

The Load-Distance Balancing Problem

Edward Bortnikov* Samir Khuller* Yishay Mansour^{†‡} Joseph (Seffi) Naor[◊]

*Yahoo! Research, Matam Park, Haifa 31905 (Israel)

*Department of Computer Science, University of Maryland, College Park, MD 20742 (USA)

[†]School of Computer Science, Tel-Aviv University, Tel-Aviv (Israel)

[‡]Google Research, New York, NY 10011 (USA)

[◊]Computer Science Department, Technion, Haifa 32000 (Israel)

Abstract

Problems dealing with assignment of clients to servers have been widely studied. However, they usually do not model the fact that the delay incurred by a client is a function of both the distance to the assigned server and the load on this server, under a given assignment. We study a problem referred to as the Load-Distance Balancing problem (or LDB), where the objective is assigning a set of clients to a set of given servers. Each client suffers a delay that is the sum of the distance to its server and the congestion delay at this server, a non-decreasing function of the number of clients assigned to the server.

We address two flavors of LDB – the first one seeking to minimize the *maximum* incurred delay, and the second one targeted for minimizing the *average* delay. For the first variation, we present hardness results, a best possible approximation algorithm, and an optimal algorithm for a special case of linear placement of clients and servers. For the second one, an optimal polynomial-time algorithm is presented.

Keywords: *Approximation Algorithms, Facility Assignment*

1 Introduction

The ever-increasing demand for large-scale real-time services to geographically dispersed user populations motivates access providers to deploy advanced services close to the network's edge. Consider, for example, wireless mesh networking (WMN) technologies, which are poised to be a next-generation platform for high-speed Internet access in urban and rural areas [1]. A WMN infrastructure consists of numerous wireless routers, which jointly forward the user traffic to (and from) a limited set of landline gateways. Today, these gateways mainly provide bandwidth sharing of their high-speed access links to the WMN users. Tomorrow, they can be envisioned as a platform for rolling out application-level services with stringent quality-of-service (QoS) requirements. For example, we foresee WMN gateways playing the role of VoIP traffic gateways, media content delivery caches, and even online game servers [4].

In a multi-server setting, *service assignment* problems naturally arise. In this context, each client session must be assigned to an application-level server. Assignment problems have been widely studied in operations research and computer science, and classical problems model the cost of assigning clients to servers as a sum of fixed client-server distances and server (facility) costs. In this context, the servers might or might not have capacities. We identify a need for a more realistic model for describing the *end-user* QoS, e.g., service delay. We model the service delay of a client session as a sum of a *network delay*,

incurred by the network connecting the user to its server, and a *congestion delay*, caused by queuing and processing at the assigned server. The delay experienced by each end user is the sum of two quantities: the distance to the assigned server, and the delay incurred at the server. The *load-distance balancing* problem, or LDB, seeks to balance between these two factors, in order to minimize the service delay among all clients. It has two flavors: (1) *maximum* delay minimization and (2) *average* delay minimization.

Summary of Results: We demonstrate that the min-max LDB problem is NP-hard and present an approximation algorithm with a factor of 2. We also show that the problem is non-approximable with a factor better than 2 for general distance and load functions assuming $P \neq NP$. In addition, we are able to show that for metric spaces (where triangle inequality is satisfied by the distance function) we cannot obtain an approximation factor better than $\frac{5}{3}$ unless $P = NP$. For the special case when the users and the servers are located on a line segment with Euclidean network distances, we present a polynomial time dynamic programming algorithm for this problem. In the sequel, we present a polynomial algorithm for min-average LDB, which applies for convex load functions, and a dynamic-programming solution for the linear setting which has an improved time complexity.

Related Work: The min-max LDB problem has been introduced in [3]. That work concentrated on solving the problem in a *distributed* setting, in which the servers jointly compute the assignment with partial local data. The protocol of [3] can use any sequential algorithm as a building block. In particular, it can use **BFlow**, the best possible approximation algorithm. Our paper studies the LDB problem in a broader context, and presents new problem variations, algorithms, and hardness results.

