal{B}\varphi^* = \varphi^*$, that is, φ^* is a fixed point of $\mathcal{B}$.*

Proof: That $\mathcal{B}\varphi^* \geq \varphi^*$ follows from Fact 0.3.2.1. The remainder of the proof verifies the more difficult inequality that $\mathcal{B}\varphi^* \leq \varphi^*$. Let $(\omega, x) \in \mathcal{W} \times M$. Then

$$\begin{aligned}
\varphi^*(\omega,x) &= \sup_n \varphi^{n+1}(\omega,x) \;=\; \sup_n \sup_{r\in\mathcal{R}} \inf_{y\in r}\{\varphi^n(\omega\wedge r,y)+xy\} \\
&= \sup_{r\in\mathcal{R}} \sup_n \inf_{y\in r}\{\varphi^n(\omega\wedge r,y)+xy\}, \text{ and} \\
\mathcal{B}\varphi^*(\omega,x) &= \sup_{r\in\mathcal{R}} \inf_{y\in r}\{\varphi^*(\omega\wedge r,y)+xy\} \;=\; \sup_{r\in\mathcal{R}} \inf_{y\in r} \sup_n\{\varphi^n(\omega\wedge r,y)+xy\}.
\end{aligned}$$

Fix an arbitrary request $r \in \mathcal{R}$. It is sufficient to show that

$$\sup_n \inf_{y\in r}\{\varphi^n(\omega\wedge r,y)+xy\} \;\geq\; \inf_{y\in r} \sup_n\{\varphi^n(\omega\wedge r,y)+xy\}. \tag{0.2}$$

Let $Z = r \cap B$, where B is the ball of radius $b = c\ \inf(\omega\wedge r) + \inf_{y\in r}\{\varphi^*(\omega\wedge r,y)+xy\}+1$, centered at x. We need the following claim.

Claim A:

$$\inf_{y\in r}\{\varphi^n(\omega\wedge r,y)+xy\} \;=\; \inf_{y\in Z}\{\varphi^n(\omega\wedge r,y)+xy\}.$$

We need only prove

$$\inf_{y\in r-Z}\{\varphi^n(\omega\wedge r,y)+xy\} \;>\; \inf_{y\in Z}\{\varphi^n(\omega\wedge r,y)+xy\}.$$

If $z \in r - Z$, then $zx \geq b$, and $\varphi^n(\omega\wedge r, z) + xz \;\geq\; -c\ \inf(\omega\wedge r) + b \;\geq\; \inf_{y\in r}\{\varphi^n(\omega\wedge r,y)+xy\}+1$. This verifies Claim A.

We now continue the proof of inequality (0.2). Pick a compact set $Y \subseteq r$ which shields Z from $\{x\}$. For each $z \in Z$ pick $y_z \in Y$ such that $xz = xy_z + y_z z$. For each fixed n, since φ^n is c-Lipschitz, we have

$$\inf_{z\in Z}\{\varphi^n(\omega\wedge r,z)+xz\} \;=\; \inf_{z\in Z}\{\varphi^n(\omega\wedge r,z)+xy_z+y_z z\} \;\geq\; \inf_{y\in Y}\{\varphi^n(\omega\wedge r,y)+xy\}.$$

Using this last inequality, Claim A, and applying Lemma 0.2.2.2 to Y, we have

$$\begin{aligned}
\sup_n \inf_{z\in r}\{\varphi^n(\omega\wedge r,z)+xz\} &\geq \sup_n \inf_{z\in Y}\{\varphi^n(\omega\wedge r,z)+xz\} \;=\; \inf_{z\in Y} \sup_n\{\varphi^n(\omega\wedge r,z)+xz\} \\
&\geq \inf_{z\in r} \sup_n\{\varphi^n(\omega\wedge r,z)+xz\},
\end{aligned}$$

completing the proof of inequality (0.2) and the theorem. □

Now define $\psi^n = \mathcal{B}^n \perp^c$ for each $n \geq 0$. In the rest of this section we will study the convergence properties of the sequence $\{\psi^n\}$. From Fact 0.3.2.1, we have the following:

Fact 0.3.2.3 *The sequence $\{\psi^n\}$ is monotone increasing, and each ψ^n is c-Lipschitz and solvent.*

Proof: By Fact 0.3.2.1, (i), $\psi^{n+1} \geq \psi^n$, for each n. By Fact 0.3.2.1, (ii) and (iii), applied to $\varphi^0 = \perp^c$, it suffices to show that $\perp^c$ is c-Lipschitz and solvent. That $\perp^c$ is obviously solvent. By Fact 0.3.1.2, $\mathcal{L}\perp^c \leq \perp^c$, and $\mathcal{L}\perp^c$ is solvent, and therefore $\mathcal{L}\perp^c = \perp^c$, by the minimality of $\perp^c$ in $\textsc{Solv}^c_{\mathcal{S}}$. Thus $\perp^c$ is Lipschitz. □

Theorem 0.3.2.4 *Assume that $\mathcal{S}$ is compactly shielded. Supposed that $\{\psi^n\}$ converges, and let $\psi^* = \lim_n \psi^n$. Then $\psi^* = \phi^*$, the minimum c-potential for $\mathcal{S}$.*

Proof: By monotonicity of $\{\psi^n\}$, ψ^* is solvent. By Lemma 0.3.2.2, ψ^* is a fixed point of $\mathcal{B}$, hence ψ^* is a c-potential. Thus the minimal potential ϕ^* exists. Since ϕ^* is solvent, $\phi^* \geq \psi^0$, and by monotonicity of $\mathcal{B}$, $\psi^n = \mathcal{B}^n\perp^c \leq \mathcal{B}^n\phi^* = \phi^*$, for each n. Therefore, passing to the limit, $\psi^* \leq \phi^*$. The minimality of ϕ^* implies that $\phi^* = \psi^*$. □

As the example below shows, Theorem 0.3.2.4 does not necessarily hold if the given MMS is not compactly shielded.

Example 1: Let $\mathcal{S}_1 = \langle M_1, \mathcal{R}_1\rangle$, where M_1 and $\mathcal{R}_1$ are defined as follows. M_1 consists of all points x_{ij}, where $i \geq j \geq 0$, and the metric function is defined by

$$x_{ij}x_{kl} = \begin{cases} 2 & \text{if } |l-j| = 1 \text{ and } i \neq k, \\ 1 & \text{otherwise.} \end{cases}$$

The set of requests is $\mathcal{R}_1 = \{r_j\}_{j\geq 0}$, where $r_j = \{x_{ij} | i \geq j\}$. For this MSS $\mathcal{S}_1$, and for $c = 1$, $\{\psi^n\}$ converges, but $\mathcal{B}\psi^* \neq \psi^*$. Consequently, ψ^* is not the minimum potential. In fact, this MSS does not have a 1-potential at all. The verification of this fact is left to the reader. ◇

The proof the the next theorem is similar to the one of Theorem 0.3.2.4, and is left to the reader.

Theorem 0.3.2.5 *Suppose that $\mathcal{S}$ is compactly shielded and has a c-potential. Then $\{\psi^n = \mathcal{B}^n\perp^c\}$ converges to the minimum potential ϕ^*.*

Like the previous theorem, Theorem 0.3.2.5 does not hold in general, as shown in the example below.

Example 2: We modify Example 1 as follows: Let $\mathcal{S}_2 = \langle M_2, \mathcal{R}_2\rangle$, where $M_2 = M_1 \cup \{x_0\}$, and $x_0x = 2$ for all $x \in M_1$. We also define $\mathcal{R}_2 = \{r \cup \{x_0\} : r \in \mathcal{R}_1\}$. Thus we add one point to M_1 and to all requests; this point is at distance 2 from all points in M_1. This MSS $\mathcal{S}_2$ has a 1-potential but, similarly as in Example 1, $\{\psi^n\}$ does not converge. ◇

Corollary 0.3.2.6 *Let M be a finite metric space with integer distances, and $\mathcal{S} = \langle M, \mathcal{R}\rangle$. $\mathcal{S}$ has a c-potential if and only if $\{\psi^n\}$ converges in a finite number of steps. Furthermore, if $\{\psi^n\}$ converges then its limit is ϕ^*, the minimum potential.*

Application. The above corollary is very useful in hand- and computer calculations. Given a finite metrical service system $\mathcal{S}$ with integer distances, we can check whether it is c-competitive by computing the minimum potential.

The algorithm for computing the minimum potential is as follows: First, note that we cannot consider all work functions, because there are infinitely many of those. Instead, we let $\mathcal{V} \subseteq \mathcal{W}_M$ be the set of all *offset functions*, defined to be those work functions whose minima are zero. Let $\xi^n = \psi^n_{|\mathcal{V}\times M}$, the function ψ^n restricted to $\mathcal{V} \times M$. We use the property that, for each given n, $\psi^n(\omega + a, x) = \psi^n(\omega, x) - ca$, for all $a \in \mathbb{R}$, $\omega \in \mathcal{W}_M$ and $x \in M$, which follows from Fact 0.3.1.4 (ii). Thus, in particular, for $(\omega, x) \in \mathcal{V} \times M$

$$\xi^{n+1}(\omega, x) \;=\; \sup_r \inf_{y\in r} \{\psi^n(\omega\wedge r, y) + xy\} \;=\; \sup_r \inf_{y\in r} \{\xi^n(\omega\triangle r, y) + xy - c\,\inf(\omega\wedge r)\}$$

where $\omega\triangle r = \omega\wedge r - \inf(\omega\wedge r)$, so that $\omega\triangle r \in \mathcal{V}$.

Initialize $\xi^0 = 0$. Compute ξ^{n+1} from ξ^n using the above formula. Note that all values of ξ^n will be non-negative integers. We halt the iteration when one of two things happens, either $\xi^{n+1} = \xi^n$, which implies that ξ^n is the minimum potential, and therefore the problem is c-competitive, or $\xi^n > 0$ everywhere, which implies that the sequence diverges (because $\min(\phi^*_{|\mathcal{V}\times M}) = 0$ if ϕ^* exists), and hence $\mathcal{S}$ is not c-competitive. Since there are only finitely many offset functions with integer values, one of these two events must occur eventually.

