Analysis of the Harmonic Algorithm for Three Servers

Marek Chrobak[*‡] Jiří Sgall[†‡]

Abstract

HARMONIC is a randomized k-server algorithm that, at each step, given a request point r, chooses the server to be moved to r with probability inversely proportional to the distance to r. For general k, it is known that the competitive ratio of HARMONIC is at least $\frac{1}{2}k(k+1)$, while the best upper bound on this ratio is exponential in k. It has been conjectured that the lower bound is tight, that is, that HARMONIC is $\frac{1}{2}k(k+1)$-competitive for all k. This conjecture has been proven in a number of special cases, including $k = 2$ and for the so-called lazy adversary.

In this paper we provide further evidence for this conjecture, by proving that HARMONIC is 6-competitive for $k = 3$. Our approach is based on the random walk techniques and their relationship to the electrical network theory. In our proof, we propose a new potential function Φ and reduce the proof of the validity of Φ to several inequalities involving hitting costs. Then, we show that these inequalities hold for $k = 3$. We believe that our potential can be used to prove that HARMONIC is $\frac{1}{2}k(k+1)$-competitive in other cases as well, and quite possibly to solve the general case.

1 Introduction

In the *k-server problem* we are given a metric space M populated by k mobile serves. At each time step a request $r \in M$ is specified, and we need to move one server to r to "satisfy" the request. The decision as to which server to move needs to be made *online*, without the knowledge about future requests.

An online algorithm $\mathbf{A}$ for the k-server problem is called *R-competitive* if its cost on any request sequence ϱ is at most R times the optimal (offline) cost on this sequence, plus possibly an additive constant independent of ϱ. The same definition applies to randomized algorithms as well, we just need to replace our cost by our *expected* cost on ϱ. The *competitive ratio* of $\mathbf{A}$ is the smallest R for which $\mathbf{A}$ is R-competitive.

Manasse, McGeoch and Sleator [MMS90] gave a 2-competitive algorithm for 2 servers and they proved that no deterministic online algorithm for k servers can be better than k-competitive. They also formulated the still-open *k-server conjecture* that a k-competitive algorithm exists for each k. The best known upper bound on the competitive ratio for $k \geq 3$ is $2k - 1$, achieved by the Work Function Algorithm, as proven by Koutsoupias and Papadimitriou [KP95]. We refer the reader to [BEY98, CL98] for more information on the k-server problem.

[*]Department of Computer Science, University of California, Riverside, CA 92521. E-mail: `marek@cs.ucr.edu`. Research supported by NSF grants CCR-9988360 and CCR-0208856.

[†]E-mail: `sgall@math.cas.cz`, `http://www.math.cas.cz/~sgall`, Mathematical Inst., AS CR, Žitná 25, CZ-11567 Praha 1, Czech Republic and Institute for Computer Science, Prague (project L00A056 of MŠMT CR). Partially supported by grant 201/01/1195 of GA ČR and grant A1019901 of GA AV ČR.

[‡]Partially supported by cooperative research grant KONTAKT-ME476/CCR-9988360-001 from MŠMT ČR and NSF.

HARMONIC is a randomized k-server algorithm proposed by Raghavan and Snir [RS94]. If the distance between server j and the request point r is d_{jr}, the algorithm moves j to r with probability proportional to its inverted distance to r, that is $d_{jr}^{-1}/(\sum_i d_{ir}^{-1})$.

As shown in [RS94], the competitive ratio of HARMONIC is at least $\frac{1}{2}k(k+1)$. This is also conjectured to be the exact competitive ratio of HARMONIC (see [BKT90, BEY98]). Although this conjecture remains open in general, there is a number of results that support this conjecture in some special cases. For $k = 2$, Raghavan and Snir [RS94] showed that the competitive ratio of HARMONIC is at most 6. This case was later settled by Chrobak and Larmore [CL92], who proved that for two servers HARMONIC is 3-competitive, thus matching the lower bound. Another, simpler proof for $k = 2$ was later given by Chrobak and Sgall [CS00]. For $k = 3$, Berman *et al.* [BKT90] proved that HARMONIC's ratio is finite (the value of the upper bound in [BKT90] was astronomical). The best currently known upper bound for arbitrary k, by Bartal and Grove [BG00], is exponential in k (see also [BEY98]). In particular, for $k = 3$, the upper bound formula from [BEY98] gives the value 25, while the lower bound is 6.

The competitive analysis of HARMONIC is considered one of the main open problems in the area of on-line algorithms. There are several reasons for this. First, unlike most other k-server algorithms studied in the literature, HARMONIC is very natural and simple. It is also memoryless, that is, the choice of the server at each step depends only on the current configuration and the location of the request, not on the past requests. It can be trivially implemented in time $O(k)$ per step, independently of n (the number of requests). The only other algorithm known to have a constant competitive ratio is the Work Function Algorithm [KP95] which, in general, requires exponential space and time per step to implement (roughly $O(n^k)$). The past research on HARMONIC [RS94, CRR$^+$89, BCNR02] revealed some close connections to the analysis of random walks and electrical networks, which are worth investigating for their own interest. Finally, it is reasonable to expect that insights gained through the analysis of HARMONIC will help with resolving the question whether there exists a memoryless k-competitive randomized algorithm for k-servers in arbitrary metric spaces (see [CDRS93].)

In this paper we prove that HARMONIC is 6-competitive for $k = 3$. This matches the lower bound, and it provides further evidence for the conjectured $\frac{1}{2}k(k+1)$-competitiveness of HARMONIC. Our approach is based on the random walk techniques and their relationship to the electrical network theory. The proof is divided into two parts. First, in Section 2, we propose a candidate potential function Φ, and we show that proving $\frac{1}{2}k(k+1)$-competitiveness of HARMONIC can be reduced to proving several inequalities involving hitting costs of random walks. In Section 5 we prove one of those inequalities for arbitrary k. Later, in Section 6, we prove that the remaining inequalities hold for $k = 3$, implying that HARMONIC is 6-competitive for 3 servers.

We note that all these upper bounds on HARMONIC's competitive ratio hold even against the so-called *adaptive online adversary* that can choose requests adaptively, based on the information on the locations of HARMONIC's servers.

2 A General Framework of the Proof

Throughout the paper, M is an arbitrary but fixed metric space. We will analyze the performance of HARMONIC in M. By d_{uv} we denote the distance between points $u, v \in M$, and by $c_{uv} = 1/d_{uv}$ we denote the inverse of the distance, also called the *conductance* of edge (u, v).

Harmonic random walks. Let G be a k-point subspace of M. Let $s \in G$ and $t \in M - G$ be two nodes called the *start node* and *target node*. The *harmonic random st-walk in G* is defined as follows:

We start at node s. At each step, if the current location is $u \neq t$, move from u to another vertex $v \in G + t$ with probability inversely proportional to the distance between u and v. (For simplicity, we write $G + t = G \cup \{t\}$ and $G - x = G \setminus \{x\}$.) In terms of conductances, this probability is $c_{uv} / \sum_{x \in G+t-u} c_{ux}$. The *cost* of move from u to v is d_{uv}. The random walk terminates when $u = t$.

By H_{st} we denote the expected total cost of a harmonic st-walk in G, and we call it the *hitting cost.* (We omit G in this notation, as it is always fixed or understood from context.)