Related min-max problems dealing with capacities and facility location were studied before [2, 5, 6]. For example, the capacitated K -center problem [2, 5] asks for K locations to be designated as centers, so as to minimize the maximum distance of a node from its assigned center. In the basic K -center problem, there are no capacities and a center can be assigned an arbitrary number of clients. In the capacitated version each center has a (uniform) load capacity of L , and thus each center can have at most L clients assigned to it. In a sense, this guarantees a bound on the delay of any client, since each is within a distance $O(d^*)$ of its assigned center (d^* is the optimal radius) and cannot suffer a long service time at the assigned center due to the load being at most L . In [5], a 5-approximation on the distance measure was presented for the capacitated K -center problem (improving on a previous bound of 10 [2]) and in addition a $(\frac{2}{c}K, cL, 2R)$ solution was presented where $c = 1 + \epsilon$ for any $0 < \epsilon < 1$. This is a solution that uses more than K centers, and allows higher than L load, yet provides a better approximation guarantee.

2 Problem Definition

Consider a set of servers $S = \{s_1, \dots, s_k\}$ and a set of clients $U = \{u_1, \dots, u_n\}$, so that $k \ll n$. The *network delay* function $D : (U \times S) \rightarrow \mathbb{R}^+$ captures the network distance between a client and a server. This function is not necessarily subject to the triangle inequality.

Consider an assignment $\lambda : U \rightarrow S$ that maps every client to a server. We assume that each client u assigned to server s adds a unit of *load* on s . We denote the load on s as $\mathcal{L}(\lambda, s) \triangleq |\{u : \lambda(u) = s\}|$. We shorten this to $\mathcal{L}(s)$ when the assignment function is clear from the context. A monotonic non-decreasing *congestion delay* function, $\delta_s : \mathbb{N} \rightarrow \mathbb{R}^+$, captures the delay incurred by server s as a function of the number of assigned clients. Different servers can have different congestion delay functions. The service delay $\Delta(u, \lambda)$ of session u in assignment λ is the sum of the two delays:

$$\Delta(u, \lambda) \triangleq D(u, \lambda(u)) + \delta_{\lambda(u)}(\mathcal{L}(\lambda, \lambda(u))).$$

The *maximum* (resp., *average*) cost of an assignment λ is the maximum (resp., average) delay it incurs for a client: $\Delta^M(\lambda(U)) \triangleq \max_{u \in U} \Delta(u, \lambda)$, and $\Delta^A(\lambda(U)) \triangleq \frac{1}{n} \sum_{u \in U} \Delta(u, \lambda)$.

The min-max (resp., min-average) load-distance balancing assignment problem (or LDB in short) is to find an assignment λ^* such that $\Delta^M(\lambda^*(U))$ (resp., $\Delta^A(\lambda^*(U))$) is minimized. An assignment that yields the minimum cost is called *optimal*.

3 Min-Max Load-Distance Balancing

3.1 NP-Hardness

We prove that the decision version of min-max LDB problem is NP-hard (referred to as LDB-D). We consider the problem of deciding whether delay Δ^* is feasible, i.e., $\Delta^M(\lambda(U)) \leq \Delta^*$. In what follows, we show a reduction from the classical *exact set cover* (XSC) problem. An instance of XSC is a collection S of subsets over a finite set U . A solution $S' \subseteq S$ is a cover for U , i.e., every element in U belongs to at least one member of S' . The decision problem is whether there is a cover such that each element belongs to precisely one set in the cover.

Theorem 1 *The Min-Max LDB problem is NP-hard.*

Proof: Consider an instance of XSC in which $|U| = n$, $|S| = k$, and each set contains exactly m elements. The problem is therefore whether there is a cover containing $\frac{n}{m}$ sets.

The transformation of this instance to an instance of LDB-D is as follows. In addition to the elements in U , we define a set U' of $M(k - \frac{n}{m})$ dummy elements, where $M > m$. We construct a bipartite graph, in which the left side contains the elements in $U \cup U'$ (the clients), and the right side contains the sets in S (the servers). The dummy clients are at distance d_1 from each server. The real clients (elements) are at distance $d_2 > d_1$ from each server (set) that covers them, and at distance ∞ from all the other servers. The capacity of each server for distance d_1 is M , and for distance d_2 is m , i.e., $\delta_s^{-1}(\Delta^* - d_1) = M$, and $\delta_s^{-1}(\Delta^* - d_2) = m$. In other words, the delay at a server for load at most m is $\Delta^* - d_2$ and for load at most M is $\Delta^* - d_1$. It is easy to see that under a feasible assignment, no client's delay exceeds Δ^* .