0.3.3 History Based Strategies

An on-line strategy $\mathcal{H}$ for a metrical service system $\mathcal{S}$ can be viewed as a function $\mathcal{H} : M \times \mathcal{R}^* \to M$, with the following interpretation. For given $x \in M$ and $\varrho \in \mathcal{R}^*$, the point $\mathcal{H}(x, \varrho) \in M$ is the position of $\mathcal{H}$'s server after serving requests from ϱ, starting from the initial position x. If $\varrho = \varepsilon$, we assume that $\mathcal{H}(x, \varepsilon) = x$. This definition captures the on-line nature of the problem, since if the current server position is $y = \mathcal{H}(x, \varrho)$, and r is the new request, then the new server position $\mathcal{H}(x, \varrho r)$ is determined only by the past requests. We will refer to an on-line strategy defined as above as a *history-based strategy*, to distinguish it from *work function based* strategies introduced in the next subsection.

Given a history-based strategy $\mathcal{H}$, a point $x \in M$ and a request sequence $\varrho \in \mathcal{R}^*$, we define the *cost of $\mathcal{H}$ on ϱ starting from x* as follows: Let $\varrho = r^1 \dots r^n$ and $x^t = \mathcal{H}(x, r^1 \dots r^t)$, for $t = 0, \dots, n$. (Thus $x = x^0$.) Then

$$cost_{\mathcal{H}}(x, \varrho) \;=\; \sum_{t=1}^{n} x^{t-1}x^t.$$

A history-based strategy $\mathcal{H}$ for $\mathcal{S}$ is said to be *c-competitive* if

$$\sup_{\varrho} \{ cost_{\mathcal{H}}(x, \varrho) - c\, opt(x, \varrho) \} \;<\; \infty,$$

for all $x \in M$. We also say that a strategy $\mathcal{H}$ is *almost c-competitive* if $\mathcal{H}$ is $(c + \epsilon)$-competitive for every $\epsilon > 0$. We will call $\mathcal{S}$ c-competitive (almost c-competitive) if it has a c-competitive (respectively, almost c-competitive) history-based strategy.

0.3.4 Work Function Based Strategies

Quasi-total functions. A partial function $\mathcal{A} : \mathcal{W} \times M \times \mathcal{R} \to M$ is called *quasi-total*, if $\mathcal{A}(\omega, x, r)$ is defined implies $\mathcal{A}(\omega, x, q)$ is defined, for all $(\omega, x) \in \mathcal{W} \times M$, and $r, q \in \mathcal{R}$. We overload notation by defining $\mathsf{Dom}\, \mathcal{A} \subseteq \mathcal{W} \times M$ to be the set of pairs (ω, x) for which $\mathcal{A}(\omega, x, r)$ is defined for all r. Note that this definition of $\mathsf{Dom}\, \mathcal{A}$ is meaningful only for quasi-total functions.

Work function based strategies. A quasi-total function $\mathcal{A}$ is called a *work-function-based strategy* for $\mathcal{S}$ (*wf-based* for short) if $\mathsf{Dom}\, \mathcal{A} \neq \emptyset$, and if $(\omega, x) \in \mathsf{Dom}\, \mathcal{A}$ and $y = \mathcal{A}(\omega, x, r)$ implies $y \in r$ and $(\omega \wedge r, y) \in \mathsf{Dom}\, \mathcal{A}$.

We extend $\mathcal{A}$ to a partial function $\mathcal{A}^* : \mathcal{W} \times M \times \mathcal{R}^* \to M$ inductively as follows: $\mathcal{A}^*(\omega, x, \varepsilon) = x$, and

$$\mathcal{A}^*(\omega, x, \varrho r) \;=\; \mathcal{A}(\, \omega \wedge \varrho \,,\, \mathcal{A}^*(\omega, x, \varrho) \,,\, r \,),$$

for all $(\omega, x) \in \mathsf{Dom}\, \mathcal{A}$, $\varrho \in \mathcal{R}^*$ and $r \in \mathcal{R}$.

Given a wf-based strategy $\mathcal{A}$, $(\omega, x) \in \mathsf{Dom}\, \mathcal{A}$ and $\varrho \in \mathcal{R}^*$, we define the *cost of $\mathcal{A}$ on ϱ starting from* (ω, x) as follows: Let $\varrho = r^1 \ldots r^n$ and $x^t = \mathcal{A}^*(\omega, x, r^1 \ldots r^t)$, for $t = 0, \ldots, n$. (Thus $x^0 = x$.) Then

$$cost_{\mathcal{A}}(\omega, x, \varrho) \;=\; \sum_{t=1}^{n} x^{t-1} x^t.$$

We say that $\mathcal{A}$ is *c-competitive*, if

$$\xi_{\mathcal{A}}(\omega, x) \;=\; \sup \{ cost_{\mathcal{A}}(\omega, x, \varrho) - c \, \inf(\omega \wedge \varrho) \} \;<\; \infty$$

for all $(\omega, x) \in \mathsf{Dom}\, \mathcal{A}$.

We say that $\phi \in \textsc{Solv}^c_{\mathcal{S}}$ is a *c-potential for* $\mathcal{A}$ if,

$$xy + \phi(\omega \wedge r, y) \leq \phi(\omega, x),$$

for all $(\omega, x) \in \mathsf{Dom}\, \mathcal{A}$, $r \in \mathcal{R}$ and $y = \mathcal{A}(\omega, x, r)$. We will sometimes write this inequality as

$$\Delta cost_{\mathcal{A}} + \Delta \phi \leq 0,$$

where $\Delta cost_{\mathcal{A}} = xy$ is the cost of $\mathcal{A}$ in the current move, and $\Delta\phi = \phi(\omega\wedge r, y) - \phi(\omega, x)$ is the potential change.

Theorem 0.3.4.1 *Let $\mathcal{A}$ be a work function based strategy. Then $\mathcal{A}$ is c-competitive iff $\mathcal{A}$ has a c-potential. More specifically, if $\mathcal{A}$ is c-competitive, then $\xi_{\mathcal{A}}$ is a c-potential for $\mathcal{A}$. Furthermore, if $\mathcal{A}$ has a c-potential ϕ, then $\xi_{\mathcal{A}} \leq \phi$.*

Proof: ($\Rightarrow$) Suppose that $\mathcal{A}$ is c-competitive. By the definition of competitiveness of $\mathcal{A}$, $\xi_{\mathcal{A}}$ is defined on $\mathsf{Dom}\,\mathcal{A}$. If we take $\varrho = \varepsilon$ in the above definition, we get $\xi_{\mathcal{A}}(\omega, x) \geq -c\inf(\omega)$, thus $\xi_{\mathcal{A}}$ is c-solvent. It remains to show that the condition $\Delta cost_{\mathcal{A}} + \Delta\xi_{\mathcal{A}} \leq 0$ holds. This can be derived as follows: Let $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, $\varrho \in \mathcal{R}^*$ and $r \in \mathcal{R}$. Let $y = \mathcal{A}(\omega, x, r)$. Then

$$\begin{aligned}\xi_{\mathcal{A}}(\omega, x) &\geq \sup_{\varrho}\{cost_{\mathcal{A}}(\omega, x, r\varrho) - c\inf(\omega\wedge r\varrho)\} \\ &= xy + \sup_{\varrho}\{cost_{\mathcal{A}}(\omega\wedge r, y, \varrho) - c\inf[(\omega\wedge r)\wedge\varrho]\} \;\geq\; xy + \xi_{\mathcal{A}}(\omega\wedge r, y),\end{aligned}$$

completing the proof of ($\Rightarrow$).

($\Leftarrow$) Let ϕ be a c-potential for $\mathcal{A}$, and fix $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, $\varrho \in \mathcal{R}^*$. Let $y = \mathcal{A}^*(\omega, x, \varrho)$. Then, by the summation of the inequalities $\Delta cost_{\mathcal{A}} + \Delta\phi \leq 0$ over the request sequence ϱ, we get

$$cost_{\mathcal{A}}(\omega, x, \varrho) \;\leq\; -\phi(\omega\wedge\varrho, y) + \phi(\omega, x) \;\leq\; c\inf(\omega\wedge\varrho) + \phi(\omega, x),$$

and therefore

$$\xi_{\mathcal{A}}(\omega, x) \;=\; \sup_{\varrho}\{cost_{\mathcal{A}}(\omega, x, \varrho) - c\inf(\omega\wedge\varrho)\} \;\leq\; \phi(\omega, x),$$

and thus $\mathcal{A}$ is c-competitive. □

Time-dependent strategies. Let $\mathbb{N}$ be the set of non-negative integers. We extend the definition of quasi-total functions to functions $\mathcal{A} : \mathcal{W} \times M \times \mathbb{N} \times \mathcal{R} \to M$ as follows. We say that $\mathcal{A}$ is *quasi-total* if $\mathcal{A}(\omega, x, t, r)$ is defined implies $\mathcal{A}(\omega, x, s, q)$ is defined, for all $(\omega, x) \in \mathcal{W} \times M$, $t, s \in \mathbb{N}$, and $r, q \in \mathcal{R}$. By the domain of $\mathcal{A}$, denoted $\mathsf{Dom}\,\mathcal{A}$, we denote the set of all (ω, x) for which $\mathcal{A}(\omega, x, t, r)$ is defined for all t and r.

We define a *time-dependent wf-based strategy* for $\mathcal{S}$ to be a quasi-total function

$$\mathcal{A} : \mathcal{W} \times M \times \mathbb{N} \times \mathcal{R} \to M$$

such that $\mathsf{Dom}\,\mathcal{A} \neq \emptyset$, and for all $t \in \mathbb{N}$ and $r \in \mathcal{R}$, $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$ implies that $(\omega\wedge r, y) \in \mathsf{Dom}\,\mathcal{A}$ for $y = \mathcal{A}(\omega, x, t, r)$.