For a given target node t, the values H_{ut}, for $u \in G$, are related to each other by the following system of linear equations:

$$\left(\sum_{v \in G+t-u} c_{uv} \right) \cdot H_{ut} \;=\; k + \sum_{v \in G-u} c_{uv} H_{vt}. \tag{1}$$

System (1) has a finite and unique solution that can be expressed as a rational function in the distances (see, for example, [DS84]). We discuss this solution in more detail in Section 3.

Doubled points. We also consider a limit case of harmonic random st-walks in G when the target node t is at distance 0 from some node $v \in G$. More formally, let t be a "copy" of v, that is, for all $u \in G - v$ we have $c_{ut} = c_{uv}$ and the value of c_{vt} is arbitrary. We define $\widehat{H}_{uv} = \lim_{c_{vt} \to \infty} H_{ut}$. Intuitively, for v infinitely close to t, this is the same a as a hitting cost of the harmonic random walk in $G - v$ with two target nodes v and t, where t is a copy of v, so that the walk starts at s and terminates if we reach either v or t (and thus c_{vt} is irrelevant).

Key inequalities. As before, let G be a k-point subspace of M, $s, v \in G$ and $t \in M - G$. We now formulate a number of inequalities involving hitting costs H_{st} and $\widehat{H}_{sv}$. Later in this section we show that these inequalities imply $\frac{1}{2}k(k+1)$-competitiveness of Harmonic. In the following sections, we give a proof of inequality (D), and we prove the remaining three inequalities for $k = 3$.

These key inequalities are:

$$\sum_{v \in G+t-s} \left| \frac{\partial H_{st}}{\partial d_{sv}} \right| \;\leq\; \tfrac{1}{2}k(k+1) \tag{A}$$

$$H_{st} - H_{vt} \;\leq\; \widehat{H}_{sv} \tag{B}$$

$$(c_{vt} - c_{vs})\widehat{H}_{sv} \;\leq\; c_{vt} H_{st} \tag{C}$$

$$(c_{vs} - c_{vt}) H_{vt} \;\leq\; c_{vs} H_{st} \tag{D}$$

Inequality (A) has a natural interpretation. The *edge stretch* of a random walk from s to t is the ratio between the expected cost of this walk and d_{st}. As shown in [BCNR02], the edge stretch of a harmonic random st-walk is at most $\frac{1}{2}k(k+1)$. Inequality (A) is an extension of this result, since it says that if the position of s changes by some distance $\epsilon > 0$, then the hitting cost changes at most by $\frac{1}{2}k(k+1)\epsilon$.

In inequality (B), notice that it becomes an equality for $t = v$. So inequality (B) addresses the following question: where should we place the target point t so that it is much more expensive to reach from s than from v? The assertion of (B) is that it should be at point v.

In inequality (D), consider the special case when s is in the interval between v and t, i.e. $d_{vs} + d_{st} = d_{vt}$. Then (D) asserts that $H_{vt}/d_{vt} \leq H_{st}/d_{st}$. This agrees with the intuition that the stretch factor increases when we move two endpoints towards each other. As we show in Section 5, this inequality is indeed true for all k. Inequality (C) has a similar interpretation when t is between s and v. (However, counter-intuitively, we can show that, although true for $k = 3$, this inequality does not hold for large values of k.)

The potential function. We use the harmonic random walks introduced above to define a potential function for the analysis of HARMONIC. By $\mathcal{A} = \{1, 2, \ldots, k\}$ and $\mathcal{H} = \{1', \ldots, k'\}$ we denote the configurations of the adversary's k servers and HARMONIC's k servers, respectively. To simplify notation, we will often use the same notation for the server (j or j') and its position in M. Without loss of generality, we assume that $|\mathcal{A}| = |\mathcal{H}| = k$, that is, at any step, the adversary's servers and the HARMONIC's servers are each on k distinct points. Of course, some of HARMONIC's servers may coincide with adversary's servers.

The random walks we now consider take place in the subspace $G = \mathcal{A}$, with the target node $t \in \mathcal{H}$. In terms of HARMONIC, H_{it} can be interpreted as follows. Suppose that HARMONIC's configuration is $\mathcal{H} = \mathcal{A} + t - i$. Then H_{it} is equal to the expected cost of HARMONIC if the adversary only makes *lazy* requests, that is requests on points in $\mathcal{A}$ (without moving its servers). This can be seen by noting that the harmonic random walk from i to t exactly describes the moves of the "hole", i.e., the point covered by the adversary but not by HARMONIC.

Consider the weighted, complete bipartite graph $\mathcal{B} \subseteq \mathcal{A} \times \mathcal{H}$ in which the weight of each edge (i, j') is $H_{ij'}$. In the degenerate case when $j' \in \mathcal{A}$, this weight becomes the cost of a harmonic random walk with a doubled target point j' and thus, by definition, $H_{ij'} = \widehat{H}_{ij'}$. In particular for $i = j'$, we have $H_{ij'} = 0$. The potential of configuration $(\mathcal{A}, \mathcal{H})$ is defined as

$$\Phi(\mathcal{A}, \mathcal{H}) \quad = \quad \text{minimum weight of a perfect matching in } \mathcal{B}.$$

Theorem 2.1 *For any $k \geq 2$, if inequalities* (A)–(D) *are true, then* HARMONIC *is $\frac{1}{2}k(k+1)$-competitive against an adaptive on-line adversary.*

Proof: We view the process as follows: At each time step, the adversary issues a request specified by a point $r \in M$. Since the adversary is adaptive, he can choose r depending on the current position of HARMONIC's servers. Once the request is issued, first the adversary serves it with one of its servers and then HARMONIC chooses randomly which server to move to r.

Consider an arbitrary, but fixed, configuration $(\mathcal{A}, \mathcal{H})$. It is sufficient to prove that

$$\Delta cost_H + \Delta\Phi \;\leq\; \tfrac{1}{2}k(k+1)\Delta cost_{adv}, \tag{2}$$

where $\Delta cost_H$, $\Delta\Phi$, and $\Delta cost_{adv}$, denote, respectively, the expected cost of HARMONIC, the expected potential change, and the adversary cost, in the current step. By definition, $\Phi \geq 0$, and the theorem follows from (2) by routine amortization (see, for example, [BEY98]).

In our proof, we divide each step into two sub-steps: the adversary move and the HARMONIC's move, and we prove (2) separately for each sub-step.

Without loss of generality, we assume that in the matching which realizes the minimum in $\Phi(\mathcal{A}, \mathcal{H})$ before each sub-step considered, the adversary server j is matched to HARMONIC's server j', for each $j = 1, \ldots, k$, and thus $\Phi = \sum_{j=1}^{k} H_{jj'}$.

The adversary move. We can assume that in this sub-step the adversary moves just one server, say server 1. (Otherwise, we can move the servers one by one and analyze each move separately). Denoting by z the distance between the old and new location of server 1, we have $\Delta cost_{adv} = z$. The weight of the initial minimum matching changes by $\Delta H_{11'}$, where $\Delta H_{11'}$ denotes the change of the hitting cost $H_{11'}$. The potential after the move cannot be greater than the new weight of this initial matching, so $\Delta\Phi \leq |\Delta H_{11'}|$.

