Analysis of Server Provisioning for Distributed Interactive Applications

Zheng, Hanying; Tang, Xueyan

2015

Zheng, H., & Tang, X. (2015). Analysis of Server Provisioning for Distributed Interactive Applications. IEEE Transactions on Computers, 64(10), 2752-2766.

<https://hdl.handle.net/10356/80933>

<https://doi.org/10.1109/TC.2014.2378252>

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [<http://dx.doi.org/10.1109/TC.2014.2378252>].

Downloaded on 26 Apr 2023 21:25:56 SGT

Analysis of Server Provisioning for Distributed Interactive Applications

Hanying Zheng and Xueyan Tang

Abstract—Increasing geographical spreads of modern distributed interactive applications (DIAs) make distributed server deployment vital for combating network latency and improving the interactivity among participants. In this paper, we investigate the server provisioning problem that concerns where to place servers for DIAs. We formulate the server provisioning problem with an objective of reducing the network latency involved in the interaction between participants. We prove that the problem is NP-hard under several scenarios. We analyze the performance of the classical k -median server placement for DIAs and propose a new greedy server provisioning heuristic for DIAs. Theoretical analysis shows that the approximation ratio of the proposed greedy algorithm is much lower than that of the k -median placement. Experiments using real Internet latency data also show that our proposed algorithm significantly outperforms the k -median and other baseline server placements.

Index Terms—Distributed interactive application, server placement, interactivity

I. INTRODUCTION

An increasing number of Distributed Interactive Applications (DIAs) are emerging in recent years to provide people with new ways of collaboration and entertainment. In these applications, participants dispersed at different locations interact with each other through the network in real time. Examples of DIAs include online gaming [2], instant messaging [3], collaborative computer-aided design and engineering [4], and web-based e-learning [5].

A critical Quality of Service measure for DIAs is the time lag experienced by the participants during their interaction. The interaction process in DIAs normally involves communicating user-initiated operations and their resultant updates to the application state (such as the virtual game worlds in online gaming and the shared workspaces in collaborative design tools) between the participants and the servers as well as executing the operations at the servers. Thus, the time lag in the interaction includes the network latency in the communication and the processing delay at the servers. The latter is generally easier to cut than the former. In particular, the emerging cloud computing paradigm enables customers to rent computing resources purely on demand for hosting their applications [6]. Although elastic computing power supply from the clouds can largely minimize the server-side processing delay in DIAs, the network latency remains as a major barrier to achieving high quality interaction experience. Excessive network latency can severely degrade the participant's quality of experience in DIAs [2].

Increasing geographical spreads of participants in modern DIAs necessitate distributed server deployment to combat network latency [7]. In this paper, we focus on reducing the network latency involved in the interaction between participants by considering where to place servers in the network. We formulate the server provisioning problem for DIAs as a combinatorial optimization problem and prove that it is NP-hard under *any* one of the following three scenarios that may be common in practice: (a) the network latency does not satisfy the triangle inequality; or (b) the choices of server locations in the network are restricted; or (c) the number of server locations to select is limited. We then propose an efficient greedy heuristic for server provisioning in DIAs. We conduct theoretical analysis to compare the greedy heuristic with the classical k -median server placement. In particular, we study the approximability of the algorithms for networks with triangle inequality violations. To the best of our knowledge, this is the first paper that analyzes the asymptotic approximability of server placement algorithms with respect to the extent of triangle inequality violations. The results suggest that our proposed algorithm has significantly better approximability than the k -median placement. The proposed algorithm is also experimentally evaluated with real Internet latency data. The results again show that our algorithm significantly outperforms the k -median and other baseline server placements.

Related Work. The classical k -median and k -center problems have been widely used to model server placement in the Internet [8]–[12]. The k -median placement aims to place a given number of k servers to minimize the total distance (latency) from the clients to their nearest servers, whereas the k -center placement aims to place k servers to minimize the maximum distance (latency) from the clients to their nearest servers. Both problems are NP-hard [13]. A variety of heuristic approaches have been explored for these problems [10], [14]–[16]. The k -median and k -center server placements are quite successful for web content delivery [8], [9]. However, DIAs are fundamentally different from web content delivery. The clients in the web just download contents from web servers. Thus, their access performance can be optimized by simply minimizing the client-to-server latency. In contrast, the clients in DIAs are engaged in mutual interactions among themselves. Each client connects to one server through which it interacts with all the other clients. Therefore, the interaction time between clients must include not only the network latencies from the clients to their connected servers but also the latencies between their connected servers. We shall analyze and evaluate the performance of the k -median server placement for DIAs. In a recent work, we have investigated client assignment

strategies for enhancing the interactivity performance of DIAs given a set of servers placed [17]. This paper complements our earlier work by studying server provisioning for DIAs.

This paper significantly extends a preliminary conference version [1]. The rest of this paper is organized as follows. Section II formulates the server provisioning problem for DIAs and Section III analyzes its hardness. Section IV studies the approximability of the k -median server placement. Section V presents the greedy server provisioning heuristic and the theoretical analysis of its approximation ratio. Experimental evaluations are elaborated in Section VI. Finally, Section VII concludes the paper.

II. PROBLEM FORMULATION

We model the network infrastructure underlying the DIA as a graph $G = (V, E)$, where V is the set of nodes and E is the set of links. For each pair of nodes $(u, v) \in V \times V$, we denote by $d(u, v)$ the latency of the network path between nodes u and v . We define $d(v, v) = 0$ for each node $v \in V$.

Without loss of generality, we assume that servers can be placed only at a particular set of nodes $Z \subseteq V$ in the network and refer to these nodes as *candidate server locations*. Z can be defined in different ways to address different server provisioning scenarios. For example, there are two typical types of models for running DIAs. In the traditional client-server model, the application state is maintained by dedicated servers. Participants, known as clients, are responsible for sending user-initiated operations to the servers for execution and receiving the state updates from the servers. In this case, the candidate server locations (e.g., the data centers operated by cloud providers) are normally separate from client locations. In the peer-to-peer model, on the other hand, the clients are responsible for executing operations and maintaining the application state by themselves. In this case, to select which clients to take up the role of operation execution, the candidate server locations Z can be modeled as the set of clients.