Given a time-dependent wf-based strategy $\mathcal{A}$, we define a partial function $\mathcal{A}^* : \mathcal{W} \times M \times \mathcal{R}^* \to M$, as follows: $\mathcal{A}^*(\omega, x, \varepsilon) = x$, and

$$\mathcal{A}^*(\omega, x, \varrho r) \;=\; \mathcal{A}(\,\omega\wedge\varrho\,,\, \mathcal{A}^*(\omega, x, \varrho)\,,\, |\varrho|\,,\, r\,),$$

for all $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, $\varrho \in \mathcal{R}^*$, $r \in \mathcal{R}$.

We extend the definition of the cost function and competitiveness to time-dependent strategies in a straightforward way: Given a time-dependent wf-based strategy $\mathcal{A}$, $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, and $\varrho \in \mathcal{R}^*$, we define the *cost of* $\mathcal{A}$ *on* ϱ *starting from* (ω, x) as follows: Let $\varrho = r^1 \ldots r^n$ and $x^t = \mathcal{A}^*(\omega, x, r^1 \ldots r^t)$, for $t = 0, \ldots, n$. Then

$$cost_{\mathcal{A}}(\omega, x, \varrho) \;=\; \sum_{t=1}^{n} x^{t-1} x^t.$$

We say that $\mathcal{A}$ is *c-competitive*, if

$$\sup_{\varrho \in \mathcal{R}^*} \{cost_{\mathcal{A}}(\omega, x, \varrho) - c \inf(\omega \wedge \varrho)\} \;<\; \infty,$$

for all $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$.

Theorem 0.3.4.2 *The following conditions are equivalent:*

(i) *$\mathcal{S}$ has a c-competitive time-dependent wf-based strategy.*
(ii) *$\mathcal{S}$ is c-competitive, that is, it has a c-competitive history based strategy.*
(iii) *$\mathcal{S}$ has a c-potential.*

Proof: (i) $\Rightarrow$ (ii) Let $\mathcal{A}$ be a time-dependent wf-based strategy that is c-competitive. We will construct a history-based strategy $\mathcal{H}$ as follows. Pick an arbitrary $(\omega^0, x^0) \in \mathsf{Dom}\,\mathcal{A}$. For any $x \in M$, let $\mathcal{H}(x, \varepsilon) = x$ and $\mathcal{H}(x, \varrho) = \mathcal{A}^*(\omega^0, x^0, \varrho)$ if $\varrho \neq \varepsilon$.

For fixed $x \in M$, let $a_x = \sup(\omega^0 - \chi_x)$ and $b_x = \sup_{\varrho \in \mathcal{R}^*} \{cost_{\mathcal{A}}(\omega, x, \varrho) - c \inf(\omega \wedge \varrho)\}$. Then

$$cost_{\mathcal{H}}(x, \varrho) \;\leq\; cost_{\mathcal{A}}(\omega^0, x^0, \varrho) + xx_0 \;\leq\; c\, opt(x, \varrho) + c\, a_x + b_x + xx^0$$

establishing that $\mathcal{H}$ is c-competitive.

(ii) $\Rightarrow$ (iii) Let $\mathcal{H}$ be a c-competitive history-based strategy. For $x \in M$ define

$$\xi(x) \;=\; \sup_{\varrho \in \mathcal{R}^*} \{cost_{\mathcal{H}}(x, \varrho) - c\, opt(x, \varrho)\}.$$

Since $\mathcal{H}$ is c-competitive, $\xi(x) < \infty$ for all $x \in M$. We define $\phi \in \mathrm{P{\scriptstyle ART}}_{\mathcal{S}}$ such that for $(\omega, y) \in \mathcal{W} \times M$

$$\phi(\omega, y) \;=\; \inf_{x \in M} \inf_{\varrho \in \mathcal{R}^*} \{\xi(x) - cost_{\mathcal{H}}(x, \varrho) \;:\; \chi_x \wedge \varrho = \omega \text{ and } \mathcal{H}(x, \varrho) = y\}$$

provided an appropriate pair (x, ϱ) exists. Since $(x, \chi_x) \in \mathsf{Dom}\,\phi$, for each $x \in M$, we have $\mathsf{Dom}\,\phi \neq \emptyset$. Thus, in order to prove that ϕ is a c-potential, we need to verify properties (pot1) and (pot2). Fix $(\omega, y) \in \mathsf{Dom}\,\phi$. Let $\epsilon > 0$. Pick (x, ϱ) such that $\chi_x \wedge \varrho = \omega$, $\mathcal{H}(x, \varrho) = y$, and $\phi(\omega, y) > \xi(x) - cost_{\mathcal{H}}(x, \varrho) - \epsilon$. Then

$$\phi(\omega, y) \;>\; \xi(x) - cost_{\mathcal{H}}(x, \varrho) - \epsilon \;\geq\; -c\, opt(x, \varrho) - \epsilon \;=\; -c \inf(\omega) - \epsilon$$

Letting $\epsilon \to 0$ we conclude that ϕ is solvent.

Given an arbitrary $r \in \mathcal{R}$, pick $z = \mathcal{H}(x, \varrho r)$. Then, by the choice of ϵ and (x, ϱ),

$$\phi(\omega \wedge r, z) + yz \;\le\; \xi(x) - cost_{\mathcal{H}}(x, \varrho r) + yz \;=\; \xi(x) - cost_{\mathcal{H}}(x, \varrho) \;<\; \phi(\omega, y) + \epsilon.$$

Letting $\epsilon \to 0$, we obtain (pot3), completing the proof of (ii) $\Rightarrow$ (iii).

(iii) $\Rightarrow$ (i) Let ϕ be a c-potential for $\mathcal{S}$, $(\omega, x) \in \mathsf{Dom}\,\phi$, and $r \in \mathcal{R}$. Since $\mathcal{B}\phi \le \phi$, we have

$$\inf_{y \in r} \{xy + \phi(\omega \wedge r, y)\} \;\le\; \phi(\omega, x).$$

We define a time-dependent wf-based strategy $\mathcal{A}$ with $\mathsf{Dom}\,\mathcal{A} = \mathsf{Dom}\,\phi$, as follows: For any given $t \in \mathbf{N}$, choose $\mathcal{A}(\omega, x, t, r) = y$ such that $xy + \phi(\omega \wedge r, y) < \phi(\omega, x) + 2^{-t}$.

Now fix $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$. Given $\varrho = r^1 \ldots r^n \in \mathcal{R}^*$, let $x^t = \mathcal{A}^*(\omega, x, r^1 \ldots r^t)$ and $\omega^t = \omega \wedge r^1 \ldots r^t$, for $t = 0, \ldots, n$. Then

$$\begin{aligned} cost_{\mathcal{A}}(\omega, x, \varrho) \;&=\; \sum_{t=0}^{n-1} x^t x^{t+1} \;<\; \sum_{t=0}^{n-1} \left[-\phi(\omega^{t+1}, x^{t+1}) + \phi(\omega^t, x^t) + 2^{-t}\right] \\ &<\; -\phi(\omega^n, x^n) + \phi(\omega^0, x^0) + 1 \;\le\; c \inf(\omega \wedge \varrho) + \phi(\omega, x) + 1, \end{aligned}$$

verifying competitiveness of $\mathcal{A}$. □

0.3.5 Competitive Strategies in Compactly Shielded Systems

As we have shown in Theorem 0.3.4.2, the existence of a c-competitive history-based strategy $\mathcal{H}$ is equivalent to the existence of a time-dependent wf-based c-competitive strategy $\mathcal{A}$. Strategy $\mathcal{A}$ keeps track only of the current work function and "time", that is the length of the past request sequence. It does not need to remember the request sequence itself. This time dependence is undesirable, it would be more convenient to deal only with work-function based strategies. Fortunately, time-dependence can be eliminated if the system is compactly shielded.

Lemma 0.3.5.1 *Let $\mathcal{S}$ be a compactly shielded metrical service system. If $\mathcal{S}$ has a c-potential, then it has a c-competitive work-function-based strategy.*

Proof: Let $\phi^* : \mathcal{W} \times M \to \mathbb{R}$ be the minimum c-potential for $\mathcal{S}$. We show that ϕ^* is a c-potential for some wf-based strategy $\mathcal{A}$.

Fix $(\omega, x) \in \mathcal{W} \times M$ and $r \in \mathcal{R}$. For convenience, for $y \in M$, write $f(y) = xy + \phi^*(\omega \wedge r, y)$. We have $\inf_{y \in r} f(y) \le \mathcal{B}\phi^*(\omega, x) = \phi^*(\omega, x)$. Let $Z = \{y \in r \;:\; f(y) \le \phi^*(\omega, x) + 1\}$. Then $Z \ne \emptyset$, and Z is bounded since

$$xz \;=\; f(z) - \phi^*(\omega \wedge r, z) \;\le\; \phi^*(\omega, x) + 1 - \phi^*(\omega \wedge r, z) \;\le\; \phi^*(\omega, x) + 1 + c \inf(\omega \wedge r),$$

for any $z \in Z$. Pick a compact set Y which shields Z from $\{x\}$. By the shielding property, for each $z \in Z$ there is a $y_z \in Y$ such that $xy_z + y_z z = xz$. Then, by the choice of Z and by the Lipschitz property of ϕ^*, we have

$$\inf_{z \in r} f(z) \;=\; \inf_{z \in Z} f(z) \;=\; \inf_{z \in Z} \{xy_z + y_z z + \phi^*(\omega \wedge r, z)\} \;\geq\; \inf_{z \in Z} \{xy_z + \phi^*(\omega \wedge r, y_z)\} \;\geq\; \inf_{y \in Y} f(y).$$

Every continuous function on a compact set has a minimum. Pick $y \in Y$ which minimizes $f(y)$ on Y and define $\mathcal{A}(\omega, x, r) = y$. By the inequality above,

$$xy + \phi^*(\omega \wedge r, y) \;=\; \inf_{z \in r} \{xz + \phi^*(\omega \wedge r, z)\} \;\leq\; \phi^*(\omega, x).$$

Thus ϕ^* is a c-potential for $\mathcal{A}$, hence $\mathcal{A}$ is c-competitive. □

Theorem 0.3.5.2 *Let $\mathcal{S}$ be a compactly shielded MSS. Then the following conditions are equivalent*

(i) *$\mathcal{S}$ has a wf-based c-competitive strategy.*
(ii) *$\mathcal{S}$ has a c-competitive history-based strategy.*
(iii) *$\mathcal{S}$ has a c-potential.*

Proof: Any wf-based c-competitive strategy can be expressed as a time-dependent wf-based c-competitive strategy. For given $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$ and $r \in \mathcal{R}$, just let $\mathcal{A}(\omega, x, t, r) = \mathcal{A}(\omega, x, r)$ for all $t \in \mathbf{N}$. By Theorem 0.3.4.2, (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii). By Lemma 0.3.5.1, (iii) $\Rightarrow$ (i). □

0.3.6 On the Existence of Optimally Competitive Strategies

Let $\mathcal{S}$ be a metrical service system, $\mathcal{S} = \langle M, \mathcal{R} \rangle$. We say that an on-line strategy for $\mathcal{S}$ is *optimally competitive* if it is c-competitive for some c, and there is no c'-competitive on-line strategy for $\mathcal{S}$ for any $c' < c$.