We think about this move as a continuous change of the position of server 1 along some line. Then we can view $H_{11'} = H_{11'}(z)$ as composed of two mappings: $z \to (d_{12}, d_{13}, \ldots, d_{1k}, d_{11'}) \to H_{11'}$. Since $|\partial d_{1i}/\partial z| \leq 1$ for all i, we have

$$\left|\frac{\partial H_{11'}}{\partial z}\right| \quad \leq \quad \left|\frac{\partial H_{11'}}{\partial d_{11'}}\right| + \sum_{i=2}^{k} \left|\frac{\partial H_{11'}}{\partial d_{1i}}\right| \quad \leq \quad \tfrac{1}{2}k(k+1),$$

using (A) with $s = 1$ and $t = 1'$. Therefore $\Delta\Phi \leq |\Delta H_{11'}| \leq \frac{1}{2}k(k+1)\Delta cost_{adv}$, and, since $\Delta cost_H = 0$, we conclude that (2) holds in the adversary move.

Harmonic's move. Without loss of generality, we assume that the adversary makes a request on server 1. We first claim that, for each $j = 2, \ldots, k$,

$$(c_{1j} - c_{1j'})(H_{11'} - H_{j1'}) \quad \leq \quad c_{1j'} H_{jj'}. \tag{3}$$

We distinguish two cases. If $c_{1j} \leq c_{1j'}$ then, using (B) with $v = 1$, $s = j$ and $t = 1'$ and (C) with $v = 1$, $s = j$, $t = j'$, we have

$$\begin{aligned}(c_{1j} - c_{1j'})(H_{11'} - H_{j1'}) &= (c_{1j'} - c_{1j})(H_{j1'} - H_{11'}) \\ &\leq (c_{1j'} - c_{1j})\widehat{H}_{j1} \\ &\leq c_{1j'} H_{jj'}.\end{aligned}$$

Otherwise $c_{1j} \geq c_{1j'}$. The minimality of the matching implies $H_{11'} + H_{jj'} \leq H_{1j'} + H_{j1'}$. Then

$$\begin{aligned}(c_{1j} - c_{1j'})(H_{11'} - H_{j1'}) &\leq (c_{1j} - c_{1j'})(H_{1j'} - H_{jj'}) \\ &= c_{1j'} H_{jj'} + [\,(c_{1j} - c_{1j'})H_{1j'} - c_{1j} H_{jj'}\,] \\ &\leq c_{1j'} H_{jj'},\end{aligned}$$

where the last inequality follows from inequality (D) with $v = 1$, $s = j$ and $t = j'$. This completes the proof of (3).

Now we prove (2). The equations for hitting costs (1), with $u = 1$ and $t = 1'$, can be rewritten as

$$k - c_{11'} H_{11'} \quad = \quad \sum_{j=2}^{k} c_{1j}(H_{11'} - H_{j1'}) \tag{4}$$

Let $\tilde{c} = \sum_{i=1}^{k} c_{1i'}$. The probability that Harmonic's server j' moves to 1 is $c_{1j'}/\tilde{c}$, and the cost of moving j' is $d_{1j'} = 1/c_{1j'}$. Thus $\Delta cost_H = k/\tilde{c}$. If $1'$ moves to 1, the potential decreases by $H_{11'}$. If $j' \neq 1'$ moves to 1, the potential increases by at most $H_{j1'} - H_{11'} - H_{jj'}$, since after the move we can match $1'$ with j and j' with 1. So

$$\begin{aligned}\tilde{c} \cdot (\Delta cost_H + \Delta\Phi) &\leq k + [\,-c_{11'}H_{11'} + \sum_{j=2}^{k} c_{1j'}(H_{j1'} - H_{11'} - H_{jj'})\,] \\ &= \sum_{j=2}^{k} \left[c_{1j}(H_{11'} - H_{j1'}) + c_{1j'}(H_{j1'} - H_{11'} - H_{jj'})\right] \\ &= \sum_{j=2}^{k} \left[(c_{1j} - c_{1j'})(H_{11'} - H_{j1'}) - c_{1j'} H_{jj'}\right] \\ &\leq 0,\end{aligned}$$

where the first equality follows from (4) and the last inequality from (3). Since $\Delta cost_{adv} = 0$, (2) holds when HARMONIC moves. This completes the proof of (2) and the theorem. $\square$

In configurations where exactly one of HARMONIC's servers is not covered by adversary servers, only a single edge in the minimum matching has non-zero cost, and Φ is equal to the cost of HARMONIC when the adversary makes only lazy requests (i.e., does not move his server). This is clearly a lower bound on a potential function, and in [BCNR02] it was shown that this potential is sufficient against the lazy adversary. Our potential is a natural generalization of this idea to all configurations.

3 Computing Hitting Costs

In this section we describe formulas for the hitting costs H_{it} as rational expressions in the conductances. We also introduce notation useful for their manipulation in subsequent analysis. Throughout this section, we assume that $G = \{1, 2, \ldots, k\}$ and $t \in M - G$. We treat $G + t$ as a complete graph with weights c_{ij} assigned to its edges (unordered pairs or nodes).

Admittance matrix. Let $\bar{c}_j = \sum_{i=1}^{k} c_{ij}$, for all j (where $c_{ii} = 0$ for all i). Define the *admittance matrix* C as

$$C \quad = \quad \begin{bmatrix} \bar{c}_1 + c_{1t} & -c_{12} & -c_{13} & \ldots & -c_{1k} \\ -c_{21} & \bar{c}_2 + c_{2t} & -c_{23} & \ldots & -c_{2k} \\ \vdots & & & \ddots & \vdots \\ -c_{k1} & -c_{k2} & -c_{k3} & \ldots & \bar{c}_k + c_{kt} \end{bmatrix}$$

Let $\overline{C}_j$ denote a matrix obtained from C by replacing column j by vector $\mathbf{1} = (1, 1, \ldots, 1)^T$. By C_{ij} we denote C with row i and column j removed. We have $\det(\overline{C}_j) = \sum_{i=1}^{k} (-1)^{i+j} \det(C_{ij})$. The equations for hitting costs (1) can be rewritten as $C\mathbf{H}_t = k\mathbf{1}$, where $\mathbf{H}_t = (H_{1t}, \ldots, H_{kt})^T$. Thus

$$H_{jt} \quad = \quad k \cdot \frac{\det(\overline{C}_j)}{\det(C)}. \tag{5}$$

ℓ-Trees and conductance products. An *ℓ-tree* is defined to be a forest with exactly ℓ connected components (trees). Any ℓ-tree has $k + 1 - \ell$ edges. If S is an ℓ-tree, then by r_{jS} we denote the size of the component of j in S.

For any set S of edges of $G + t$, let $\xi_S = \prod_{(i,j) \in S} c_{ij}$. The theorem below is a classical result from the theory of electrical networks (see [Per53, SR61]). It characterizes determinants and cofactors of admittance matrices in terms of certain sums of conductance products ξ_S.

Theorem 3.1 (a) $\det(C) = \sum_T \xi_T$, *where the sum is taken over all spanning trees T of $G + t$.*
(b) $\det(C_{ij}) = (-1)^{i+j} \sum_S \xi_S$, *where the sum is taken over all spanning 2-trees S of $G + t$ in which i, j are in one component and t in the other.*

Corollary 3.2 $\det(\overline{C}_j) = \sum_S r_{jS} \xi_S$, *where the sum is taken over all spanning 2-trees S of $G + t$ in which j and t are in different components.*

Proof: Applying Theorem 3.1.b, we have

$$\begin{aligned}\det(\overline{C}_j) &= \sum_{i=1}^{k}(-1)^{i+j}\det(C_{ij}) \\ &= \sum_{i=1}^{k}\sum_{S_i}\xi_{S_i},\end{aligned}$$

where, for each fixed i, the second summation is over all spanning 2-trees S_i of $G+t$ in which i,j are in one component and t in the other. Thus each ξ_S is counted r_{jS} times, and the corollary follows. □

Using Theorem 3.1 and Corollary 3.2, we can express the hitting costs H_{it} and $\widehat{H}_{ij}$ in terms of ℓ-tree conductance products.