Let $C \subseteq V$ be the set of clients in the network. To participate in the DIA, each client $c_i \in C$ needs to connect to one server. The clients are often autonomous in deciding which servers to connect to. An intuitive and widely used strategy in many applications is for the clients to connect to their nearest servers, i.e., the servers with the shortest network latency to them [11], [17]–[19]. Suppose that a set of locations $S \subseteq Z$ are selected to place servers. For each client $c_i \in C$, we denote by $n(c_i, S)$ the nearest server to c_i , i.e., $d(c_i, n(c_i, S)) = \min_{s \in S} d(c_i, s)$. Then, all the operations issued by client c_i would be sent to $n(c_i, S)$.

The interaction between two clients c_i and c_j goes through their connected servers. On receiving an operation issued by c_i , its server $n(c_i, S)$ forwards the operation to c_j 's connected server $n(c_j, S)$ if they are different. Then, $n(c_j, S)$ executes the operation¹ and delivers the resultant state update to c_j to present the effect of c_i 's operation. Thus, the interaction process involves the paths from c_i to $n(c_i, S)$, from

$n(c_i, S)$ to $n(c_j, S)$, and from $n(c_j, S)$ to c_j . We refer to the concatenation of these paths as the *interaction path* from c_i to c_j . The length of the interaction path is given by $d(c_i, n(c_i, S)) + d(n(c_i, S), n(c_j, S)) + d(n(c_j, S), c_j)$ and represents the network latency involved in the interaction between c_i and c_j . Note that the interaction path from a client c_i to itself is the round trip between c_i and its connected server $n(c_i, S)$, whose length indicates the network latency involved for c_i to see the effect of its own operation.

We measure the interactivity performance of the DIA by the average interaction path length between all client pairs:

$$\frac{1}{|C|^2} \sum_{c_i \in C} \sum_{c_j \in C} \left(d(c_i, n(c_i, S)) + d(n(c_i, S), n(c_j, S)) + d(n(c_j, S), c_j) \right).$$

The shorter the average interaction path length, the higher the interactivity performance of the DIA. Since the total number of client pairs is fixed given the client set, to minimize the average interaction path length, it is equivalent to minimize the total interaction path length between all client pairs. Therefore, the server provisioning problem for DIAs is defined as follows.

DIA Server Provisioning Problem. *Given a set of clients C and a set of candidate server locations Z in the network, select a set of locations $S \subseteq Z$ to place servers for the DIA that minimizes the total interaction path length between all client pairs, i.e.,*

$$\text{minimize } \sum_{c_i \in C} \sum_{c_j \in C} \left(d(c_i, n(c_i, S)) + d(n(c_i, S), n(c_j, S)) + d(n(c_j, S), c_j) \right),$$

where $n(c_i, S)$ refers to the nearest server in S to client c_i .

In the above formulation, we focus on minimizing the network latencies involved in the interaction. No server capacity limitation is assumed at the candidate server locations. We shall discuss and evaluate how to deal with server capacity constraints in our proposed algorithms in Section VI-B.

III. HARDNESS ANALYSIS

The DIA server provisioning problem is trivial if (1) the network latencies among the nodes satisfy the triangle inequality; (2) all the nodes in the network are candidate server locations; and (3) there is no limit on the number of server locations to select (or more precisely, the number of server locations selected can be as large as the number of clients because there is certainly no need to place more servers than the number of clients). The triangle inequality implies that for any two clients c_i and c_j , and any two servers s_a and s_b , we have $d(c_i, s_a) + d(s_a, s_b) + d(s_b, c_j) \geq d(c_i, c_j)$. Therefore, under the above three assumptions, the optimal server provisioning solution is to place servers at all the nodes where the clients are located. In this way, each client connects to the server co-located with it so that the latencies between all clients and their nearest servers are 0. Thus, the interaction path length between any two clients c_i and c_j is simply $d(c_i, c_j)$, which is the shortest possible.

¹In some applications such as instant messaging, operation execution is as simple as forwarding the message typed by the user. In other applications like online gaming, operation execution could involve more complex computation and modification to the application state.

Fig. 1. Example instances of the minimum set cover problem and the DIA server provisioning problem

Interestingly, if *any* one of the above assumptions is relaxed, the DIA server provisioning problem becomes NP-hard.

A. Networks without Triangle Inequality

Triangle inequality violations are not uncommon as Internet routing is often based on business policies and is not optimal in terms of network latency [20]. If the network latency does not satisfy the triangle inequality, we can show the NP-hardness of the DIA server provisioning problem by a polynomial reduction from the minimum set cover problem [13]. Given a finite set P and a collection $\mathcal{Q}$ of its subsets, and a positive integer $k \leq |\mathcal{Q}|$, the decision version of the minimum set cover problem is to find out whether $\mathcal{Q}$ contains a subcollection $\mathcal{Q}'$ of at most k subsets such that $\bigcup_{Q \in \mathcal{Q}'} Q = P$.

Consider an instance $\mathcal{R}$ of the minimum set cover problem. Suppose that set P contains n elements: $P = \{p_1, p_2, \dots, p_n\}$, and collection $\mathcal{Q}$ contains m subsets: $\mathcal{Q} = \{Q_1, Q_2, \dots, Q_m\}$. We construct a network $G_{\mathcal{R}}$ consisting of the node set $V_C \cup V_S$, where V_C contains n nodes $V_C = \{c_1, c_2, \dots, c_n\}$, and V_S contains k groups of nodes $V_S = \bigcup_{i=1}^k V^i$. Each node $c_i \in V_C$ corresponds to an element p_i in set P . Each group V^i contains m clusters of nodes $V^i = \bigcup_{j=1}^m V_j^i$. Each cluster V_j^i corresponds to a subset Q_j , and contains $|Q_j|$ nodes that correspond to the elements of Q_j . Thus, there are a total of $k \cdot \sum_{j=1}^m |Q_j|$ nodes in V_S .

Among the nodes in V_S , the latency between any two nodes in the same cluster or in different groups is set to 1. The latency between any two nodes in different clusters of the same group is set to L , where $L > 3n^2$. The latency between any two nodes in V_C is also set to L . The latency from a node $c_i \in V_C$ to a node $v \in V_S$ is set to 1 if v corresponds to element $p_i \in P$, and is set to L otherwise. Figure 1 illustrates an example of network $G_{\mathcal{R}}$, where each node pair connected by a link has latency 1, and each pair not connected by a link has latency L .