In this section, we prove that if $\mathcal{S}$ is compactly shielded, bounded, and has a competitive strategy, then there exists an optimally competitive on-line strategy for $\mathcal{S}$. This does not hold in general. We present an example of an MSS that has an almost 1-competitive strategy, but does not have a 1-competitive strategy.

Lemma 0.3.6.1 *If M is bounded, and if $\mathcal{S}$ is c-competitive, then $\phi_c^* \leq \perp^c + (c+1)\,\mathrm{diam}(M)$.*

Proof: Suppose not. Without loss of generality, $\mathrm{diam}(M) = 1$. Pick $(\omega, x) \in \mathcal{W} \times M$ such that $\phi_c^*(\omega, x) = -c\,\inf(\omega) + 1 + c + \epsilon$, for $\epsilon > 0$.

For any pair $(\eta, y) \in \mathcal{W} \times M$, by the Lipschitz property of the minimum potential, and by the inequality $\sup(\eta - \omega) \leq \sup(\eta) - \inf(\omega) \leq 1 + \inf(\eta) - \inf(\omega)$, we have

$$\begin{aligned} \phi_c^*(\eta, y) \;&\geq\; \phi_c^*(\omega, x) - xy - c\,\sup(\eta - \omega) \;\geq\; \phi_c^*(\omega, x) - xy - c - c\,\inf(\eta) + c\,\inf(\omega) \\ &=\; 1 + \epsilon - xy - c\,\inf(\eta) \;\geq\; \perp^c(\eta, y) + \epsilon. \end{aligned}$$

Take $\xi = \phi_c^* - \epsilon$. By the argument above, $\xi \geq \perp^c$. Also, $\mathcal{B}\xi = \mathcal{B}\phi_c^* - \epsilon = \xi$, proving that ξ is a c-potential smaller than ϕ_c^*, contradiction. □

Lemma 0.3.6.2 *If $c < d$, and if $\mathcal{S}$ is c-competitive, then $\mathcal{S}$ is d-competitive.*

Proof: Any c-competitive strategy is d-competitive. □

Theorem 0.3.6.3 *Suppose that $\mathcal{S} = \langle M, \mathcal{R}\rangle$ is compactly shielded, and that M is bounded. If $\mathcal{S}$ is c-competitive, for some $c \geq 1$, then it has an optimally competitive strategy.*

Proof: Without loss of generality, $\operatorname{diam}(M) = 1$. Let $c^* = \inf\{c : \mathcal{S} \text{ is } c\text{-competitive}\}$. By Lemma 0.3.6.2, $\mathcal{S}$ is d-competitive for $d = c^* + 1$. For any integer $n \geq 0$, let $\Gamma^n = \min\left\{\perp^{c^*}, \perp^d + n\right\}$.

Claim A: $\mathcal{B}\Gamma^n \geq \Gamma^n$, for each $n \geq 0$.

Proof of Claim A: From Fact 0.3.1.1 (i), we know that $\mathcal{B}\perp^{c^*} \geq \perp^{c^*}$ and $\mathcal{B}\perp^d \geq \perp^d$. Fix $(\omega, x) \in \mathcal{W} \times M$ and $\epsilon > 0$, and pick $r \in \mathcal{R}$ such that $\inf(\omega \wedge r) < \inf \omega + \epsilon$. This is possible since M contains no useless points. Then

$$\begin{aligned}
\mathcal{B}\Gamma^n(\omega, x) \;&\geq\; \inf_{y \in r}\{xy + \Gamma^n(\omega \wedge r, y)\} \;\geq\; \min\{-c^* \inf(\omega \wedge r), -d \inf(\omega \wedge r) + n\} \\
&\geq\; \min\{-c^* \inf \omega, -d \inf \omega + n\} - d\epsilon \;=\; \Gamma^n(\omega, x) - d\epsilon.
\end{aligned}$$

Letting $\epsilon \to 0$, we verify Claim A.

For each $n \geq 0$, let $\psi^n = \lim_{m\to\infty} \mathcal{B}^m \Gamma^n$. By Claim A, $\psi^n \geq \Gamma^n$. Since $\mathcal{B}$ is monotone and $\Gamma^n \leq \perp^d + n$, we have $\psi^n \leq \lim_{m\to\infty} \mathcal{B}^m(\perp^d + n) \leq \phi_d^* + n$. Also, since $\Gamma^{n+1} \geq \Gamma^n$ and $\mathcal{B}$ is monotone, $\psi^{n+1} \geq \psi^n$, for all n.

Claim B: For any $n \geq 0$, $\psi^n \leq \perp^{c^*} + c^* + 1$.

Proof of Claim B: Fix n and consider any $c^* < c \leq d$. By Lemma 0.3.6.2, $\mathcal{S}$ is c-competitive. By Lemma 0.3.6.1, since $\mathcal{B}$ is monotone and $\Gamma^n \leq \perp^c + (c - c^*)n$, for each fixed $(\omega, x) \in \mathcal{W} \times M$ we have

$$\begin{aligned}
\psi^n(\omega, x) \;&=\; \lim_{m\to\infty} \mathcal{B}^m \Gamma^n(\omega, x) \;\leq\; \lim_{m\to\infty} \mathcal{B}^m \perp^c(\omega, x) + (c - c^*)n \\
&\leq\; \phi_c^*(\omega, x) + (c - c^*)n \;\leq\; \perp^c(\omega, x) + c + 1 + (c - c^*)n,
\end{aligned}$$

and letting $c \to c^*$ we verify Claim B.

Let $\psi^* = \lim_{n\to\infty} \psi^n$, the pointwise limit. By Claim B, ψ^* is well-defined.

Claim C: ψ^* is a c^*-potential for $\mathcal{S}$.

Proof of Claim C: We have $\mathsf{Dom}\,\psi^* = \mathcal{W} \times M$. Thus we need to verify conditions (pot2) and (pot3). For any $(\omega, x) \in \mathcal{W} \times M$, pick $n \geq \inf \omega$. Then

$$\psi^*(\omega, x) \;\geq\; \psi^n(\omega, x) \;\geq\; \Gamma^n(\omega, x) \;=\; \perp^{c^*}(\omega, x),$$

proving (pot2). To verify (pot3), we need only show that for fixed $(\omega, x) \in \mathcal{W} \times M$ and $r \in \mathcal{R}$

$$\inf_{z \in r} \{xz + \psi^*(\omega \wedge r, z)\} \leq \psi^*(\omega, x).$$

Note that ψ^* is d-Lipschitz. This follows from the fact that each Γ^n is d-Lipschitz, that $\mathcal{B}$ preserves the Lipschitz property (See Fact 0.3.1.4, (i)), and that ψ^* is a pointwise limit of functions $\mathcal{B}^m\Gamma^n$.

Now pick a compact set $Y \subseteq r$ which shields r from $\{x\}$. For each $z \in r$ choose a $y_z \in Y$ such that $xz = xy_z + y_z z$. By Lemma 0.3.2.2, ψ^n is a fixed point of $\mathcal{B}$ for each n. Using the d-Lipschitz property of ψ^*, and Lemma 0.2.2.2 for $\lambda = 2$,

$$\begin{aligned}
\inf_{z \in r} \{xz + \psi^*(\omega \wedge r, z)\} &\leq \inf_{y \in Y} \{xy + \psi^*(\omega \wedge r, y)\} = \inf_{y \in Y} \sup_n \{xy + \psi^n(\omega \wedge r, y)\} \\
&= \sup_n \inf_{y \in Y} \{xy + \psi^n(\omega \wedge r, y)\} \leq \sup_n \inf_{z \in r} \{xy_z + y_z z + \psi^n(\omega \wedge r, z)\} \\
&= \sup_n \inf_{z \in r} \{xz + \psi^n(\omega \wedge r, z)\} \leq \sup_n \psi^n(\omega, x) = \psi^*(\omega, x).
\end{aligned}$$

This completes the proof of Claim C and the theorem. □

Now we present an example that shows that Theorem 0.3.6.3 does not necessarily hold if $\mathcal{S}$ is not compactly shielded.

Example 2: Let $\mathcal{S} = \langle M, \mathcal{R} \rangle$, where M and $\mathcal{R}$ are defined as follows. Denote $I = (0, 1]$, and take $M = X \cup Y \cup Z \cup T$, for $X = \{x_\delta\}_{\delta \in I}$ $Y = \{y_\delta\}_{\delta \in I}$, $Z = \{z_\delta\}_{\delta \in I}$, $T = \{t_\delta\}_{\delta \in I}$, and the metric is defined as follows: for all $\delta, \epsilon \in (0, 1]$, $\delta \neq \epsilon$,

$$\begin{aligned}
x_\delta y_\delta &= x_\delta z_\delta = x_\delta x_\epsilon = 1 & \qquad y_\delta y_\epsilon &= y_\delta z_\epsilon = z_\delta z_\epsilon = 3 \\
x_\delta t_\delta &= 1 + \delta & y_\delta t_\epsilon &= z_\delta t_\epsilon = 3 + \epsilon \\
x_\delta y_\epsilon &= x_\delta z_\epsilon = 2 & t_\delta t_\epsilon &= 3 + \epsilon + \delta \\
x_\delta t_\epsilon &= 2 + \epsilon & &
\end{aligned}$$

Pictorially, the space consists of a uniform subspace X, with the 4-cycle $(x_\delta, y_\delta, z_\delta, t_\delta)$ attached to each $x_\delta \in X$ (see Fig. 0.1).