Corollary 3.3 (a) *H_{it} can be written as*

$$H_{it} = k\cdot\frac{\sum_S r_{iS}\xi_S}{\sum_T \xi_T}, \tag{6}$$

where T ranges over all spanning trees of $G+t$ and S ranges over all spanning 2-trees S of $G+t$ in which i and t are in different components.
(b) *For $i \neq j$, $\widehat{H}_{ij}$ can be written as*

$$\widehat{H}_{ij} = k\cdot\frac{\sum_Y r_{iY}2^{\deg_Y(j)}\xi_Y}{\sum_Z 2^{\deg_Z(j)}\xi_Z}, \tag{7}$$

where Z ranges over all spanning trees of G, Y ranges over all spanning 2-trees of G that separate i from j, and $\deg_Q(u)$ denotes the degree of u in a graph Q.

Proof: Part (a) follows immediately from the formula (5) for hitting costs, Theorem 3.1.a, and Corollary 3.2.

It remains to prove (b). The proof is by letting t converge to j and applying (a). Let t be a copy of j and $c_{jt}\to\infty$. Then, using notation from (6), we have

$$\begin{aligned}H_{it} &= k\cdot\frac{\sum_S r_{iS}\xi_S}{\sum_T \xi_T} \\ &\to k\cdot\frac{\sum_{S\ni(j,t)} r_{iS}\xi_{S-(j,t)}}{\sum_{T\ni(j,t)}\xi_{T-(j,t)}} \\ &= k\cdot\frac{\sum_P r_{iP}\xi_P}{\sum_R \xi_R}\end{aligned}$$

where P ranges over all spanning 2-trees of $G+t$ that separate i,j from t, and R ranges over all spanning 2-trees of $G+t$ that separate j from t.

For each such R, by contracting j with t we get a spanning tree Z of G with $\xi_Z = \xi_R$. Further, each Z can be obtained from such a contraction in exactly $2^{\deg_Z(j)}$ ways. So $\sum_R \xi_R = \sum_Z 2^{\deg_Z(j)}\xi_Z$. By the same argument, $\sum_P r_{iP}\xi_P = \sum_Y r_{iY}2^{\deg_Y(j)}\xi_Y$, where Y is as in (7). This completes the proof of part (b). □

$\mathcal{X}$-Expressions. We use the symbol $\mathcal{X}\langle\cdot\rangle$ to denote sums of terms ξ_S or $r_{jS}\xi_S$ over various sets S specified by appropriate parameters of $\mathcal{X}\langle\cdot\rangle$. The argument of $\mathcal{X}\langle\cdot\rangle$ is a list of nodes separated into

$\ell \geq 0$ fields by symbols $|$ or $\|$. The sum in $\mathcal{X}\langle\cdot\rangle$ is taken over all spanning ℓ-trees of $G+t$ in which the nodes in different fields are in different components and the nodes in the same field are in the same component. The separator $\|$ appears at most once; and if present, each term ξ_S is multiplied by the total number of vertices in the components corresponding to the fields before $\|$. For example:

$\mathcal{X}\langle\rangle \ = \ \sum_S \xi_S$, where the sum is taken over all spanning trees S of $G+t$.

$\mathcal{X}\langle 1|2\rangle \ = \ \sum_S \xi_S$, where the sum is taken over all spanning 2-trees S of $G+t$ in which 1 and 2 are in different components.

$\mathcal{X}\langle 1,2\|3\rangle \ = \ \sum_S r_{1S}\xi_S$, where the sum is taken over all spanning 2-trees S of $G+t$ in which $1,2$ are in the same component and 3 in the other component.

$\mathcal{X}\langle 1|2\|3\rangle \ = \ \sum_S (r_{1S}+r_{2S})\xi_S$, where the sum is taken over all spanning 3-trees S of $G+t$ in which $1,2$ and 3 are in different components.

By Theorem 3.1 and Corollaries 3.2 and 3.3, we can express $\det(C)$, $\det(\overline{C}_j)$, and H_{it} in terms of $\mathcal{X}$-expressions as follows:

$$\det(C) \ = \ \mathcal{X}\langle\rangle \tag{8}$$

$$\det(\overline{C}_j) \ = \ \mathcal{X}\langle j\|t\rangle \tag{9}$$

$$H_{it} \ = \ k\cdot\frac{\mathcal{X}\langle i\|t\rangle}{\mathcal{X}\langle\rangle} \tag{10}$$

Partitioning $\mathcal{X}$-expressions. We now introduce two useful identities for $\mathcal{X}$-expressions. Let $i \neq j$ be two arbitrary nodes in $G+t$. Then:

$$\mathcal{X}\langle\rangle \ = \ c_{ij}\mathcal{X}\langle i|j\rangle + \sum_{\ell\neq i,j} c_{i\ell}\mathcal{X}\langle i|j,\ell\rangle \tag{11}$$

$$(k-1)\mathcal{X}\langle\rangle \ = \ \sum_{\ell\neq i} c_{i\ell}\mathcal{X}\langle \ell\|i\rangle \tag{12}$$

The equation (11) is obtained by grouping all spanning trees according to the first node $\ell \neq i$ on the path from i to j.

Equation (12) can be derived as follows. The LHS(12) is the sum of all terms $(k-1)\xi_T$, where T is a spanning tree of $G+t$. Given a spanning tree T, we think of T as being rooted at i. For each neighbor ℓ of i, let k_ℓ be the size of the subtree rooted at ℓ. Then $(k-1)\xi_T = \sum_{(i,\ell)\in T} k_\ell c_{i\ell}\xi_{T-(i,\ell)}$. Thus LHS(12) is the sum of the above expressions over all trees T. By exchanging the order of summations and grouping the terms according to $c_{i\ell}$, we obtain the RHS(12).

These two identities play crucial role in our proofs. They allow us to manipulate $\mathcal{X}$-expressions by breaking them into smaller expressions. Both (11) and (12) can be extended to more complex $\mathcal{X}$-expressions and, in fact, we will usually use these more general forms later in the paper. We refer to these two methods (11) and (12) of partitioning $\mathcal{X}$-expressions as *cut-partitioning* and *neighbor-partitioning*, respectively.

4 An Auxiliary Inequality

In this section we focus on inequality (A). We show that, for any k, inequality (A) can be reduced to a statement about certain tree products.