An instance $\mathcal{T}$ of the DIA server provisioning problem in the decision version is then defined on the constructed network $G_{\mathcal{R}}$ as follows: Suppose that a client is located at each node in V_C , all the nodes in the network are candidate server locations, and there is no limit on the number of server locations to select. Can we select a set of server locations such that the total interaction path length between all client pairs is bounded by $B = 3n^2$?

We first prove that if $\mathcal{Q}$ contains a set cover of size at most k for instance $\mathcal{R}$, there must exist a server placement with total interaction path length bounded by B for instance $\mathcal{T}$. Let $\mathcal{Q}' = \{Q_{x_1}, Q_{x_2}, \dots, Q_{x_l}\}$ (where $1 \leq l \leq k$) be a set cover of size not exceeding k for instance $\mathcal{R}$. Then, placing servers at all the nodes in clusters $V_{x_1}^1, V_{x_2}^2, \dots, V_{x_l}^l$ is a valid server provisioning solution for instance $\mathcal{T}$. In fact, since $\mathcal{Q}'$ is a set cover, each element $p_i \in P$ is contained in at least one subset among $Q_{x_1}, Q_{x_2}, \dots, Q_{x_l}$. Thus, for each client $c_i \in V_C$, there exists at least one node in $V_{x_1}^1 \cup V_{x_2}^2 \cup \dots \cup V_{x_l}^l$ having latency 1 to c_i . As a result, the latencies between all the clients and their nearest servers are 1. In addition, the latency between any two servers in $V_{x_1}^1 \cup V_{x_2}^2 \cup \dots \cup V_{x_l}^l$ is 1 because they either come from the same cluster or from different groups. Therefore, the interaction path length between each pair of clients is bounded by $1 + 1 + 1 = 3$. Since there are n^2 interaction paths in total, the total interaction path length is bounded by $3n^2 = B$.

Next, we prove that if a valid server placement solution can be found for instance $\mathcal{T}$, there must exist a set cover of size at most k for instance $\mathcal{R}$. Let $S \subseteq V_C \cup V_S$ be the set of server locations selected in a valid solution of instance $\mathcal{T}$. Without loss of generality, we assume that the server at each location in S is connected by at least one client.² Since each client connects to its nearest server, if a server is placed at a node in V_C , the client at that node should connect to this server. Note that the latency between any two nodes in V_C is L . If servers are placed at more than one node in V_C , the interaction path length between the clients at these nodes becomes $L > 3n^2 = B$. Therefore, at most one node in V_C can be selected to place a server, i.e., $|S \cap V_C| \leq 1$.

If $|S \cap V_C| = 1$, there is exactly one server located in V_C and the remaining servers are all located in V_S . Let the server in V_C be located at node c_x . Since the latency between any two nodes in V_C is L , to bound the total interaction path length by $B < L$, the clients in $V_C \setminus \{c_x\}$ should all connect to the servers in V_S . Moreover, each of these clients must have latency 1 to its nearest server because the latency from a client to a server in V_S is either 1 or L . Thus, the clients in $V_C \setminus \{c_x\}$ must connect to servers not corresponding to element $p_x \in P$. Therefore, the latencies from c_x to these servers are L , exceeding the bound B . So, there is no valid solution of instance $\mathcal{T}$ satisfying $|S \cap V_C| = 1$.

If $|S \cap V_C| = 0$, all the servers are placed in V_S . Similar to the former case, since the total interaction path length does not exceed B , each client c_i must have latency 1 to its nearest server. This implies that each element $p_i \in P$ is covered by

²The server locations not connected by any clients can simply be removed from S .

the subset Q_j corresponding to the node cluster of c_i 's nearest server. Thus, $\mathbb{Q}' = \{Q_j \mid \exists i, V_j^i \cap S \neq \emptyset\}$ is a set cover for P . Furthermore, to cap the total interaction path length by B , all the servers must have latency 1 from each other. Thus, for each group of nodes V^i , all the servers in $V^i \cap S$ must come from the same cluster of V^i . Since there are k groups of nodes $V^1, V^2, \dots, V^k$ in V_S , we have $|\mathbb{Q}'| \leq k$. Therefore, $\mathbb{Q}'$ is a set cover of size at most k .

Hence, a set cover of size at most k can be found for instance $\mathcal{R}$ if and only if there exists a server placement with total interaction path length bounded by B for instance $\mathcal{T}$. Thus, the DIA server provisioning problem for networks without the triangle inequality is NP-hard.

The above hardness analysis also leads to the following non-approximability result.

Theorem 1. *For networks without the triangle inequality, the DIA server provisioning problem cannot be approximated by any constant factor unless $P = NP$.*

Proof: For any instance of the minimum set cover problem, if there exists a set cover of size at most k , the total interaction path length of an optimal server placement must be bounded by $B = 3n^2$ in the network constructed above. Any non-optimal server placement, however, produces a total interaction path length greater than L . This is because in any non-optimal placement, either a client connects to a server that has a latency L to it, or the latency between a certain pair of servers is L . In either case, there is at least one interaction path longer than L , so the total interaction path length exceeds L . Similarly, if there does not exist any set cover of size at most k in the set cover problem, the total interaction path length of an optimal server placement in the constructed network is also greater than L .

Assume on the contrary that there exists a polynomial-time server provisioning algorithm with a constant approximation ratio of β . For any instance of the set cover problem, we set $L = \beta B = 3n^2\beta$ in the constructed network and then run the β -approximation algorithm. If the output server placement produces a total interaction path length larger than L , it indicates that the optimal server placement has a total interaction path length larger than B . So, there does not exist any set cover of size at most k . On the other hand, if the output server placement has a total interaction path length within L , it must be bounded by B according to the above analysis. Thus, the optimal server placement would also have a total interaction path length within B so that there exists a set cover of size at most k . Therefore, this implies a polynomial-time algorithm for the set cover problem, which contradicts to $P \neq NP$.