The request set is $\mathcal{R} = \{X, Y \cup T, Z \cup T, Y \cup Z \cup T\}$. We claim that $\mathcal{S}$ has a $(1+\epsilon)$-competitive strategy for each $\epsilon > 0$, but it does not have a 1-competitive strategy.

Fix $\epsilon > 0$. Our $(1+\epsilon)$-competitive strategy $\mathcal{A}_\epsilon$ is defined as follows: If the request is X, move to x_ϵ. If the request is $Y \cup T$, $Z \cup T$ or $Y \cup Z \cup T$, move the server to t_ϵ.

That $\mathcal{A}_\epsilon$ is $(1+\epsilon)$-competitive follows from the following two observations: (1) Given any two requests $r, r' \in \mathcal{R}$ such that $r \subseteq r'$, without loss of generality we can assume that r' precedes all occurrences of r in each request sequence. (2) Without loss of generality each request does not contain the current server position. These two observations imply that when we use strategy $\mathcal{A}_\epsilon$, we need to consider only the request sequence that alternates X and $Y \cup Z \cup T$. In this request sequence, at every step, the optimal cost is 1 and our cost is $1 + \epsilon$.

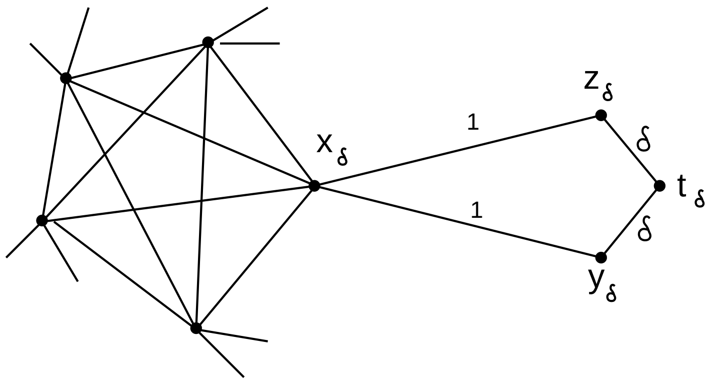

Figure 0.1: The space M from Example 2.

Non-existence of a 1-competitive strategy can be proven using the following adversary strategy: request X and $Y \cup Z \cup T$. If, after the second request, the server is in Y, request $Z \cup T$. If the server is in Z, request $Y \cup T$. Then request X again, and so on. The details are left to the reader. $\diamond$

The strategies $\mathcal{A}_\epsilon$ constructed in Example 2 can be combined into a *single* strategy $\mathcal{A}$ which works in phases, simulating strategy $\mathcal{A}_{1/2^i}$ in phase number i, for $i = 1, 2, \ldots$, where phase i lasts long enough so that the optimal cost in this phase is at least 2^i. Then $\mathcal{A}$ is almost 1-competitive. The detailed proof is rather complicated, so we present it below for interested readers, in a more general setting.

Theorem 0.3.6.4 *Let $\mathcal{S} = \langle M, \mathcal{R} \rangle$ be an MSS where M is bounded, and suppose that $\mathcal{S}$ is $(c+\epsilon)$-competitive, for each $\epsilon > 0$. Then $\mathcal{S}$ has an almost c-competitive strategy.*

Proof: Without loss of generality, $\mathrm{diam}(M) = 1$. For each $i \geq 0$, let $D_i = 2i\,(c+2)$ and $A_i = 2^i\,(c+2)$. Let $c_i = c + 2^{-i}$ for all $i \geq 0$. Define $\gamma \in \textsc{Part}_{\mathcal{S}}$ as follows:

$$\gamma(\omega, x) = \begin{cases} \phi^*_{c_0}(\omega, x) - D_0 & \text{if } \inf\omega \leq A_1 \\ \phi^*_{c_i}(\omega, x) - D_i & \text{if } A_i < \inf\omega \leq A_{i+1} \text{ for } i > 0 \end{cases}$$

Claim A: If $j < i$ and $\inf(\omega) \leq A_{j+1}$ then $\perp^{c_j}(\omega, x) - D_j \geq \perp^{c_i}(\omega, x) + c + 2 - D_i$.

Proof of Claim A: Let $d = i - j - 1 \geq 0$. Then

$$\begin{aligned}(\perp^{c_j}(\omega, x) - D_j) - (\perp^{c_i}(\omega, x) + c + 2 - D_i) &= -(2^{-j} - 2^{-i})\inf(\omega) + (2i - 2j - 1)(c+2) \\ &\geq (-2 + 2^{-i+j+1} + 2i - 2j - 1)(c+2) \\ &= (2d - 1 + 2^{-d})(c+2) \geq 0.\end{aligned}$$

Claim B: $\gamma(\omega, x) \geq \phi^*_{c_i}(\omega, x) - D_i$ if $i \geq 0$ and $\inf \omega \leq A_{i+1}$.

Proof of Claim B: Choose the largest integer j such that $\inf \omega \leq A_{j+1}$. If $j = i$, we are done by the definition of γ, so we may assume $j \leq i - 1$. Then, by Lemma 0.3.6.1 and Claim A,

$$\begin{aligned} \gamma(\omega, x) &= \phi^*_{c_j}(\omega, x) - D_j \;\geq\; \perp^{c_j}(\omega, x) - D_j \\ &\geq \perp^{c_i}(\omega, x) + c + 2 - D_i \;\geq\; \phi^*_{c_i}(\omega, x) - D_i. \end{aligned}$$

Claim C: $\gamma \geq \perp^{c_i} - D_i - c - 2$, for each i.

Proof of Claim C: Let $(\omega, x) \in \mathcal{W} \times M$. Choose the largest integer j such that $\inf \omega \leq A_{j+1}$. If $j \leq i$, we are done by Claim B, since $\phi^*_{c_i} \geq \perp^{c_i}$. Thus, we may assume $j \geq i + 1$. Note that $\inf \omega > A_j = 2^{-j}(c + 2)$, and that $2^d \geq 2d$ for any positive integer d. Then

$$\begin{aligned} \gamma(\omega, x) - \perp^{c_i}(\omega, x) + c + 2 + D_i &\geq -c_j \inf \omega - D_j + c_i \inf \omega + c + 2 + D_i \\ &= (2^{j-i} - 1)2^{-j} \inf \omega - (2j - 2i - 1)(c + 2) \\ &\geq (2^{j-i} - 2j + 2i)(c + 2) \;\geq\; 0. \end{aligned}$$

Claim D: $\mathcal{B}\gamma \leq \gamma$.

Proof of Claim D: We need only show that, for all $(\omega, x) \in \mathcal{W} \times M$, and all $r \in \mathcal{R}$,

$$\gamma(\omega, x) \geq \inf_{y \in r} \{xy + \gamma(\omega \wedge r, y)\}$$

Pick the smallest integer i such that $\inf(\omega \wedge r) \leq A_{i+1}$. By Claim B, since $\inf(\omega) \leq \inf(\omega \wedge r)$, and since $\phi^*_{c_i}$ is a fixed point of $\mathcal{B}$,

$$\gamma(\omega, x) \;\geq\; \phi^*_{c_i}(\omega, x) - D_i \;\geq\; \inf_{y \in r} \left\{xy + \phi^*_{c_i}(\omega \wedge r, y)\right\} - D_i \;=\; \inf_{y \in r} \{xy + \gamma(\omega \wedge r, y)\},$$

completing the proof of Claim D.

We now use γ to define a history-based strategy $\mathcal{A}$. Let $x \in M$ and $\varrho \in \mathcal{R}^*$. We define $\mathcal{A}(x, \varrho)$ by induction on the length of ϱ. The basis is $\mathcal{A}(x, \varepsilon) = x$. Suppose $y = \mathcal{A}(x, \varrho)$. By Claim D, we can pick $z \in r$ such that

$$yz + \gamma(\chi_x \wedge \varrho r, z) < \gamma(\chi_x \wedge \varrho, y) + 2^{-|\varrho|} \tag{0.3}$$

Then let $\mathcal{A}(x, \varrho r) = z$. Summing (0.3) over all moves, we obtain

$$cost_{\mathcal{A}}(x, \varrho) + \gamma(\chi_x \wedge \varrho, \mathcal{A}(x, \varrho)) < \gamma(\chi_x, x) + 1.$$

By Claim C and Lemma 0.2.3.3, $cost_{\mathcal{A}}(x, \varrho) \leq c_i\, opt(x, \varrho) + D_i + c + 3 + \gamma(\chi_x, x)$. Thus, $\mathcal{A}$ is c_i-competitive. By Lemma 0.3.6.2, $\mathcal{A}$ is $(c + \epsilon)$-competitive for all $\epsilon > 0$. $\square$

0.4 The k-Point Request Problem in a Uniform Space

Recall that the *k-point request problem* for a metric space M, denoted $k-\mathrm{PRP}_M$, is defined to be the metrical service system $\langle M, \mathcal{R}\rangle$, where $\mathcal{R} = \mathcal{P}_k(M)$ is the set of all non-empty subsets of M of cardinality at most k.

A metric space M is called *uniform* if all distances in M are equal. (Without loss of generality, all distances are 1.) In this section we present a deterministic k-competitive strategy and prove that it is optimally competitive.

Theorem 0.4.0.5 *There is a k-competitive strategy for the $k-\mathrm{PRP}$ in any uniform metric space M, and the constant k is optimal if M has at least $k+1$ points.*

Proof: The upper bound. Let $\mathcal{V}$ be the set of all work functions $\omega \in \mathcal{W}_M$ such that $\omega(M) \subseteq \{a, a+1\}$ for some integer a, and such that the set $S_\omega = \{x : \omega(x) = a\}$ has cardinality at most k. For all $\omega \in \mathcal{V}$ and $r \in \mathcal{R}$, if $S_\omega \cap r = \emptyset$, then $S_{\omega\wedge r} = r$, otherwise $S_{\omega\wedge r} = S_\omega \cap r$.