Each server can cover either M dummy clients, or any combination of $0 < m' \leq m$ original clients and $m - m'$ dummy clients. If both real and dummy clients are assigned to at least one server, the total number of servers that have real clients assigned to them is $k' > \frac{n}{m}$. All these servers have capacity m , and hence, they serve at most $mk' - n$ dummy clients. The remaining servers can host $M(k - k')$ dummy clients. Hence, the total number of assigned dummy clients is bounded by $M(k - k') + mk' - n = M(k - \frac{n}{m}) - M(k' - \frac{n}{m}) + m(k' - \frac{n}{m}) < M(k - \frac{n}{m})$, that is, the assignment is not feasible. Hence, exactly $\frac{n}{m}$ servers must be allocated to real clients, thus solving the XSC instance. \square

Hardness of Approximation: We simply specify some of the parameters in the above reduction to obtain the hardness results. In particular, we show that if there is a solution to the exact cover problem, then there is a solution for the load-balance problem with cost Δ^* . If there is no solution to the exact cover problem, then all solutions for LDB have a high cost of $(2(\Delta^* - \epsilon))$, or $\frac{5}{3}\Delta^*$ in metric spaces).

Consider $d_1 = \epsilon$ and $d_2 = \Delta^* - \epsilon$. If an element is not a member of a set, the distance to that server is very high. If there is no solution for exact cover, then any collection of $\frac{n}{m}$ sets will leave some element uncovered. The corresponding client will have to be assigned to a server that is also serving $M - 1$ dummy clients. The delay experienced by this client is thus $d_2 + (\Delta^* - d_1)$. This could be as high as $2(\Delta^* - \epsilon)$. Note, however, that this distance function does not satisfy triangle inequality.

The choice of $d_1 = \frac{\Delta^*}{3}$ and $d_2 = \Delta^*$ preserves the triangle inequality. The distance of a client in C' to a server is either Δ^* or $\frac{5}{3}\Delta^*$. If there is no solution to exact set cover, then the best assignment can have delay no lower than $\frac{5}{3}\Delta^*$.

3.2 A 2-Approximation Algorithm

We now present an algorithm, called BFlow, which computes a 2-approximate solution for min-max LDB. Algorithm BFlow works in phases; in each phase it guesses $\Delta^* = \Delta^M(\lambda^*(U))$, and checks the feasibility of a specific assignment in which neither the network nor the congestion delay exceeds Δ^* , and hence its cost is bounded by $2\Delta^*$. BFlow performs a binary search on the value of Δ^* . A single phase is as follows:

1. Each client u marks all servers s that are at distance $D(u, s) \leq \Delta^*$. These are its feasible servers.
2. Each server s announces how many clients it can serve by computing the inverse of $\delta_s^{-1}(\Delta^*)$.

- Define a bipartite client-server graph where an edge specifies that a server is feasible for the client. We need to determine if there is a matching in which the degree of each client is exactly one, and the degree of server s is at most $\delta_s^{-1}(\Delta^*)$. A feasible solution can be found via any flow algorithm.

Theorem 2 *BFlow computes a 2-approximation of an optimal assignment for min-max LDB.*

Proof : Consider an optimal assignment λ^* with cost Δ^* . It holds that $\Delta_1 = \max_u D(u, \lambda^*(u)) \leq \Delta^*$, and $\Delta_2 = \max_s \delta_s(\mathcal{L}(s)) \leq \Delta^*$. A phase of BFlow that tests an estimate $\Delta = \max(\Delta_1, \Delta_2)$ is guaranteed to find a feasible solution with cost $\Delta' \leq \Delta_1 + \Delta_2 \leq 2\Delta^*$. \square

Since there are at most kn distinct D values, the number of binary search phases is logarithmic in n . The number of phases needed for covering all possible capacity values of server s is $O(\log \delta_s(n))$, which is $O(n)$ for any reasonable δ_s .

3.3 Optimal Assignment on a Line with Euclidean Distances

In this section, we consider the case when the users and the servers are located on a line segment $[0, L]$, and the network delays are Euclidean distances. We show that min-max LDB is polynomially solvable in this model through dynamic programming.