Hence, the theorem is proven. $\blacksquare$

B. Restricted Choices of Server Locations

If not all the nodes in the network are candidate server locations, the DIA server provisioning problem becomes NP-hard as well. We can again prove it by a polynomial reduction from the minimum set cover problem.

Given an instance $\mathcal{R}$ of the minimum set cover problem, we construct the same network $G_{\mathcal{R}}$ as in the previous section,

except that the node pairs that had latency L earlier have latency 2 now to satisfy the triangle inequality. That is, in the illustration of Figure 1, a node pair has latency 1 if they are connected by a link and has latency 2 otherwise. Suppose that a client is located at each node in V_C , only the nodes in V_S are candidate server locations, and there is no limit on the number of server locations to select. An instance $\mathcal{Y}$ of the DIA server provisioning problem is then defined on network $G_{\mathcal{R}}$ by setting a bound $H = 3n^2 - n$ on the total interaction path length between all client pairs.

We first show that a set cover of size at most k for instance $\mathcal{R}$ gives rise to a valid server provisioning solution for instance $\mathcal{Y}$. Again, let $\mathbb{Q}' = \{Q_{x_1}, Q_{x_2}, \dots, Q_{x_l}\}$ (where $1 \leq l \leq k$) be a set cover of size not exceeding k . Then, a valid solution for instance $\mathcal{Y}$ is to place servers at all the nodes in clusters $V_{x_1}^1, V_{x_2}^2, \dots, V_{x_l}^l$. Similar to the argument in Section III-A, under such server placement, the latencies between all clients and their nearest servers are 1, and the latency between any two servers is also 1. Thus, the interaction path from a client to itself has length $1 + 1 = 2$, and the interaction path length between two distinct clients is bounded by $1 + 1 + 1 = 3$. Note that there are $n(n-1)$ pairs of distinct clients. Therefore, the total interaction path length is bounded by $3n(n-1) + 2n = 3n^2 - n = H$.

Next, we show that a valid server provisioning solution for instance $\mathcal{Y}$ gives rise to a set cover of size at most k for instance $\mathcal{R}$. Let $S \subseteq V_S$ be a set of server locations that produces a total interaction path length not exceeding H . Without loss of generality, assume that the server at each location in S is connected by at least one client. Note that each node in V_S has latency 1 to exactly one client in V_C and has latency 2 to all the other clients. As a result, if two distinct clients in V_C connect to the same server in S , their interaction path length is at least $1 + 2 = 3$. On the other hand, the latency between any two nodes in V_S is at least 1. Thus, if two distinct clients connect to different servers in S , their interaction path length is at least $1 + 1 + 1 = 3$. Therefore, regardless of server placement, the interaction path length between any two distinct clients is at least 3. Since there are $n(n-1)$ pairs of distinct clients, this implies that the total interaction path length from all clients to themselves under server placement S cannot exceed $H - 3n(n-1) = 2n$.

Since the shortest latency from each client to the nodes in V_S is 1, it follows that the latencies from all the clients to their nearest servers must be exactly 1. Therefore, each element $p_i \in P$ is covered by the subset Q_j corresponding to the node cluster of c_i 's nearest server. That is, $\mathbb{Q}' = \{Q_j \mid \exists i, V_j^i \cap S \neq \emptyset\}$ is a set cover for P .

To limit the total interaction path length by H , the interaction path length between any two distinct clients must now be exactly 3. This implies that the latency between any two servers in S must be 1. Thus, for each group of nodes V^i , all the servers in $V^i \cap S$ must come from the same cluster of V^i . It follows that $|\mathbb{Q}'| \leq k$. Therefore, $\mathbb{Q}'$ is a set cover of size at most k .

In summary, there exists a set cover of size at most k for instance $\mathcal{R}$ if and only if a server placement with total interaction path length bounded by H can be found for

Fig. 2. Example instances of the 3DM problem and the DIA server provisioning problem

instance $\mathcal{Y}$. Thus, the DIA server provisioning problem with restricted choices of server locations is NP-hard.

C. Limited Number of Server Locations to Select

If the DIA operator can deploy only a limited number of servers due to budget restriction, a cap can be set on the number of server locations to select. If the cap is less than the number of clients, the DIA server provisioning problem is also NP-hard. This can be proved by a polynomial reduction from the 3-Dimensional Matching (3DM) problem [13]. The decision version of the 3DM problem is defined as follows: Given three disjoint sets W , X and Y each having k elements, and a set of triples $M \subseteq W \times X \times Y$, find out whether M contains a 3-dimensional matching, i.e., whether there exists a subset $M' \subseteq M$ such that $|M'| = k$ and any two triples (w_a, x_a, y_a) and $(w_b, x_b, y_b) \in M'$ satisfy $w_a \neq w_b$, $x_a \neq x_b$ and $y_a \neq y_b$.

Given an instance $\mathcal{U}$ of the 3DM problem, we construct a network G_U consisting of the node set $V_C \cup V_M$. The set V_C contains $3k$ nodes, each corresponding to an element in $W \cup X \cup Y$ of instance $\mathcal{U}$. The set V_M contains the same number of nodes as the size of M in instance $\mathcal{U}$. Each node in V_M corresponds to a triple in M , and establishes links of latency 2 to the three nodes in V_C corresponding to its three coordinates. For example, in Figure 2, node v_4^m , which corresponds to the triple $m_4 = (w_1, x_2, y_3)$, has three links to nodes v_1^w , v_2^x and v_3^y in V_C . In addition, all the nodes in V_M are inter-connected with each other via links of latency 1. Suppose that a client is located at each node in V_C , and all the nodes in network G_U are candidate server locations. An instance $\mathcal{F}$ of the DIA server provisioning problem is then defined on network G_U by capping the number of server locations to select at k and setting a bound $A = 45k^2 - 9k$ on the total interaction path length between all client pairs.