Now we define our strategy $\mathcal{A}$. The domain of $\mathcal{A}$ is $\mathsf{Dom}\,\mathcal{A} = \{(\omega, x) : \omega \in \mathcal{V} \text{ and } x \in S_\omega\}$. Given $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, and $r \in \mathcal{R}$, $\mathcal{A}(\omega, x, r)$ is defined as follows. If $x \in r$, then stay on x, that is $\mathcal{A}(\omega, x, r) = x$. Otherwise, let $\mathcal{A}(\omega, x, r) = y$, where $y \in S_{\omega\wedge r}$ is chosen arbitrarily.

We define a potential function $\phi(\omega, x) = |S_\omega| - k\,\omega(x)$, for all $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$. Since ϕ is k-solvent, it suffices to show that $\Delta cost_{\mathcal{A}} + \Delta\phi \le 0$. In our notation, for $(\omega, x) \in \mathsf{Dom}\,\mathcal{A}$, we need

$$xy + |S_\mu| - |S_\omega| \ \le\ k\,[\mu(y) - \omega(x)], \tag{0.4}$$

where $\mu = \omega\wedge r$ and $y = \mathcal{A}(\omega, x, r)$ is the new server position.

We consider two cases. Suppose $r \cap S_\omega \ne \emptyset$. Then $|S_\mu| \le |S_\omega|$ and $\mu(y) = \omega(x)$. If $x \in r$ then $xy = 0$, and (0.4) is true. Otherwise, if $x \notin r$, we have $xy = 1$ and $|S_\mu| < |S_\omega|$ (since $x \in S_\omega - S_\mu$), and therefore (0.4) holds.

Suppose now that $r \cap S_\omega = \emptyset$. Then $xy = 1$, $|S_\omega| \ge 1$, $|S_\mu| \le k$ and $\mu(y) - \omega(x) = 1$, and therefore (0.4) holds in this case as well.

The lower bound. Let $|M| = k+1$. For $X \subseteq M$, let χ_X be a function that is 0 on X and 1 elsewhere. Consider the minimum c-potential function ϕ^*. For any $X, Y \subseteq M$ such that $|X| = |Y|$, and for any $x \in X$ and $y \in Y$, we have $\phi^*(\chi_X, x) = \phi^*(\chi_Y, y)$, by Fact 0.3.1.8, since we can choose a symmetry of M which takes X to Y and x to y. Define $f_j = \phi^*(\chi_X, x)$ for $|X| = j$ and $x \in X$. For $j \ge 2$, by considering the request $X - \{x\}$, we have $f_j \ge f_{j-1} + 1$. For $j = 1$, by considering the request $M - \{x\}$, we have that $f_1 \ge f_k - c + 1$. It follows that $f_1 \ge f_1 - c + k$ and, consequently, $c \ge k$. □

The potential-based lower bound proof can be also expressed in terms of an adversary strategy as follows. Suppose that there exists a c-competitive strategy $\mathcal{A}$, with $c < k$. Let the server be initially at x^0. We describe a k-step adversarial strategy, consisting of requests $r^1, \ldots, r^k$, that forces any strategy to expend k, but which can be served optimally at a cost of 1. Let $r^1 \subseteq M$ be any set of k points which does not contain x^0. Let $x^1 \in r^1$ be the new server position. For any $t = 2, \ldots, k$, the adversary chooses $r^t = r^{t-1} - \{x^{t-1}\}$, and the strategy must choose some $x^t \in r^t$. The strategy expends $\sum_{t=1}^{k} x^{t-1}x^t = k$, while the optimal cost to service this sequence is 1, since all requests can be served from x^k. By repeating the above sequence as many times as necessary, the value of $cost_{\mathcal{A}} - c\, opt$ can be made to increase without bound, contradiction.

0.5 The Two-Point Request Problem

The CL_λ strategy. Our strategy for $k-$PRP is an example of a more general strategy that we believe is applicable to many other metrical service systems. The strategy is to balance against two different "greedy" approaches to the problem:

> *Lazy Strategy:* At every step, given a request r, move to a closest point in r.
> *Cheap Strategy:* At every step, if the current work function is ω, move to a point $y \in r$ that minimizes $\omega \wedge r(y) = \omega(y)$.

Neither of these two strategies is competitive for non-trivial metrical service systems. Consider, however, the following strategy that balances the Lazy and Cheap approaches. Let $\lambda \geq 0$.

Strategy λ-Cheap-and-Lazy, CL_λ*:* If the current work function is ω, the server position is x, and the new request is r, move to a point $y \in r$ that minimizes the value of $xy + \lambda\, \omega(y)$.

Note that the Lazy Strategy is CL_0, while the Cheap Strategy is CL_∞.

More formally, CL_λ is a request-wise total function from $\mathcal{W} \times M \times \mathcal{R} \to M$, and $\mathsf{Dom}\, \mathrm{CL}_\lambda$ is defined to be the set of all pairs $(\omega, x) \in \mathcal{W} \times M$ such that ω is supported by some $r \in \mathcal{R}$ containing x, and $\lambda\, [\omega(x) - \omega(u)] \leq xu$ for all $u \in M$. If $(\omega, x) \in \mathsf{Dom}\, \mathrm{CL}_\lambda$ and $r \in \mathcal{R}$, let $\mathrm{CL}_\lambda(\omega, x, r) = y$, where $y \in r$ is chosen to minimize $xy + \lambda\, \omega(y)$.

A careful reader may have noticed that, in general, CL_λ may not be well-defined. The first problem is, there may not exist a y that minimizes $xy + \lambda\, \omega(y)$, unless the given MSS satisfies some topological conditions. Next, we also need to show that the domain is well-defined, that is $(\omega, x) \in \mathsf{Dom}\, \mathrm{CL}_\lambda$ and $y = \mathrm{CL}_\lambda(\omega, x, r)$ implies $(\omega \wedge r, y) \in \mathsf{Dom}\, \mathrm{CL}_\lambda$. Below we will prove that CL_λ is well-defined for systems where all requests are finite. This enables us to apply the CL_λ to $k-$PRP. However, it can also be proven that CL_λ is well-defined for any MSS with the compact shielding property. We postpone that generalization to a later publication.

Lemma 0.5.0.6 *If each $r \in \mathcal{R}$ is finite then CL_λ is well-defined. That is, if $(\omega, x) \in \mathsf{Dom}\, \mathrm{CL}_\lambda$ and $r \in \mathcal{R}$, and $y = \mathrm{CL}_\lambda(\omega, x, r)$, then $(\omega \wedge r, y) \in \mathsf{Dom}\, \mathrm{CL}_k$.*

Proof: We need only show that $\lambda\,[\omega(x) - \omega(u)] \le xu$ for all $u \in M$. If $\lambda \le 1$, we are done by (wf1).

Suppose that $\lambda > 1$. We know that $xv + \lambda\,\omega(v) \ge xy + \lambda\,\omega(y)$ for any $v \in r$. Thus, for any $u \in M$

$$\begin{aligned} yu + \lambda\,\omega \wedge r(u) &= \inf_{v \in r} \{yu + \lambda\,uv + \lambda\,\omega(v)\} \ \ge\ \inf_{v \in r} \{yu + xy - xv + \lambda\,uv + \lambda\,\omega(y)\} \\ &\ge \inf_{v \in r} \{yu + xy - xv + uv + \lambda\,\omega(y)\} \ \ge\ \lambda\,\omega(y) \ =\ \lambda\,\omega \wedge r(y), \end{aligned}$$

Finally, $\omega \wedge r$ is supported by r, completing the proof. □

The strategy given in the previous section for $k-$PRP on uniform spaces is actually identical to CL_λ for every $\lambda > 1$. Below, we prove that CL_3 is 9-competitive for $2-$PRP in any metric space, and that 9 is a lower bound on the general competitiveness of $2-$PRP. Another important example is the strategy proposed for the k-server problem. In [7] it was proven that CL_1 (which is called the "work function algorithm" in that paper) is 2-competitive for two servers, and it has been conjectured that CL_1 is k-competitive for k servers, for any $k \ge 3$.

Theorem 0.5.0.7 CL_3 *is 9-competitive for* $2-$PRP *in any metric space.*

Proof: Throughout this proof, we use the following simplifying notation. For $a \in \mathbb{R}$, let $[a]^+ = \max\{a, 0\}$.

Fix a metric space M. We apply a potential argument to show that CL_3 is 9-competitive in M. Let $(\omega, x) \in \mathrm{CL}_3$. Then ω is supported by $\{x, y\}$ for some y. We define a potential function by

$$\phi(\omega, x) \quad = \quad 2\ [xy + 3\,\omega(x) - 3\,\omega(y)]^+ - 9\,\omega(x).$$

If $\omega(x) \le \omega(y)$ then $\phi(\omega, x) \ge -9\,\omega(x) = -9\ \inf(\omega)$. If $\omega(x) \ge \omega(y)$ then, $xy + 3\,\omega(y) - 3\,\omega(x) \ge 0$, and we obtain

$$\phi(\omega, x) \ = \ xy + [xy + 3\,\omega(y) - 3\,\omega(x)] - 9\,\omega(y) \ \ge\ -9\,\omega(y) \ = \ -9\ \inf(\omega).$$

Thus ϕ is solvent. It remains to prove that $\Delta cost_{\mathrm{CL3}} + \Delta\phi \le 0$, in every move of CL_3.