We start with some definitions. For simplicity of presentation, we assume that every user or server i has a distinct location x_i . The distance between user u and server s is therefore $D(u, s) = |x_s - x_u|$. Assignment λ is called *order-preserving* if for every pair of users u_1 and u_2 such that $x_{u_1} < x_{u_2}$ it holds that $x_{\lambda(u_1)} \leq x_{\lambda(u_2)}$. Otherwise, both λ and every pair (u_1, u_2) for which this condition does not hold are called *order-violating*.

Every order-preserving assignment partitions the line into a series of non-overlapping segments such that every user within segment i is assigned to server s_i . Segment i is located to the left of segment j if and only if $i < j$. Note that s_i is not necessarily located inside segment i .

Theorem 3 *The min-max LDB problem on a line has an order-preserving optimal assignment.*

Proof : Consider an order-violating assignment λ . We show how it can be transformed into an order-preserving assignment that incurs smaller or equal cost.

Since λ is order-violating, there exists a pair of users u_1 and u_2 assigned to servers s_2 and s_1 such that $x_{u_1} < x_{u_2}$ but $x_{s_2} > x_{s_1}$. We transform λ to a new assignment λ' from by switching the assignments of u_1 and u_2 , i.e., $\lambda'(u_1) = s_1$ and $\lambda'(u_2) = s_2$. Since this switch does not affect the load on s_1 and s_2 , no change is incurred to any user's processing delay. Therefore, only the network delays incurred to u_1 and u_2 are affected. We therefore need to show that λ' does not incur greater maximum network delay values than λ , that is, we need to show that $\max(D(u_1, s_1), D(u_2, s_2)) \leq \max(D(u_1, s_2), D(u_2, s_1))$. To this end, consider the following cases:

- $x_{u_1} < x_{u_2} < x_{s_1} < x_{s_2}$ (Figure 1(a)). Then, $D(u_1, s_1) < D(u_1, s_2)$ and $D(u_2, s_2) < D(u_1, s_2)$, hence, $\max(D(u_1, s_1), D(u_2, s_2)) < \max(D(u_1, s_2), D(u_2, s_1))$.
- $x_{u_1} < x_{s_1} < x_{u_2} < x_{s_2}$ (Figure 1(b)). Then, $D(u_1, s_1) < D(u_1, s_2)$ and $D(u_2, s_2) < D(u_1, s_2)$, hence, $\max(D(u_1, s_1), D(u_2, s_2)) < \max(D(u_1, s_2), D(u_2, s_1))$.
- $x_{s_1} < x_{u_1} < x_{u_2} < x_{s_2}$ (Figure 1(c)). Then, $D(u_1, s_1) < D(u_2, s_1)$ and $D(u_2, s_2) < D(u_1, s_2)$, hence, $\max(D(u_1, s_1), D(u_2, s_2)) < \max(D(u_1, s_2), D(u_2, s_1))$.
- $x_{s_1} < x_{u_1} < x_{s_2} < x_{u_2}$. Symmetric to case (2).
- $x_{s_1} < x_{s_2} < x_{u_1} < x_{u_2}$. Symmetric to case (1).

Thus, we switch the assignment of every order-violating pair of users until an order-preserving assignment is obtained. We conclude that every optimal assignment for min-max LDB is either order-preserving, or can be transformed into an order-preserving assignment that incurs an equal service delay. \square

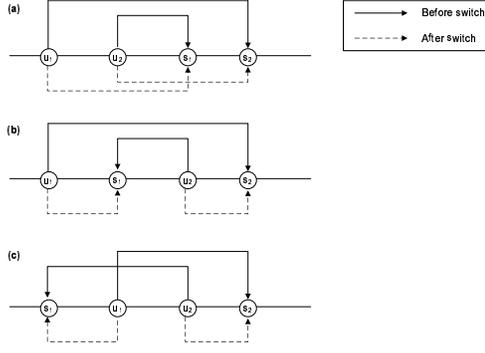


Figure 1: **Switching the assignment of an order-violating pair** (u_1, u_2) .