We first prove that if M contains a matching for instance $\mathcal{U}$, there must exist a server placement with total interaction path length bounded by A for instance $\mathcal{F}$. Suppose that $M' \subseteq M$ is a matching. Let S be the set of k nodes in V_M that represent the triples in M' (for example, nodes v_1^m , v_2^m and v_3^m in Figure 2). We construct a server placement by selecting all the k nodes in S as the server locations. Then, each client in V_C has exactly one adjacent node in S . Thus, each client connects to the server placed at its adjacent node in S and the

latency between them is 2. Since all the servers in S are inter-connected by links of latency 1, and each server has exactly three adjacent clients, the total interaction path length under placement S is given by

$$(2 + 0 + 2) \cdot 9k + (2 + 1 + 2) \cdot (9k^2 - 9k) = 45k^2 - 9k = A.$$

Next, we prove that M contains a matching for instance $\mathcal{U}$ if a valid server placement solution can be found for instance $\mathcal{F}$. Consider a set S of up to k server locations in network G_U . Let b_i ($i \geq 0$) be the number of clients that have latency i to their nearest servers in S . Since there are a total of $3k$ clients and all the links incident on clients have latency 2, we have $\sum_{i \geq 0} b_i = 3k$ and $b_1 = 0$. Note that there are b_0 servers of S located in V_C . So, the number of servers located in V_M is at most $k - b_0$. If a client has latency 2 to its nearest server, this server must be adjacent to the client and thus is located in V_M . Since each server in V_M has three adjacent clients, we have $b_2 \leq 3 \cdot (k - b_0)$. Thus, the total length of the interaction paths from the clients to themselves is

$$\begin{aligned} \sum_{i \geq 0} 2i \cdot b_i &= 0 \cdot b_0 + 4 \cdot b_2 + \sum_{i \geq 3} 2i \cdot b_i \\ &\geq 4 \cdot b_2 + 6 \cdot \sum_{i \geq 3} b_i \end{aligned} \quad (1)$$

$$\begin{aligned} &= 4 \cdot b_2 + 6 \cdot (3k - b_2 - b_0) \\ &= 18k - 2 \cdot b_2 - 6 \cdot b_0 \\ &\geq 18k - 2 \cdot (3 \cdot (k - b_0)) - 6 \cdot b_0 \quad (2) \\ &= 12k. \end{aligned}$$

Since all the links incident on clients have latency 2 and the clients are not adjacent to each other, the shortest possible interaction path between two distinct clients has length 4. Let p_i ($i \geq 4$) be the number of pairs of distinct clients whose interaction path length is i under server placement S . Since there are $3k(3k - 1)$ pairs of distinct clients in total, we have $\sum_{i \geq 4} p_i = 3k(3k - 1)$. If any two distinct clients have an interaction path length of 4, there must be a server in V_M that is adjacent to both clients. Each server in V_M has three adjacent clients and can thus support at most $3 \cdot 2 = 6$ interaction paths of length 4 between distinct clients. Since there are at most k servers, we have $p_4 \leq 6k$. Therefore, the total length of the interaction paths between distinct clients is

$$\begin{aligned} \sum_{i \geq 4} i \cdot p_i &= 4 \cdot p_4 + \sum_{i \geq 5} i \cdot p_i \\ &\geq 4 \cdot p_4 + 5 \cdot \sum_{i \geq 5} p_i \end{aligned} \quad (3)$$

$$\begin{aligned} &= 4 \cdot p_4 + 5 \cdot (3k(3k - 1) - p_4) \\ &= 45k^2 - 15k - p_4 \\ &\geq 45k^2 - 15k - 6k \quad (4) \\ &= 45k^2 - 21k. \end{aligned}$$

Thus, the total interaction path length is

$$\begin{aligned} \sum_{i \geq 0} 2i \cdot b_i + \sum_{i \geq 4} i \cdot p_i &\geq 12k + 45k^2 - 21k \\ &= 45k^2 - 9k = A. \end{aligned}$$

Therefore, if server placement S has a total interaction path length bounded by A , the equality must be satisfied at all the steps (1), (2), (3) and (4) in the above derivation.

Equality at step (3) implies that $\sum_{i>5} p_i = 0$. So, at most one client can have latency 3 to its nearest server in S , i.e., $b_3 \leq 1$. Equality at step (1) indicates $\sum_{i>3} b_i = 0$. Thus, $b_0 + b_2 + b_3 = 3k$. It follows from equality at step (2) (i.e., $b_2 = 3(k - b_0)$) that $b_3 = 2 \cdot b_0$, so b_3 is an even number. Therefore, b_3 must be 0, and b_0 is also 0. As a result, we have $b_2 = 3k$, i.e., the latency from each client to its nearest server in S must be 2. Thus, all the server locations in S are selected from node set V_M . Since each node in V_M has only three adjacent clients, exactly k server locations must be selected from V_M and these server locations cannot have any adjacent client in common. Therefore, the k triples of M corresponding to the k server locations selected in V_M form a matching.

Hence, M contains a matching for instance $\mathcal{U}$ if and only if there exists a server placement with total interaction path length bounded by A for instance $\mathcal{F}$. Thus, the DIA server provisioning problem with limited number of server locations to select is NP-hard.

IV. ANALYSIS OF k -MEDIAN PLACEMENT

Since the DIA server provisioning problem is NP-hard under several scenarios, we shall focus on investigating heuristic algorithms and their approximability. We first study the performance of the classical k -median server placement for DIAs. The k -median problem has been widely used to model server placement in the Internet [8], [9], [12]. The k -median problem seeks to select a given number of k locations in the network to host servers such that the total distance (latency) from the clients to their nearest servers is minimized. Our DIA server provisioning problem shares some similarity with the k -median problem in that we also aim to find server placement that minimizes an aggregate latency metric. However, in our problem, the interaction time between the clients in DIAs includes not only the latencies from the clients to their connected servers but also the latencies between their connected servers.

While many works in the field of network coordinate systems [21], [22] have assumed that the triangle inequality holds in the Internet, several other works have argued that triangle inequality violations are not uncommon due to the restrictions of the Internet structure and the routing policies of ISPs [20], [23]. Without loss of generality, our analysis concerns the performance of server placement not only in networks with the triangle inequality but also in networks with triangle inequality violations. To this end, we relax the triangle inequality constraint by introducing an additional parameter α and define the α -triangle inequality as follows.