Lemma 4.1 *Let $s \in G$ and $t \notin G$. Suppose that for every $v \in G - s$,*

$$|\mathcal{X}\langle s|v\|t\rangle\mathcal{X}\langle\rangle - \mathcal{X}\langle s|v\rangle\mathcal{X}\langle s\|t\rangle| \;\; \leq \;\; \tfrac{1}{2}(k+1)\mathcal{X}\langle s|v,t\rangle\mathcal{X}\langle s|v\rangle. \tag{13}$$

Then inequality (A) *holds.*

Proof: By equation (10) we have $H_{st} = k \cdot \mathcal{X}\langle s\|t\rangle/\mathcal{X}\langle\rangle$. Thus inequality (A) can be written as

$$\sum_{v \in G+t-s} \left| \frac{\partial\mathcal{X}\langle s\|t\rangle}{\partial d_{sv}}\mathcal{X}\langle\rangle - \frac{\partial\mathcal{X}\langle\rangle}{\partial d_{sv}}\mathcal{X}\langle s\|t\rangle \right| \;\; \leq \;\; \tfrac{1}{2}(k+1)(\mathcal{X}\langle\rangle)^2. \tag{14}$$

Since $\partial c_{sv}/\partial d_{sv} = -c_{sv}^2$, for any set of edges Q we have

$$\frac{\partial \xi_Q}{\partial d_{sv}} \;\; = \;\; \begin{cases} -c_{sv}^2 \xi_{Q-(s,v)} & (s,v) \in Q \\ 0 & (s,v) \notin Q \end{cases},$$

and thus

$$\begin{aligned} \frac{\partial\mathcal{X}\langle\rangle}{\partial d_{sv}} &= -c_{sv}^2\mathcal{X}\langle s|v\rangle, \\ \frac{\partial\mathcal{X}\langle s\|t\rangle}{\partial d_{sv}} &= -c_{sv}^2\mathcal{X}\langle s|v\|t\rangle \qquad \text{for } v \neq t, \text{ and} \\ \frac{\partial\mathcal{X}\langle s\|t\rangle}{\partial d_{st}} &= 0. \end{aligned}$$

Substituting into (14) we obtain

$$\begin{aligned} \text{LHS(14)} &= c_{st}^2\mathcal{X}\langle s|t\rangle\mathcal{X}\langle s\|t\rangle + \sum_{v \in G-s} c_{sv}^2|\mathcal{X}\langle s|v\|t\rangle\mathcal{X}\langle\rangle - \mathcal{X}\langle s|v\rangle\mathcal{X}\langle s\|t\rangle| \\ &\leq \tfrac{1}{2}(k+1)c_{st}\mathcal{X}\langle s|t\rangle\mathcal{X}\langle\rangle + \sum_{v \in G-s} \tfrac{1}{2}(k+1)c_{sv}^2\mathcal{X}\langle s|v,t\rangle\mathcal{X}\langle s|v\rangle \\ &\leq \tfrac{1}{2}(k+1)\mathcal{X}\langle\rangle\left[c_{st}\mathcal{X}\langle s|t\rangle + \sum_{v \in G-s} c_{sv}\mathcal{X}\langle s|v,t\rangle\right] \\ &= \text{RHS(14)}. \end{aligned}$$

The first inequality follows from (13), and from the inequality $c_{st}\mathcal{X}\langle s\|t\rangle \leq \frac{1}{2}(k+1)\mathcal{X}\langle\rangle$ that was proved in [BCNR02] (see Theorem 2.b in [BCNR02]). The second inequality follows from $c_{sv}\mathcal{X}\langle s|v\rangle \leq \mathcal{X}\langle\rangle$. In the last step we used the equality $\mathcal{X}\langle\rangle = c_{st}\mathcal{X}\langle s|t\rangle + \sum_{v \in G-s} c_{sv}\mathcal{X}\langle s|v,t\rangle$, which is the cut-partitioning identity (11). $\square$

5 Proof of Inequality (D)

In this section we prove inequality (D), that is, $(c_{vs} - c_{vt})H_{vt} \leq c_{vs}H_{st}$. By applying the formula (10) for the hitting costs, multiplying by $\mathcal{X}\langle\rangle/k$, and applying the equation $\mathcal{X}\langle v\|t\rangle - \mathcal{X}\langle s\|t\rangle = \mathcal{X}\langle v\|s,t\rangle - \mathcal{X}\langle s\|v,t\rangle$, we reduce (D) to the following inequality:

$$c_{vs}(\mathcal{X}\langle v\|s,t\rangle - \mathcal{X}\langle s\|v,t\rangle) \;\; \leq \;\; c_{vt}\mathcal{X}\langle v\|t\rangle. \tag{15}$$

To prove (15), we prove the following, slightly stronger lemma.

Lemma 5.1 *Let $s, v \in G$ and $t \in M - G$. Then*

$$c_{vs}|\mathcal{X}\langle v\|s,t\rangle - \mathcal{X}\langle s\|v,t\rangle| \leq c_{vt}\mathcal{X}\langle v\|t,s\rangle \tag{16}$$

Proof: The proof is by induction on k, the size of G. In the base step, for $k = 2$, inequality (16) reduces to the triangle inequality $c_{vs}|c_{st} - c_{vt}| \leq c_{vt}c_{st}$.

Suppose now that $k \geq 3$ and that the lemma holds for all values $k' < k$. First, we claim that the following identities hold:

$$\mathcal{X}\langle v\|s,t\rangle = c_{st}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{ss'}\mathcal{X}\big\langle v\|s',t|s\big\rangle \tag{17}$$

$$\mathcal{X}\langle s\|v,t\rangle = \mathcal{X}\langle v,t|s\rangle + \sum_{s'\in G-v-s} c_{ss'}\mathcal{X}\big\langle s'\|v,t|s\big\rangle \tag{18}$$

$$c_{vt}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{vs'}\mathcal{X}\big\langle v\|s',t|s\big\rangle = \mathcal{X}\langle v,t|s\rangle + \sum_{s'\in G-v-s} c_{vs'}\mathcal{X}\big\langle s'\|v,t|s\big\rangle \tag{19}$$

The equality (17) is obtained using the cut-partitioning formula (11), applied to expression $\mathcal{X}\langle v\|s,t\rangle$, with s' denoting the first node on the path from v to s. Similarly, the equality (18) is obtained using the neighbor-partitioning formula (12), applied to expression $\mathcal{X}\langle v\|s,t\rangle$, with s' ranging over the neighbors of s.

The justification for (19) is a bit more complicated. It is again derived using cut- and neighbor-partitioning, but this time the expression we partition is not expressible using the $\mathcal{X}$-notation. For a 2-tree S that separates v, t from s, denote by q_S the size of the component of v in a 3-tree obtained from S by removing the first edge on the path from v to t. We obtain (19) by partitioning the expression $\sum_S q_S\xi_S$ (where S ranges over all 2-trees separating v, t from s) in two different ways: the LHS(19) is obtained using cut-partitioning, and RHS(19) is obtained using neighbor-partitioning.

We derive (16) in several steps:

$$\begin{aligned}
\text{LHS(16)} &= \Big| c_{vs}c_{st}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{vs}c_{ss'}\big(\mathcal{X}\big\langle v\|s',t|s\big\rangle - \mathcal{X}\big\langle s'\|v,t|s\big\rangle\big) - c_{vs}\mathcal{X}\langle v,t|s\rangle \Big| && (20)\\
&\leq \Big| c_{vs}c_{vt}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{vs}c_{vs'}\big(\mathcal{X}\big\langle v\|s',t|s\big\rangle - \mathcal{X}\big\langle s'\|v,t|s\big\rangle\big) - c_{vs}\mathcal{X}\langle v,t|s\rangle \Big| && (21)\\
&\qquad + c_{vt}c_{st}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{vs'}c_{ss'}\big|\mathcal{X}\big\langle v\|s',t|s\big\rangle - \mathcal{X}\big\langle s'\|v,t|s\big\rangle\big| \\
&\leq c_{vt}c_{st}\mathcal{X}\langle v\|t|s\rangle + \sum_{s'\in G-v-s} c_{ss'}c_{vt}\mathcal{X}\langle v\|t|s\rangle && (22)\\
&= \text{RHS(16)} && (23)
\end{aligned}$$

We now explain the steps in the above derivation. The first two steps are straighforward. The equality (20) is obtained by substituting (17) and (18) into LHS(16).