We now identify the recursive structure of an optimal assignment λ^* . Let $\lambda_{i,j}^*$ for $1 \leq i \leq n$ and $1 \leq j \leq k$ be an optimal assignment for users $\{u_i, \dots, u_n\}$ that employs servers $\{s_j, \dots, s_k\}$. We can assign $\ell = 0, \dots, n - i + 1$ leftmost users to server s_j . This assignment defines the maximum delay among the leftmost users. From the optimality of $\lambda_{i,j}^*$, the assignment $\lambda_{i+\ell, j+1}^*$ of the remaining users to the remaining servers is also an optimal one. Hence,

$$\Delta^M(\lambda_{i,j}^*) = \min_{0 \leq \ell \leq n-i+1} [\max(\delta_{s_j}(\ell) + \max_{0 \leq \ell' < \ell} |x_{s_j} - x_{u_{i+\ell'}}|, \Delta^M(\lambda_{i+\ell, j+1}^*))], \quad (1)$$

The boundary conditions are: $\Delta^M(\lambda^* n + 1, j) = 0$ (no users), and $\Delta^M(\lambda^* i, k + 1) = \infty$ (no servers), for $1 \leq i \leq n$ and $1 \leq j \leq k$. The global optimal assignment cost is $\Delta^M(\lambda^*(U, S)) = \Delta^M(\lambda_{1,1}^*)$.

Optimal assignments can be computed through dynamic programming using the above recurrence. An optimal algorithm employs a two-dimensional table `Table[1..n+1, 1..k+1]`, where an entry `Table[i, j]` holds the value of $\Delta^M(\lambda_{i,j}^*)$, and the number of users assigned to s_j . Note that

$$\max_{0 \leq \ell' < \ell} |x_{s_j} - x_{u_{i+\ell'}}| = \max(|x_{s_j} - x_{u_i}|, |x_{s_j} - x_{u_{i+\ell-1}}|),$$

and hence, the computation of a single entry `Table[i, j]` incurs $O(1)$ operations for each examined entry `Table[i + \ell, j + 1]`. A naive implementation examines $O(n)$ such entries, and therefore, the time complexity of filling the whole table is $O(kn^2)$. This result can be improved by noting that Eq. (1) defines a min-max among the value pairs of $f_{i,j}(\ell) = \delta_{s_j}(\ell) + \max_{0 \leq \ell' < \ell} |x_{s_j} - x_{u_{i+\ell'}}|$ (a non-decreasing function of ℓ) and $g_{i,j}(\ell) = \Delta^M(\lambda_{i+\ell, j+1}^*)$ (a non-increasing function of ℓ). Hence, the min-max is achieved for the value of ℓ for which $f_{i,j}(\ell) - g_{i,j}(\ell)$ is closest to zero. It can be efficiently found through binary search, which yields $O(\log n)$ operations for a single table entry, and $O(kn \log n)$ operations altogether.

4 Min-Average Load-Distance Balancing

We now present a polynomial-time algorithm for min-average load-distance balancing. Contrary to the min-max problem, the goal is to minimize the *sum* of delays among all users. We require that for each server s , the function $x\delta_s(x)$ is convex (most practical congestion delay functions satisfy this requirement).

The algorithm reduces the assignment problem to minimum-cost matching in a bipartite graph. The left part contains n clients, and the right part contains n copies of each server (i.e., nk nodes). The cost of connecting user u to the i 'th instance of server s is defined as

$$\Delta_i(u, s) = D(u, s) + i\delta_s(i) - (i - 1)\delta_s(i - 1).$$

Intuitively, these costs are *marginal* costs in the assignment, that is, $\Delta_i(u, s)$ is the cost of connecting user u to server s after $i - 1$ other users.

The algorithm computes a minimum-cost matching in the constructed graph (i.e., each user is assigned to exactly one server copy), and turns this matching to a legal assignment by assigning each user to the server it is matched to, regardless of the instance number.

Theorem 4 *The algorithm computes an optimal assignment for min-average LDB.*

Proof : We first claim that if the copy s_i of server s is utilized by the matching, then all the copies s_j for $j \leq i$ are used too. Indeed, suppose by contradiction that user u is matched to some copy s_i ($i > 1$), and s_{i-1} is not used. If u is switched from s_i to s_{i-1} , the matching cost can be reduced by

$$\Delta_i(u, s) - \Delta_{i-1}(u, s) = i\delta_s(i) + (i-2)\delta_s(i-2) - 2\delta_s(i-1),$$

which is a positive value since $x\delta_s(x)$ is a convex function. Hence, the matching's cost can be improved, in contradiction to optimality.