α -Triangle Inequality. *Given a constant α ($\alpha \geq 1$), a network is said to satisfy the α -triangle inequality if for any three different nodes u , v and w in the network, it holds that*

$$d(u, v) \leq \alpha \cdot (d(u, w) + d(w, v)).$$

Note that for any two different nodes u and v in the network, it is obvious that $d(u, v) \leq \alpha \cdot d(u, v) = \alpha \cdot (d(u, v) + d(v, v))$.

Therefore, if a network satisfies the α -triangle inequality, any nodes u , v and w in the network fulfills $d(u, v) \leq \alpha \cdot (d(u, w) + d(w, v))$.

The α -triangle inequality is generic because for any network, a value of α can always be found for the network latencies to satisfy the α -triangle inequality. Suppose that V is the set of nodes in a network. Then, α can be set at

$$\max \left\{ 1, \max_{\substack{u, v, w \in V \\ u \neq v \neq w}} \frac{d(u, v)}{d(u, w) + d(w, v)} \right\}.$$

In essence, the parameter α characterizes the extent of triangle inequality violations. The larger the value of α , the more severe the triangle inequality violations. When $\alpha = 1$, the α -triangle inequality degenerates to the original triangle inequality. Given a constant $\alpha > 1$, the networks satisfying the α -triangle inequality but not the original 1-triangle inequality constitute a *subset* of all the networks with triangle inequality violations. Although Theorem 1 has given a non-approximability result for networks with triangle inequality violations, it is possible to approximate the DIA server provisioning problem for networks satisfying the α -triangle inequality.

Next, we analyze the performance of the k -median server placement under the α -triangle inequality. When a limit k is set on the number of server locations to select in the DIA server provisioning problem, we have the following approximability results of the k -median server placement.

Theorem 2. *For networks satisfying the α -triangle inequality ($\alpha \geq 1$), the k -median server placement has an approximation ratio of $\max\{\frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}\}$ for the DIA server provisioning problem with a limit k on the number of server locations to select.*

Proof: Consider a network satisfying the α -triangle inequality ($\alpha \geq 1$). Assume that there are n clients $c_1, c_2, \dots, c_n$ in the network. Let S be a set of candidate server locations in the network. Suppose that the set of optimal server locations for the DIA server provisioning problem is $S_P \subseteq S$. We denote by $p_i \in S_P$ the nearest server of each client c_i ($1 \leq i \leq n$) under placement S_P . Suppose that the k -median server placement is $S_M \subseteq S$. We denote by $m_i \in S_M$ the nearest server of each client c_i ($1 \leq i \leq n$) under placement S_M .

Define a new variable

$$x = \frac{\sum_{i=1}^n \sum_{j=1}^n d(p_i, p_j)}{n \cdot \sum_{i=1}^n d(c_i, p_i)}.$$

Then, it can be shown that the total interaction path length T_M achieved by the k -median server placement S_M satisfies

$$T_M \leq n \sum_{i=1}^n d(c_i, p_i) \cdot \left(2 + 2\alpha^2 + \alpha^3 + \alpha^2 \cdot \min \left\{ x + \alpha, \alpha x + 1 \right\} \right). \quad (5)$$

Please refer to Appendix A in the supplementary material for the detailed derivation.

If $x \geq 1$, since $\alpha \geq 1$, we have $x - 1 \leq \alpha(x - 1)$. It follows that $\min\{x + \alpha, \alpha x + 1\} = x + \alpha$. Thus,

$$T_M \leq n \sum_{i=1}^n d(c_i, p_i) \cdot (2 + 2\alpha^2 + \alpha^3 + \alpha^2(x + \alpha)).$$

Note that the total interaction path length under placement S_P (the optimal solution of the server provisioning problem) is

$$\begin{aligned} T_P &= 2n \sum_{i=1}^n d(c_i, p_i) + \sum_{i=1}^n \sum_{j=1}^n d(p_i, p_j) \\ &= n \sum_{i=1}^n d(c_i, p_i) \cdot (2 + x). \end{aligned}$$

Therefore, the largest possible ratio between T_M and T_P is

$$\begin{aligned} r(x \geq 1) &= \max_{x \geq 1} \frac{2 + 2\alpha^2 + \alpha^3 + \alpha^2(x + \alpha)}{2 + x} \\ &= \max_{x \geq 1} \left(\alpha^2 + \frac{2\alpha^3 + 2}{2 + x} \right) \\ &= \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}. \end{aligned} \quad (6)$$

Similarly, if $0 \leq x \leq 1$, since $\alpha \geq 1$, we have $x - 1 \geq \alpha(x - 1)$ and thus $\min\{x + \alpha, \alpha x + 1\} = \alpha x + 1$. In this case, the largest possible ratio between T_M and T_P is

$$\begin{aligned} r(0 \leq x \leq 1) &= \max_{0 \leq x \leq 1} \frac{2 + 2\alpha^2 + \alpha^3 + \alpha^2(\alpha x + 1)}{2 + x} \\ &= \max_{0 \leq x \leq 1} \left(\alpha^3 + \frac{2 + 3\alpha^2 - \alpha^3}{2 + x} \right). \end{aligned}$$

If $2 + 3\alpha^2 - \alpha^3 \geq 0$, i.e., $1 \leq \alpha \leq 1 + \sqrt[3]{2 + \sqrt{3}} + \sqrt[3]{2 - \sqrt{3}}$, we have

$$\max_{0 \leq x \leq 1} \left(\alpha^3 + \frac{2 + 3\alpha^2 - \alpha^3}{2 + x} \right) = \frac{2 + 3\alpha^2 + \alpha^3}{2}.$$

If $2 + 3\alpha^2 - \alpha^3 < 0$, i.e., $\alpha > 1 + \sqrt[3]{2 + \sqrt{3}} + \sqrt[3]{2 - \sqrt{3}}$, we have

$$\max_{0 \leq x \leq 1} \left(\alpha^3 + \frac{2 + 3\alpha^2 - \alpha^3}{2 + x} \right) = \frac{2 + 3\alpha^2 + 2\alpha^3}{3}.$$

Thus,

$$r(0 \leq x \leq 1) = \begin{cases} \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1 & \text{if } 2 + 3\alpha^2 \geq \alpha^3, \\ \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3} & \text{if } 2 + 3\alpha^2 < \alpha^3, \end{cases} \quad (7)$$

which can actually be rewritten as

$$r(0 \leq x \leq 1) = \max \left\{ \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3} \right\}.$$

Combining (6) and (7), if $\alpha \geq 1$, the approximation ratio of the k -median placement for the server provisioning problem is given by

$$\max \left\{ \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3} \right\}.$$

Hence, the theorem is proven. $\blacksquare$

When $\alpha = 1$, $\max \left\{ \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3} \right\} = \max\{3, \frac{7}{3}\} = 3$. Thus, Theorem 2 implies the following corollary.