In step (21) we use triangle inequalities. This need to be done with care, because applying a triangle inequality inside an absolute value may increase or decrease the value of the whole expression. For the first term, we apply $|c_{vs}c_{st} - c_{vs}c_{vt}| \leq c_{vt}c_{st}$ as follows: this triangle inequality implies that $|\alpha c_{vs}c_{st} + \beta| \leq |\alpha c_{vs}c_{vt} + \beta| + |\alpha|c_{vt}c_{st}$ for any reals α and β. We apply this form with $\alpha = \mathcal{X}\langle v\|t|s\rangle$ and with β equal to the sum of the remaining terms within the absolute value. Similarly, for each term in the sum, one at a time, we use $|c_{vs}c_{ss'} - c_{vs}c_{vs'}| \leq c_{vs'}c_{ss'}$ as follows: this triangle inequality implies that $|\alpha c_{vs} c_{vs'} + \beta| \leq |\alpha c_{vs} c_{vs'} + \beta| + |\alpha| c_{vs'} c_{ss'}$ for any reals α and β. We apply this form with $\alpha = \mathcal{X}\langle v \| s', t|s\rangle - \mathcal{X}\langle s' \| v, t|s\rangle$ and β equal to the sum of the remaining terms within the absolute value.

In step (22), all terms within the absolute value in the first line cancel, by identity (19). Then we use the inductive assumption on each element of the sum in the second line of (21), with the component of s removed. For every fixed tree in the component of s, both terms contain the same factor of ξ_S and the remainder is just the left-hand side of LHS(16) for a smaller set of vertices (note that we removed at least one node, namely s).

Finally, the equality (23) is obtained by applying yet again the cut-partitioning method (11). This completes the inductive step and the proof of the lemma. □

6 Analysis of Harmonic for $k = 3$

We now use the framework developed in the previous sections to prove that Harmonic is 6-competitive for 3 servers. We prove that the inequalities (A)–(C) hold for $k = 3$.

We remark that the inequalities that do not require the triangle inequality, for example (24), can be easily verified in Maple or Mathematica. Other inequalities can also be verified in this manner, but this process is complicated by the fact that the choice of the triangle inequalities and the order of their application is not always easy to determine.

Throughout this section, to simplify the notation, we assume that $G = \{1, 2, 3\}$, $s = 1$, $v = 2$ and $t \notin G$.

Proof of inequality (A) for $k = 3$. By Lemma 4.1, it is sufficient to prove that

$$|\mathcal{X}\langle 1|2\|t\rangle \mathcal{X}\langle\rangle - \mathcal{X}\langle 1|2\rangle \mathcal{X}\langle 1\|t\rangle| \quad \leq \quad 2\mathcal{X}\langle 1|2, t\rangle \mathcal{X}\langle 1|2\rangle. \tag{24}$$

As noted earlier, this inequality can be verified without using any triangle inequality, i.e., only by comparing the coefficients of equal terms on both sides. The terms containing c_{12} on left-hand side cancel (this is in fact true for any k). Then we deal separately with those terms that contain c_{1t} and those that don't.

We use subscripts $i \not\sim j$ in $\mathcal{X}\langle\cdot\rangle$ to restrict the summation range to those S in which i is not adjacent to j. For example, $\mathcal{X}_{1\not\sim 2}\langle 1|t\rangle$ is the sum of all ξ_S, where S is a 2-tree with $1, t$ in different components and with no edge $(1, 2)$.

We claim that

$$\begin{aligned}
|\, \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1|t\rangle - \mathcal{X}\langle 1|2|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle \,| \quad &\leq \quad 2\mathcal{X}\langle 1|2, t\rangle \mathcal{X}\langle 1|2|t\rangle & (25)\\
|\, \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}_{1\not\sim 2,t}\langle\rangle - \mathcal{X}_{1\not\sim t}\langle 1|2\rangle \mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle \,| \quad &\leq \quad 2\mathcal{X}\langle 1|2, t\rangle \mathcal{X}_{1\not\sim t}\langle 1|2\rangle & (26)
\end{aligned}$$

We first derive inequality (25). We have $\mathcal{X}_{1\not\sim 2}\langle 1|t\rangle = \mathcal{X}\langle 1|2, t\rangle + \mathcal{X}_{1\not\sim 2}\langle 1, 2|t\rangle$ and a similar equality holds for $\mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle$. By substituting and canceling some terms, we get

$$\begin{aligned}
\text{LHS}(25) \quad &= \quad |\, \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}\langle 1|2, t\rangle + \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1, 2|t\rangle \\
&\qquad - \mathcal{X}\langle 1|2|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1, 2\|t\rangle - \mathcal{X}\langle 1|2|t\rangle \mathcal{X}\langle 1\|2, t\rangle \,| \\
&= \quad |\, \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}\langle 1|2, t\rangle - \mathcal{X}\langle 1|2|t\rangle \mathcal{X}\langle 1\|2, t\rangle - c_{13}c_{23}c_{3t} \,| & (27)\\
&\leq \quad \text{RHS}(25)
\end{aligned}$$

The last inequality can be justified as follows. All the non-zero terms in the absolute value in (27) and in RHS(25) contain exactly the same products $\xi_S \xi_R$ (where S and R are 3-trees and 2-trees), but with possibly different coefficients. In (27), all coefficients in $\mathcal{X}\langle 1|2\|t\rangle \mathcal{X}\langle 1|2,t\rangle$ are 1, 2 or 3. All coefficients in $\mathcal{X}\langle 1|2|t\rangle \mathcal{X}\langle 1\|2,t\rangle$ are 1 or 2. So in the difference (taking into account the term $c_{13}c_{23}c_{3t}$) each coefficient is between -2 and 2. Consequently, (27) is bounded by RHS(25).

Now we prove (26). Let $\Upsilon = c_{2t}c_{3t} + c_{2t}c_{23} + c_{23}c_{3t}$. In other words, Υ is the sum of the ξ_S over the spanning trees of $\{2,3,t\}$. Then

$$\begin{aligned}
\mathcal{X}\langle 1|2\|t\rangle &= 3c_{13} + 3c_{23} + 2c_{3t} \\
\mathcal{X}_{1\not\sim 2,t}\langle\rangle &= c_{13}\Upsilon \\
\mathcal{X}_{1\not\sim t}\langle 1|2\rangle &= c_{13}(c_{3t} + c_{2t}) + \Upsilon \\
\mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle &= c_{13}(3c_{23} + 2c_{2t}) + \Upsilon \\
\mathcal{X}\langle 1|2,t\rangle &= c_{13}c_{2t} + \Upsilon
\end{aligned}$$

Estimating the absolute value of a sum by a sum of absolute values, we have

$$\begin{aligned}
\text{LHS}(26) &= \left|c_{13}^2 c_{2t}(c_{3t} - 2c_{2t}) + c_{13}\Upsilon(c_{3t} - 3c_{2t}) - \Upsilon^2\right| \\
&\leq 2[\, c_{13}^2 c_{2t}(c_{3t} + c_{2t}) + c_{13}\Upsilon(c_{3t} + 2c_{2t}) + \Upsilon^2\,] \\
&= \text{RHS}(26)
\end{aligned}$$