Let $(\omega, x) \in \mathsf{Dom}\,\mathrm{CL}_3$. Suppose that ω is supported by $\{x, y\}$. Thus ω is the current work function, x is the current server position, and y is the "other point" in the last request $\{x, y\}$. Without loss of generality, $xu + 3\,\omega(u) \le xv + 3\,\omega(v)$, *i.e.*, CL_3 will move its server to u.

$$\Delta cost_{\mathrm{CL3}} + \Delta\phi \ = \ xu + 2\ [uv + 3\,\omega(u) - 3\,\omega(v)]^+ - 9\,\omega(u) - 2\ [xy + 3\,\omega(x) - 3\,\omega(y)]^+ + 9\,\omega(x).$$

The proof that $\Delta cost_{\mathrm{CL3}} + \Delta\phi \le 0$ is by analyzing a number of cases. In calculation we use the triangle inequality in M, and the inequalities: $3\,[\omega(x) - \omega(y)] \le xy$ and $3\,[\omega(u) - \omega(v)] \le xv - xu$. We also use the fact that $[a]^+ - [b]^+ \le [a-b]^+$ for all a, b, and $[a]^+ - b \le [a-b]^+$ if $b \ge 0$.

Case 1: $\omega(u) = \omega(x) + xu$ and $\omega(v) = \omega(x) + xv$. Then

$$\begin{aligned} \Delta cost_{\mathrm{CL3}} + \Delta\phi &= -8\,xu + 2\ [uv + 3\,xu - 3\,xv]^+ - 2\ [xy + 3\,\omega(x) - 3\,\omega(y)]^+ \\ &\le -8\,xu + 2\ [uv + 3\,xu - 3\,xv]^+ \ \le\ 2\ [uv -\ xu - 3\,xv]^+ \ = \ 0. \end{aligned}$$

Case 2: $\omega(u) = \omega(y) + yu$ and $\omega(v) = \omega(y) + yv$. Then

$$\begin{aligned}
\Delta cost_{\mathrm{CL3}} + \Delta\phi &= xu + 2\,[uv + 3\,yu - 3\,yv]^+ - 9\,yu - 2\,[xy + 3\,\omega(x) - 3\,\omega(y)]^+ + 9\,[\omega(x) - \omega(y)] \\
&\le xu - yu + 2\,[uv - \,yu - 3\,yv]^+ - 2\,(xy + 3\,\omega(x) - 3\,\omega(y)) + 9\,[\omega(x) - \omega(y)] \\
&\le -xy + 3\,[\omega(x) - \omega(y)] \;\le\; 0.
\end{aligned}$$

Case 3: $\omega(u) = \omega(x) + xu$ and $\omega(v) = \omega(y) + yv$. Then

$$\begin{aligned}
\Delta cost_{\mathrm{CL3}} + \Delta\phi &= -8\,xu + 2\,[uv + 3\,xu - 3\,yv + 3\,\omega(x) - 3\,\omega(y)]^+ - 2\,[xy + 3\,\omega(x) - 3\,\omega(y)]^+ \\
&\le -8\,xu + 2\,[uv + 3\,xu - 3\,yv - xy]^+ \;\le\; 2\,[uv - xu - 3\,yv - xy]^+ \;=\; 0.
\end{aligned}$$

Case 4.1: $\omega(u) = \omega(y) + yu$, $\omega(v) = \omega(x) + xv$ and $3\,\omega(u) \ge 3\,\omega(v) - uv$. Then

$$\begin{aligned}
\Delta cost_{\mathrm{CL3}} + \Delta\phi &\le xu + 2\,(uv + 3\,\omega(u) - 3\,\omega(v)) - 9\,\omega(u) - 2\,(xy + 3\,\omega(x) - 3\,\omega(y)) + 9\,\omega(x) \\
&= xu + 2\,uv - 2\,xy - 3\,xv - 6\,yu + 3\,[\omega(u) - \omega(v)] \\
&\le 2\,uv - 2\,xy - 2\,xv - 6\,yu \;\le\; 0.
\end{aligned}$$

Case 4.2: $\omega(u) = \omega(y) + yu$, $\omega(v) = \omega(x) + xv$ and $3\,\omega(u) \le 3\,\omega(v) - uv$. Then

$$\begin{aligned}
\Delta cost_{\mathrm{CL3}} + \Delta\phi &\le xu - 9\,\omega(u) - 2\,(xy + 3\,\omega(x) - 3\,\omega(y)) + 9\,\omega(x) \\
&= xu - 9\,yu - 2\,xy + 3\,[\omega(x) - \omega(y)] \;\le\; xu - 9\,yu - xy \;\le\; 0.
\end{aligned}$$

This completes the analysis of all cases and the proof of the theorem. □

Lower Bound. We now give a lower bound of 9 for the competitiveness constant of any deterministic strategy for $2-\mathrm{PRP}$ in $M = \mathbb{R}$, the real line. More formally, we deal with the metrical service system $2-\mathrm{PRP}_{\mathbb{R}} = \langle \mathbb{R}, \mathcal{P}_2(\mathbb{R})\rangle$, where $\mathbb{R}$ is viewed as a metric space in which the distance function is $xy = |x - y|$ for all $x, y \in \mathbb{R}$, and $\mathcal{P}_2(\mathbb{R})$ is the set of all 2-element subsets of $\mathbb{R}$. This result shows that CL_3 is optimally competitive for general metric spaces. Our method is to assume that there exists a c-competitive strategy for $2-\mathrm{PRP}_{\mathbb{R}}$, and reach a contradiction if $c < 9$.

We first introduce some specialized notation. For any $a, b \in \mathbb{R}^+$, let $\omega_{a,b} \in \mathcal{W} = \mathcal{W}_{\mathbb{R}}$ be the work function with support $\{-a, b\}$ defined by $\omega_{a,b}(-a) = a$ and $\omega_{a,b}(b) = b$. Let $\phi^* = \phi^*_c$ be the minimum c-potential, which must exist if a problem is c-competitive, by Theorem 0.3.4.2 and Theorem 0.3.1.6. Let $f(a,b) = \phi^*(\omega_{a,b}, -a)$.

In the lemmas below, we will translate the properties of the minimum potential function into properties of f, which, being simply a function from the positive quadrant of $\mathbb{R}^2$ into $\mathbb{R}$, is easier to deal with than ϕ^*. Then we will derive certain inequalities involving c and values of f at certain points. This system of inequalities will not have a solution if $c < 9$.

Lemma 0.5.0.8 *Suppose that $x, y \in \mathbb{R}$, and that $\omega \in \mathcal{W}$ is supported by $\{x, y\}$. Then $\phi^*(\omega, x) = f(a,b) - c\,[\omega(x) - a]$, where a and b are defined by*

$$a + b \;=\; |x - y|, \quad \text{and} \quad a - b \;=\; \omega(x) - \omega(y).$$

Proof: Let $F : \mathbb{R} \to \mathbb{R}$ be the unique symmetry such that $F(x) = -a$ and $F(y) = b$. By Fact 0.3.1.7 and by Fact 0.3.1.8 (iii), we are done. □

Lemma 0.5.0.9 *Let $a, b \in \mathbb{R}^+$. Then*

(i) $f(a,b) \le f(b,a) + a + b$.
(ii) $\lambda f(a,b) = f(\lambda a, \lambda b)$, *for all* $\lambda \in \mathbb{R}^+$.
(iii) $f(a+\epsilon, b+\delta) \le f(a,b) + (c+1)\,|\epsilon| + c\,|\delta|$.

Proof: (i) Let $F(t) = -t$ for $t \in \mathbb{R}$. Then F is a symmetry, and $F(\omega_{a,b}) = \omega_{b,a}$. Since ϕ_c^* is a fixed point of F, by Fact 0.3.1.8 (iii), and since ϕ_c^* is Lipschitz

$$f(a,b) \;=\; \phi_c^*(\omega_{a,b}, -a) \;=\; \phi_c^*(\omega_{b,a}, a) \;\le\; \phi_c^*(\omega_{b,a}, -b) + a + b \;=\; f(b,a) + a + b.$$

(ii) Let $F(t) = \lambda\, t$ for all $t \in \mathbb{R}$. Then F is a λ-symmetry, and $F(\omega_{a,b}) = \omega_{\lambda\, a, \lambda\, b}$. Since ϕ_c^* is a fixed point of F, by Fact 0.3.1.8 (iii)

$$f(\lambda\, a, \lambda\, b) = \phi_c^*(\omega_{\lambda\, a, \lambda\, b}) = \lambda\, \phi_c^*(\omega_{a,b}) = \lambda\, f(a,b).$$

(iii) Since ϕ_c^* is Lipschitz, and since $\sup(\omega_{a,b} - \omega_{a+\epsilon, b+\delta}) \le |\epsilon| + |\delta|$,

$$\begin{aligned} f(a+\epsilon, b+\delta) - f(a,b) \;&=\; \phi^*(\omega_{a+\epsilon,b+\delta}, -a-\epsilon) - \phi^*(\omega_{a,b}, -a) \\ &\le\; |\epsilon| + c\, \sup(\omega_{a,b} - \omega_{a+\epsilon,b+\delta}) \;\le\; (c+1)\,|\epsilon| + c\,|\delta|. \end{aligned}$$

This completes the proof. □

Lemma 0.5.0.10 *Fix $a_2 > a_1 \ge 0$ and $b \ge 0$. Suppose that $f(a,b) < f(b,a) + a + b$ for all $a \in (a_1, a_2)$. Then $f(a_1, b) \ge f(a_2, b) + a_2 - a_1$.*

Proof: Let $g(a) = f(a,b)$ for $a \in (a_1, a_2)$. By Lemma 0.5.0.9 (iii), the derivative of g is defined for almost all a.[2] For fixed $a \in (a_1, a_2) \subseteq \mathbb{R}^+$, let $r = \{-a-\epsilon, b\}$. Note that $\omega_{a,b} \wedge r = \omega_{a+\epsilon,b}$. Then

$$\begin{aligned} f(a,b) \;&=\; \phi^*(\omega_{a,b}, -a) \;=\; \mathcal{B}\phi^*(\omega_{a,b}, -a) \;\ge\; \min_{y \in r} \{\phi^*(\omega_{a,b} \wedge r, y) + |y + a|\} \\ &=\; \min \left\{ \begin{array}{l} \phi^*(\omega_{a+\epsilon,b}, -a-\epsilon) + \epsilon \\ \phi^*(\omega_{a+\epsilon,b}, b) + a + b \end{array} \right\} \;=\; \min \left\{ \begin{array}{l} f(a+\epsilon, b) + \epsilon \\ f(b, a+\epsilon) + a + b \end{array} \right\}. \end{aligned}$$

[2] The function g satisfies a classical *Lipschitz condition*, namely that $|g(a) - g(a')| \le \lambda\,|a - a'|$ for some constant λ. It follows that the derivative of g exists everywhere, and that g is an indefinite integral of g', allowing us to use the Fundamental Theorem of Calculus.