Corollary 1. For networks satisfying the 1-triangle inequality, the k -median server placement has an approximation ratio of 3 for the DIA server provisioning problem with a limit k on the number of server locations to select.

The proof of Theorem 2 also gives rise to the following result for the DIA server provisioning problem when there is no limit on the number of server locations to select.

Corollary 2. For each instance of the DIA server provisioning problem on a network satisfying the α -triangle inequality, if an optimal server placement selects k server locations, then the k -median server placement produces a total interaction path length within $\max\{\frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}\}$ times of that in the optimal server placement.

The approximation ratio derived in Theorem 2 is tight. To demonstrate, we present an example in which the ratio between the total interaction path lengths of the k -median placement and optimal placement can be made arbitrarily close to $\frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1$, and another example in which the ratio can be made arbitrarily close to $\frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}$. Both examples are constructed based on a network topology consisting of three disjoint subsets of nodes $V_C = \{c_1, c_2, \dots, c_n\}$, $V_M = \{m_1, m_2, \dots, m_n\}$ and $V_P = \{p_1, p_2, \dots, p_n\}$. Figure 3(a) illustrates the first example. Due to space limitations, only the network latencies between some node pairs are shown in the figure. Table I lists the network latencies between all pairs of nodes, where $\alpha \geq 1$ and $0 < \varepsilon < 1$. It is easy to verify that the network latencies satisfy the α -triangle inequality. Suppose that a client is located at each node in V_C , and all the nodes in $V_M \cup V_P$ are candidate server locations. Then, the optimal solution to the DIA server provisioning problem is to place servers at all the n nodes in V_P . In this case, each client c_i connects to the server at p_i , and this gives the minimum total interaction path length of

$$2n \sum_{i=1}^n d(c_i, p_i) + \sum_{i=1}^n \sum_{j=1}^n d(p_i, p_j) = 2n^2(1 + \varepsilon) + n(n - 1)\varepsilon.$$

On the other hand, the n -median placement is to place servers at all the n nodes in V_M . Under such placement, the nearest server of each client c_i is at m_i , so the total latency from the clients to their nearest servers is minimized at n . In this case, the total interaction path length is given by

$$\begin{aligned} 2n \sum_{i=1}^n d(c_i, m_i) + \sum_{i=1}^n \sum_{j=1}^n d(m_i, m_j) \\ = 2n^2 + n(n - 1)(\alpha^3 + 3\alpha^2). \end{aligned}$$

Therefore, the ratio between the results of the n -median and optimal placements is

$$\frac{2n^2 + n(n - 1)(\alpha^3 + 3\alpha^2)}{2n^2(1 + \varepsilon) + n(n - 1)\varepsilon}.$$

As the number of clients n goes toward infinity and ε approaches 0, this ratio approaches

$$\begin{aligned} \lim_{n \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{2n^2 + n(n - 1)(\alpha^3 + 3\alpha^2)}{2n^2(1 + \varepsilon) + n(n - 1)\varepsilon} \\ = \lim_{n \rightarrow \infty} \frac{2n^2 + n(n - 1)(\alpha^3 + 3\alpha^2)}{2n^2} = \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1. \end{aligned}$$

Fig. 3. The approximation ratio of the k -median placement is tight.

TABLE I
NETWORK LATENCY MATRIX FOR TIGHT EXAMPLE OF $\frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1$.

$\forall i, j \ (i \neq j)$	c_i	m_i	p_i
c_i	0	1	$1 + \varepsilon$
m_i	1	0	2α
p_i	$1 + \varepsilon$	2α	0
c_j	$\alpha^2 + \alpha$	$2\alpha^2 + \alpha$	$\alpha(1 + 2\varepsilon)$
m_j	$2\alpha^2 + \alpha$	$\alpha^3 + 3\alpha^2$	$\alpha^2 + \alpha$
p_j	$\alpha(1 + 2\varepsilon)$	$\alpha^2 + \alpha$	ε

TABLE II
NETWORK LATENCY MATRIX FOR TIGHT EXAMPLE OF $\frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}$.

$\forall i, j \ (i \neq j)$	c_i	m_i	p_i
c_i	0	1	$1 + \varepsilon$
m_i	1	0	2α
p_i	$1 + \varepsilon$	2α	0
c_j	$2\alpha^2$	$2\alpha^2 + 3\alpha$	2α
m_j	$2\alpha^2 + 3\alpha$	$2\alpha^3 + 3\alpha^2$	$2\alpha^2 + \alpha$
p_j	2α	$2\alpha^2 + \alpha$	1

Similarly, Figure 3(b) illustrates the second example and Table II shows the network latencies between all pairs of nodes in this example, where $\alpha \geq 1$ and $0 < \varepsilon < 1$. It can be checked that the network latencies satisfy the α -triangle inequality. Again, suppose that a client is located at each node in V_C , and all the nodes in $V_M \cup V_P$ are candidate server locations. The optimal placement for the DIA server provisioning problem is to place a server at each node in V_P so that each client c_i connects to the server at p_i . This gives the minimum total interaction path length of

$$2n \sum_{i=1}^n d(c_i, p_i) + \sum_{i=1}^n \sum_{j=1}^n d(p_i, p_j) = 2n^2(1 + \varepsilon) + n(n - 1).$$

Meanwhile, the n -median placement is to select all the n nodes in V_M as server locations. In this case, each client c_i connects

to its nearest server at m_i . Thus, the total interaction path length is

$$\begin{aligned} & 2n \sum_{i=1}^n d(c_i, m_i) + \sum_{i=1}^n \sum_{j=1}^n d(m_i, m_j) \\ &= 2n^2 + n(n - 1)(2\alpha^3 + 3\alpha^2). \end{aligned}$$

Therefore, the ratio between the results of the n -median and optimal placements can be made arbitrarily close to

$$\begin{aligned} & \lim_{n \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{2n^2 + n(n - 1)(2\alpha^3 + 3\alpha^2)}{2n^2(1 + \varepsilon) + n(n - 1)} \\ &= \lim_{n \rightarrow \infty} \frac{2n^2 + n(n - 1)(2\alpha^3 + 3\alpha^2)}{2n^2 + n(n - 1)} = \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}, \end{aligned}$$

as n goes toward infinity and ε approaches 0.