It remains to show how inequalities (25) and (26) imply inequality (24). Substituting

$$\begin{aligned}
\mathcal{X}\langle\rangle &= c_{12}\mathcal{X}\langle 1|2\rangle + c_{1t}\mathcal{X}_{1\not\sim 2}\langle 1|t\rangle + \mathcal{X}_{1\not\sim 2,t}\langle\rangle \\
\mathcal{X}\langle 1|2\rangle &= c_{1t}\mathcal{X}\langle 1|2|t\rangle + \mathcal{X}_{1\not\sim t}\langle 1|2\rangle \\
\mathcal{X}\langle 1\|t\rangle &= c_{12}\mathcal{X}\langle 1|2\|t\rangle + \mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle
\end{aligned}$$

into LHS(24) and canceling, we get

$$\begin{aligned}
\text{LHS}(24) &= |c_{1t}\mathcal{X}\langle 1|2\|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1|t\rangle + \mathcal{X}\langle 1|2\|t\rangle \mathcal{X}_{1\not\sim 2,t}\langle\rangle \\
&\quad - c_{1t}\mathcal{X}\langle 1|2|t\rangle \mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle - \mathcal{X}_{1\not\sim t}\langle 1|2\rangle \mathcal{X}_{1\not\sim 2}\langle 1\|t\rangle| \\
&\leq 2c_{1t}\mathcal{X}\langle 1|2,t\rangle \mathcal{X}\langle 1|2|t\rangle + 2\mathcal{X}\langle 1|2,t\rangle \mathcal{X}_{1\not\sim t}\langle 1|2\rangle \\
&= 2\mathcal{X}\langle 1|2,t\rangle[\, c_{1t}\mathcal{X}\langle 1|2|t\rangle + \mathcal{X}_{1\not\sim t}\langle 1|2\rangle\,] \\
&= \text{RHS}(24).
\end{aligned}$$

This completes the proof of (24) and the key inequality (A) for $k = 3$. $\square$

Proof of inequality (C) for $k = 3$. We need to prove that

$$(c_{2t} - c_{12})\widehat{H}_{12} \quad \leq \quad c_{2t}H_{1t}. \tag{28}$$

Using Lemma 3.3 and substituting (see Corollary 3.3)

$$\begin{aligned}
H_{1t} &= 3\mathcal{X}\langle 1\|t\rangle / \mathcal{X}\langle\rangle \quad \text{and} \\
\widehat{H}_{12} &= 3(c_{13} + c_{23})/(2c_{12}c_{23} + c_{12}c_{13} + c_{13}c_{23}),
\end{aligned}$$

we reduce (28) to

$$(c_{2t} - c_{12})(c_{13} + c_{23})\mathcal{X}\langle\rangle \quad \le \quad c_{2t}(2c_{12}c_{23} + c_{12}c_{13} + c_{13}c_{23})\mathcal{X}\langle 1\|t\rangle. \tag{29}$$

Let $\Upsilon = c_{12}c_{23} + c_{12}c_{13} + c_{23}c_{13}$ and $\Upsilon' = c_{2t}c_{23} + c_{2t}c_{3t} + c_{23}c_{3t}$. Note that

$$\mathcal{X}\langle 1|t\rangle \quad = \quad \Upsilon + c_{12}c_{3t} + c_{13}c_{2t} + \Upsilon' \tag{30}$$

$$\mathcal{X}\langle 1\|t\rangle \quad = \quad 3\Upsilon + 2(c_{12}c_{3t} + c_{13}c_{2t}) + \Upsilon' \;=\; \mathcal{X}\langle 1|t\rangle + 2\Upsilon + c_{12}c_{3t} + c_{13}c_{2t} \tag{31}$$

$$\mathcal{X}\langle\rangle \quad = \quad c_{1t}\mathcal{X}\langle 1|t\rangle + (c_{12} + c_{13})\Upsilon' + c_{12}c_{13}(c_{2t} + c_{3t}) \tag{32}$$

We first estimate the two sides of (29). Substituting (32) into LHS(29), rearranging, using the triangle inequalities $(c_{2t} - c_{12})c_{1t} \le c_{2t}c_{12}$, $c_{12}(c_{13} + c_{23}) \ge c_{13}c_{23}$, and dropping some negative terms, we get

$$\begin{aligned}
\mathrm{LHS}(29) \quad &= \quad (c_{13} + c_{23})[(c_{2t} - c_{12})c_{1t}\mathcal{X}\langle 1|t\rangle + c_{2t}(c_{12} + c_{13})\Upsilon' + c_{12}c_{13}c_{2t}^2] \\
&\quad - c_{12}c_{3t}(c_{13} + c_{23})[c_{12}(c_{13} + c_{23}) + c_{13}c_{23}] \\
&\quad - c_{12}c_{2t}(c_{13} + c_{23})(c_{12}c_{13} + c_{12}c_{23} + c_{12}c_{3t} + c_{13}c_{23}) \\
&\le \quad c_{2t}(c_{13} + c_{23})[\, c_{12}\mathcal{X}\langle 1|t\rangle + (c_{12} + c_{13})\Upsilon' + c_{12}c_{13}c_{2t}\,] - 2c_{13}^2c_{23}^2c_{3t}.
\end{aligned} \tag{33}$$

Substituting (31) into RHS(29), rearranging, and then dropping some positive terms, we get

$$\begin{aligned}
\mathrm{RHS}(29) \quad &= \quad c_{2t}c_{12}(c_{13} + c_{23})\mathcal{X}\langle 1\|t\rangle + c_{23}c_{2t}(c_{12} + c_{13})\mathcal{X}\langle 1\|t\rangle \\
&\ge \quad c_{2t}c_{12}(c_{13} + c_{23})[\,\mathcal{X}\langle 1|t\rangle + c_{13}c_{2t}\,] + c_{23}c_{2t}(c_{12} + c_{13})\Upsilon' \\
&\qquad + 2c_{13}c_{23}c_{2t}(c_{13}c_{2t} + c_{13}c_{23} + c_{12}c_{3t} + c_{12}c_{2t})
\end{aligned} \tag{34}$$

Now, from (33) and (34), we obtain

$$\begin{aligned}
\mathrm{LHS}(29) - \mathrm{RHS}(29) \quad &\le \quad c_{13}[c_{2t}(c_{12} + c_{13})\Upsilon' - 2c_{13}c_{23}^2c_{3t} \\
&\qquad \text{-}2c_{23}c_{2t}(c_{13}c_{2t} + c_{13}c_{23} + c_{12}c_{3t} + c_{12}c_{2t})] \\
&= \quad c_{13}(c_{12}c_{2t} + c_{13}c_{2t} + 2c_{13}c_{23})(c_{2t}c_{3t} - c_{23}c_{2t} - c_{23}c_{3t}) \;\le\; 0.
\end{aligned}$$

This completes the proof of (29) and inequality (C) for $k = 3$. □

Proof of inequality (B) for $k = 3$. In this derivation, in order to make calculations easier to follow, we represent the tree products graphically as follows. Label the nodes as

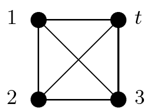

.

In the calculations, we represent each ξ_S by a subgraph consisting of the edges whose conductances appear in S. For example,  $= c_{12}c_{23}c_{2t}$,  $= c_{1t}c_{23}c_{3t}$,  $= c_{13}c_{2t}c_{3t}$, etc.

In terms of conductances, the triangle inequality takes the form $c_{xy}c_{xz} \le c_{xy}c_{yz} + c_{xz}c_{yz}$. For example, in the convention described above, the application of the triangle inequality $c_{12}c_{23}c_{2t} \le c_{1t}c_{23}c_{12} + c_{1t}c_{23}c_{2t}$ can be expressed as  $\le$  $+$ .