Consider a matching μ in the bipartite graph for which the set of used instances of each server is contiguous, and the corresponding assignment λ for the original problem. We denote the set of users assigned to some instance of server s by $\mu(s)$, and the user assigned to the i 'th copy of server s by $\mu_i(s)$. Since the used set is contiguous, the sum of individual matching cost of the users in $\mu(s)$ telescopes to

$$\sum_{i=1}^{|\mu(s)|} \Delta_i(\mu_i(s), s) = |\mu(s)|\delta_s(|\mu(s)|) + \sum_{i=1}^{|\mu(s)|} D(\mu_i(s), s) = \sum_{i=1}^{|\mu(s)|} [D(\mu_i(s), s) + \delta_s(|\mu(s)|)] = \sum_{u:\lambda(u)=s} \Delta(u, \lambda).$$

Hence, the cost of the matching is equal to the cost of an assignment for the original problem. Therefore, since the minimum-cost matching μ^* has the desired property of contiguity, it produces a minimum-cost assignment λ^* . \square

Optimal Assignment on a Line with Euclidean Distances: The fastest known minimum-cost flow algorithm on a graph $G(V, E)$ runs in $O(|E| \log |V| (|E| + |V| \log |V|))$ time [7]. We construct a bipartite graph in which $|V| = O(nk)$ and $|E| = O(kn^2)$, hence the running time is $O(kn^2 \log(nk)(kn^2 + nk \log(nk))) = O(k^2 n^4 \log n)$. In the special case when users and servers are located on a line segment, and network delays are modeled as Euclidean distances, this running time can be significantly improved. Similarly to the min-max LDB problem, the min-average LDB on a line has an order-preserving optimal assignment. Hence, a polynomial time dynamic programming algorithm similar to the one presented in Section 3.3 is applicable in this case. The algorithm's running time is $O(kn^2)$ (in contrast to min-max LDB, the binary search optimization to reduce the number of operations on a single table entry to $\log n$ cannot be applied).

5 Conclusions and Future Work

We studied two variations of the load-distance balancing (LDB) problem, namely, min-max LDB and min-average LDB. For the first problem, we proved hardness of approximation for general cost functions, and presented the best possible approximation algorithm, as well as an optimal algorithm for the case of linear placement of clients and servers. For the second problem, we presented a polynomial algorithm, and its performance optimization for the linear case.

A more general question might be to optimize over the choice of servers, rather than fixing the set of servers. For example, how can we choose a subset of k servers to open, and find an assignment of clients to open servers so as to minimize the average cost, or maximum cost of a client? Alternatively, each server may have a cost, and there may be a budget on the total cost of open servers. These and other questions would be interesting to study.

Acknowledgements

We thank Israel Cidon, Uri Feige, Idit Keidar and Isaac Keslassy for stimulating discussions.

References

- [1] I.F. Akylidiz, X. Wang, and W. Wang. Wireless Mesh Networks: a Survey. *Computer Networks Journal*, 2005.
- [2] J. Bar-Ilan, G. Kortsarz, and D. Peleg. How to allocate network centers. *Journal of Algorithms*, 15:385–415.
- [3] E. Bortnikov, I. Cidon, and I. Keidar. Scalable Load-Distance Balancing. In *International Symposium on Distributed Computing (DISC)*, 2007.
- [4] J. Chen, B. Knutsson, B. Wu, H. Lu, M. Delap, and C. Amza. Locality Aware Dynamic Load Management form Massively Multiplayer Games. *PPoPP*, 2005.
- [5] S. Khuller and Y. J. Sussmann. The Capacitated K-Center Problem. *SIAM Journal on Discrete Mathematics*, 13:403–418, 2000.
- [6] P. B. Mirchandani and R. L. Francis. *Discrete Location Theory*. John Wiley & Sons Inc., 1990.
- [7] J. Orlin. A Faster Strongly Polynomial Minimum Cost Flow Algorithm. *ACM STOC*, 1988.