By Lemma 0.5.0.9 (iii), f is continuous. Thus, the first of the two choices is smaller if $\epsilon \in (0, \epsilon_a)$, for some sufficiently small $\epsilon_a > 0$. Thus $f(a,b) \geq f(a+\epsilon, b) + \epsilon$ for each $\epsilon \in (0, \epsilon_a)$. Since this holds for each $a \in (a_1, a_2)$, we conclude that $g'(a) \leq -1$ almost everywhere, and by the Fundamental Theorem of Calculus,

$$f(a_2, b) - f(a_1, b) \;=\; g(a_2) - g(a_1) \;=\; \int_{a_1}^{a_2} g'(a)\, da \;\leq\; -(a_2 - a_1).$$

This completes the proof. □

Theorem 0.5.0.11 *If $c < 9$ then there is no c-competitive strategy for $2-\mathrm{PRP}_{\mathbb{R}}$.*

Proof: Choose the smallest $\alpha \geq 1$ such that $f(\alpha, 1) = f(1, \alpha) + \alpha + 1$. Such an α must exist, since otherwise, by Lemma 0.5.0.9 (i), Lemma 0.5.0.10, $f(1,1) \geq f(a,1) + a - 1$ for all $a \geq 1$, while $f(a,1) = \phi^*(\omega_{a,1}, -a) \geq -c \inf(\omega_{a,1}) = -c$, a contradiction for a sufficiently large a. Denote $f_1 = f(\alpha, 1)$ and $f_2 = f(1, \alpha) = f_1 - \alpha - 1$. By Lemma 0.5.0.9 (i), we only need consider two cases:

Case I: For any $\frac{1}{\alpha} < a < 1$, $f(a,1) < f(1,a) + a + 1$.

Let $g = f(1/\alpha, 1)$. By the choice of α and Lemma 0.5.0.10, $g \geq f_1 + \alpha - 1/\alpha$. By Lemma 0.5.0.9, $\alpha\, g = f_2$. Thus $(\alpha - 1) g \;\leq\; -2\alpha - 1 + 1/\alpha$. Now, since ϕ^* is a c-potential, $g \geq -c \inf(\omega_{1/\alpha,1}) = -c/\alpha$. Therefore

$$-c(\alpha - 1)/\alpha \;\leq\; -2\alpha - 1 + 1/\alpha,$$

from which, assuming $c < 9$, we obtain $\alpha^2 - 4\alpha + 4 < 0$, contradiction.

Case II: For some $\frac{1}{\alpha} < a < 1$, $f(a,1) = f(1,a) + a + 1$.

Let $\beta = 1/a$, for the minimum $a \in (1/\alpha, 1)$ such that $f(a,1) = f(1,a) + a + 1$. By definition, $1 < \beta \leq \alpha$. Let $h = f(1, \beta)$, $i = f(\beta, 1)$, and $j = f(\beta/\alpha, \beta)$. Then by Lemma 0.5.0.9 (ii), $h = \beta f(1/\beta, 1) = \beta\,[f(1, 1/\beta) + 1/\beta + 1] = i + 1 + \beta$. By applying Lemma 0.5.0.10 to intervals $(1/\alpha, 1/\beta)$ and (β, α), and by Lemma 0.5.0.9, we have

$$\begin{aligned} f_2 &= f_1 - \alpha - 1 \;\leq\; [i - \alpha + \beta)] - \alpha - 1 \;=\; h - 2\alpha - 2, \\ h &= \beta f(1/\beta, 1) \;\leq\; \beta\,[f(1/\alpha, 1) - 1/\beta + 1/\alpha)] \;=\; j + \beta/\alpha - 1. \end{aligned}$$

By Lemma 0.5.0.9, $f_2 = (\alpha/\beta) j$. Also, since ϕ^* is c-solvent, we have $j \geq -c(\beta/\alpha)$. Combining all these equalities and inequalities, we obtain

$$-(1 - \beta/\alpha)\, c \;\leq\; -3 - 2\alpha + \beta/\alpha.$$

Finally, since $c < 9$ and $\beta > 1$, we conclude $\alpha^2 - 3\alpha + 4 < 0$, contradiction. □

It may be of some interest to note that the second case in the proof of Lemma 0.5.0.11 is "more impossible" than the first one. One can reach a contradiction for $c < 7 + \sqrt{31}$.

0.6 Final Comments

We have presented several general results about metrical service systems. The most significant, in our view, are (1) the result showing that the existence of a (history-based) c-competitive strategy is equivalent to the existence of a c-potential function, and (2) the introduction of the lower-bound technique based on the minimum c-potential.

Metrical service systems provide a convenient model for investigating various on-line problems. They are sufficiently general to encompass many known on-line optimization problems, including the k-server problem.

Open problems. The major problem involving metrical service systems that remains to be solved is the k-server problem, *i.e.*, whether there exists a k-competitive strategy for the k servers in an arbitrary metric space. Other open problems are:

(1) By [9], $k-$PRP has a competitive strategy in arbitrary metric spaces. What is the optimal competitive constant?

(2) Give a better characterization of metrical service systems for which an existence of a competitive strategy implies the existence of an optimally competitive strategy. (See Theorem 0.3.6.3). Similarly, is the boundedness hypothesis of Theorem 0.3.6.4 necessary?

(3) Investigate the CL_λ strategies. What is the relationship between λ and the competitiveness constant of CL_λ? In particular, for which triples (k, λ, c) is CL_λ c-competitive for $k-$PRP?

(4) Investigate the lower bound technique, introduced in [8], based on the *adversary potential*. Informally, in terms of metrical service systems, the idea would be to present, for each $\epsilon > 0$, a function ψ_ϵ with the property that (a) $-(c-\epsilon)\inf(\omega) \leq \psi_\epsilon(\omega, x) \leq -(c-\epsilon)\inf(\omega) + a_\epsilon$, for some constant a_ϵ, and (b) in each possible "situation" (ω, x), the adversary can make a request r such that, independent of where our strategy $\mathcal{A}$ moves, the following inequality holds: $\Delta cost_{\mathcal{A}} + \Delta\psi_\epsilon \geq \epsilon$. If (a) and (b) are true, an MSS cannot have a c-competitive strategy.

(5) Investigate the competitiveness of the MSS $\mathrm{CR}_n = \langle \mathbb{R}^n, \mathcal{C}\rangle$, $n \geq 2$, where $\mathcal{C}$ is the set of convex subsets of $\mathbb{R}^n$. It is known that for $n = 2$ the problem is competitive, see [11]. Determine the optimal competitive constant for CR_2. Are the CR_n competitive for $n \geq 3$?

Randomized algorithms. Our next paper, [6], deals with randomized algorithms for metrical service systems. We show how to reduce the randomized problem to a deterministic one, and using this reduction we extend results from the present paper to randomized algorithms. We will also provide optimally competitive algorithms for the two examples of metrical service systems considered in this paper: an H_k-competitive strategy for $k-$PRP in uniform spaces (where H_k is the k-th harmonic number), and a c-competitive strategy for $2-$PRP in arbitrary metric spaces, where $c \approx 4.59112$ is the solution to the equation $\ln(c-1) = c/(c-1)$.

Bibliography

[1] R. Baeza-Yates, J. Culberson, and G.Rawlins. Searching with uncertainty. Technical Report CS-87-68, University of Waterloo, Department of Computer Science, October 1987.

[2] S. Ben-Davida, A. Borodin, R. Karp, G. Tardos, and A. Widgerson. On the power of randomization in on-line algorithms. In *Proc. 22nd Symposium on Theory of Algorithms*, pages 379–386, 1990.

[3] A. Borodin, N. Linial, and M. Saks. An optimal online algorithm for metrical task systems. In *Proc. 19th Annual ACM Symposium on Theory of Computing*, pages 373–382, 1987.

[4] M. Chrobak and L. Larmore. A new approach to the server problem. *SIAM Journal on Discrete Mathematics*, 4:323–328, 1991.

[5] M. Chrobak and L. Larmore. Generosity helps or an 11-competitive algorithm for three servers. In *Proceedings of ACM/SIAM Symposium on Discrete Algorithms*, pages 196–202, 1992.

[6] M. Chrobak and L. Larmore. Metrical service systems: Randomized strategies. Submitted for publication, 1992.

[7] M. Chrobak and L. Larmore. The server problem and on-line games. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, volume 7, pages 11–64, 1992.

[8] M. Chrobak, L. Larmore, N. Reingold, and J. Westbrook. Page migration algorithms using work functions. Submitted for publication, 1991.

[9] A. Fiat, D. Foster, H. Karloff, Y. Rabani, Y. Ravid, and S. Vishwanathan. Competitive algorithms for layered graph traversal. In *Proc. 32nd Symposium on Foundations of Computer Science, Puerto Rico*, pages 288–297, 1991.

[10] A. Fiat, Y. Rabani, and Y. Ravid. Competitive k-server algorithms. In *Proc. 22nd Symposium on Foundations of Computer Science*, pages 454–463, 1990.

[11] J. Friedman and N. Linial. On convex body chasing. Manuscript.

[12] E. Grove. The harmonic k-server algorithm is competitive. In *Proc. 23rd ACM Symposium on Theory of Computing*, pages 260–266, 1991.

[13] M. Manasse, L. A. McGeoch, and D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11:208–230, 1990. Also in Proc. 20th Annual ACM Symposium on Theory of Computing, 1988, pp. 322-333.

[14] C.H. Papadimitriou and M. Yannakakis. Shortest paths without a map. In *16th International Colloquium on Automata, Languages, and Programming, Lecture Notes in Computer Science vol. 372*, pages 610–620. Springer-Verlag, 1989.