We need to show that:

$$H_{1t} - H_{2t} \quad \le \quad \widehat{H}_{12} \tag{35}$$

Substituting (see Corollary 3.3)

$$
\begin{aligned}
H_{1t} &= 3\mathcal{X}\langle 1\|t\rangle / \mathcal{X}\langle\rangle, \\
H_{2t} &= 3\mathcal{X}\langle 2\|t\rangle / \mathcal{X}\langle\rangle, \\
\widehat{H}_{12} &= 3(c_{13} + c_{23})/(2c_{12}c_{23} + c_{12}c_{13} + c_{13}c_{23}), \quad \text{and} \\
\mathcal{X}\langle 1\|t\rangle - \mathcal{X}\langle 2\|t\rangle &= \mathcal{X}\langle 1\|2,t\rangle - \mathcal{X}\langle 2\|1,t\rangle,
\end{aligned}
$$

we reduce inequality (35) to:

$$(2c_{12}c_{23} + c_{12}c_{13} + c_{13}c_{23})[\mathcal{X}\langle 1\|2,t\rangle - \mathcal{X}\langle 2\|1,t\rangle] \;\leq\; (c_{13} + c_{23})\mathcal{X}\langle\rangle. \tag{36}$$

We now prove (36). The idea of the proof is to represent each product on the left-hand side as a product of c_{13} or c_{23} times some product ξ_T, for a spanning tree T of $\{1, 2, 3, t\}$. Using our graphical notation, we have

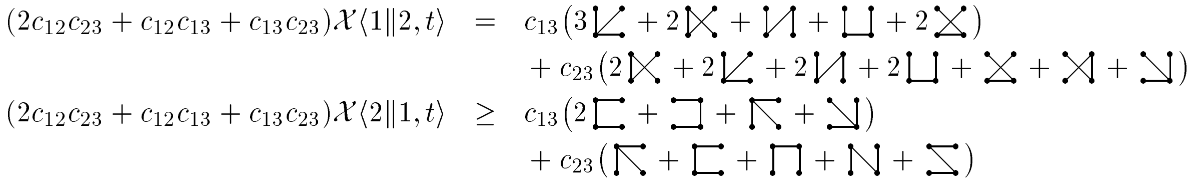

(37)

(38)

Substituting (37), (38) into (36) and applying some triangle inequalities, we get:

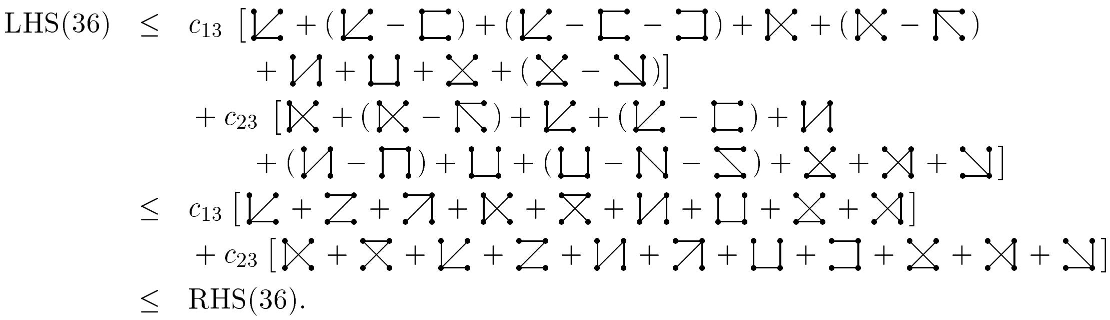

This completes the proof of (36) and inequality (B) for $k = 3$. □

We have now proved inequalities (A), (B) and (C) for $k = 3$. As shown in Section 5, (D) holds for all k. Therefore, from Theorem 2.1, we obtain the main result of this paper:

Theorem 6.1 *For $k = 3$,* Harmonic *is 6-competitive.*

7 Final Comments

We proved that Harmonic is 6-competitive for 3 servers. It is quite easy to verify that inequalities (A)–(C) also hold for $k = 2$, giving yet another proof that Harmonic is 3-competitive for $k = 2$. We believe that inequalities (A) and (B) are valid for arbitrary k, and that our potential can be applied to show that Harmonic is $\frac{1}{2}k(k+1)$-competitive in other cases, and quite possibly for the general case as well. However, we can also show that inequality (C) fails for large values of k, and thus even if Φ works in general, the proof will require a different approach.

The inequalities (A) and (B) represent interesting properties of harmonic random walks and they are of independent interest. They also may be applicable to obtain other results on HARMONIC. For example, inequality (B) can be thought of as a triangle inequality for harmonic random walks. Proving inequality (A) alone would show that HARMONIC is $\frac{1}{2}k(k+1)$ competitive against a *semi-lazy* adversary, which is like the lazy adversary with the additional freedom of moving the server that is not covered by our servers.

References

[BCNR02] Yair Bartal, Marek Chrobak, John Noga, and Prabhakar Raghavan. More on random walks, electrical networks, and the harmonic k-server algorithm. To appear in IPL, 2002.

[BEY98] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[BG00] Yair Bartal and Edward Grove. The harmonic k-server algorithm is competitive. *Journal of the ACM*, 47(1):1–15, 2000.

[BKT90] Piotr Berman, Howard Karloff, and Gabor Tardos. A competitive algorithm for three servers. In *Proc. 1st Symp. on Discrete Algorithms*, pages 280–290, 1990.

[CDRS93] Don Coppersmith, Peter G. Doyle, Prabhakar Raghavan, and Marc Snir. Random walks on weighted graphs and applications to on-line algorithms. *Journal of the ACM*, 40:421–453, 1993.

[CL92] Marek Chrobak and Lawrence L. Larmore. HARMONIC is three-competitive for two servers. *Theoretical Computer Science*, 98:339–346, 1992.

[CL98] Marek Chrobak and Lawrence L. Larmore. Metrical task systems, the server problem, and the work function algorithm. In *Online Algorithms: State of the Art*, pages 74–94. Springer-Verlag, 1998.

[CRR+89] Ashok K. Chandra, Prabhakar Raghavan, Walter L. Ruzzo, Roman Smolensky, and Prasoon Tiwari. The electrical resistance of a graph captures its commute and cover times. In *Proc. 21st Symp. Theory of Computing*, pages 574–586, 1989.

[CS00] Marek Chrobak and Jiří Sgall. A simple analysis of the harmonic algorithm for two servers. *Information Processing Letters*, 75:75–77, 2000.

[DS84] Peter G. Doyle and J. Laurie Snell. *Random Walks and Electrical Networks*. Mathematical Association of America, 1984.

[KP95] Elias Koutsoupias and Christos Papadimitriou. On the k-server conjecture. *Journal of the ACM*, 42:971–983, 1995.

[MMS90] Mark Manasse, Lyle A. McGeoch, and Daniel Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11:208–230, 1990.

[Per53] W.S. Percival. Solution of passive electrical networks by means of mathematical trees. *J. Inst. Elect. Engrs.*, 100:143–150, 1953.

[RS94] Prabhakar Raghavan and Marc Snir. Memory versus randomization in on-line algorithms. *IBM Journal on Research and Development*, 38, 1994.

[SR61] Sundaram Seshu and Myril B. Reed. *Linear Graphs and Electrical Networks*. Addison-Wesley, 1961.