V. GREEDY ALGORITHM

The cubical growth of its approximation ratio with α implies that the k -median placement may not perform well in networks with triangle inequality violations. Moreover, it is difficult to find the k -median placement due to its NP-hardness [13]. In this section, we propose an efficient GREEDY algorithm that has much better approximability than the k -median placement. The computation of GREEDY is based simply on the network latencies between clients and candidate server locations, which can be acquired with existing tools like ping and King [24].

Without loss of generality, the GREEDY algorithm takes a parameter k that indicates the cap on the number of server locations to select. In the case where there is no limit on the number of server locations to select, k may simply be set to the number of candidate server locations. The GREEDY algorithm starts with an empty set of server locations S . In each iteration, the algorithm adds to S an unselected candidate location that leads to the maximum reduction in the total interaction path length. The algorithm terminates when no new server location can be further selected to reduce the total interaction path length or the number of server locations selected reaches k .

Algorithm 1 shows the pseudo code of the GREEDY algorithm. In each iteration, the algorithm evaluates each unselected candidate location $s \in Z \setminus S$ and calculates the total interaction path length D if s is added to S (lines 6 to 17). The candidate location that results in the minimum value of D and the corresponding D value are recorded in s^* and D^* respectively (lines 16 to 17). At the end of each iteration, the algorithm keeps the nearest servers of all clients under the current placement S in an array $n[\cdot]$ (lines 21 to 23). When evaluating a candidate server location s in the next iteration, we can find out the nearest server of each client c by simply comparing the latency from c to $n[c]$ with the latency from c to s (lines 11 to 14). In this way, the efficiency of the algorithm is improved by avoiding repeatedly comparing the latencies between each client and all the selected server locations.

Suppose that under a server placement S , the number of clients connecting to each server $s_i \in S$ is $m(s_i)$. Then, the total interaction path length between all client pairs can be rewritten as

$$\begin{aligned}
& \sum_{c \in C} \sum_{c' \in C} \left(d(c, n(c, S)) + d(n(c, S), n(c', S)) \right. \\
& \quad \left. + d(n(c', S), c') \right) \\
&= |C| \cdot \sum_{c \in C} d(c, n(c, S)) + |C| \cdot \sum_{c' \in C} d(n(c', S), c') \\
& \quad + \sum_{c \in C} \sum_{c' \in C} d(n(c, S), n(c', S)) \\
&= 2|C| \sum_{c \in C} d(c, n(c, S)) \\
& \quad + \sum_{s_i \in S} \sum_{s_j \in S} m(s_i) \cdot m(s_j) \cdot d(s_i, s_j). \tag{8}
\end{aligned}$$

Thus, to calculate the total interaction path length resulting from adding a candidate location s to S , the GREEDY algorithm first counts the number of clients connecting to each

```

1:  $S \leftarrow \emptyset, D_{\text{current}} \leftarrow \infty$ ;
2: for all  $c \in C$  do
3:    $n[c] \leftarrow \emptyset$ ; //nearest server to  $c$ 
4: for  $i = 1$  to  $k$  do
5:    $D^* \leftarrow \infty$ ;
6:   for all  $s \in Z \setminus S$  do
7:     for all  $s_i \in S \cup \{s\}$  do
8:        $a[s_i] \leftarrow 0$ ; //number of clients connecting to  $s_i$ 
9:        $D \leftarrow 0$ ; //total interaction path length when evaluating  $s$ 
10:    for all  $c \in C$  do
11:      if  $n[c] = \emptyset$  or  $d(c, s) < d(c, n[c])$  then
12:         $a[s] \leftarrow a[s] + 1$ ;  $D \leftarrow D + 2 \cdot |C| \cdot d(c, s)$ ;
13:      else
14:         $a[n[c]] \leftarrow a[n[c]] + 1$ ;  $D \leftarrow D + 2 \cdot |C| \cdot d(c, n[c])$ ;
15:       $D \leftarrow D + \sum_{s_i \in S \cup \{s\}} \sum_{s_j \in S \cup \{s\}} a[s_i] \cdot a[s_j] \cdot d(s_i, s_j)$ ;
16:      if  $D < D^*$  then
17:         $D^* \leftarrow D$ ;  $s^* \leftarrow s$ ;
18:    if  $D^* \geq D_{\text{current}}$  then
19:      break; //stop if cannot further reduce interaction path length
20:     $S \leftarrow S \cup \{s^*\}$ ;  $D_{\text{current}} \leftarrow D^*$ ;
21:    for all  $c \in C$  do
22:      if  $n[c] = \emptyset$  or  $d(c, s^*) < d(c, n[c])$  then
23:         $n[c] \leftarrow s^*$ ; //update the nearest server of client  $c$ 
24:  output  $S$ ;

```

Algorithm 1: GREEDY Server Provisioning Algorithm

server $s_i \in S \cup \{s\}$ and records it in $a[s_i]$. Then, the total interaction path length is calculated according to the above formula (8) (lines 10 to 15). The computational complexity of the total interaction path length is $O(|C| + k^2)$. Since the number of candidate server locations to evaluate in each iteration is capped by $|Z|$ (where Z is the set of candidate server locations) and there are at most k iterations, the total time complexity of the GREEDY algorithm is $O(k|Z|(|C| + k^2))$.

We can show that the above GREEDY algorithm has an approximation ratio of 2α for networks with the α -triangle inequality, irrespective of whether the choices of server locations are restricted and the cap on the number of server locations to select (if any).

Theorem 3. *For networks satisfying the α -triangle inequality ($\alpha \geq 1$), the GREEDY algorithm produces a total interaction path length within 2α times of that in an optimal server placement.*

Proof: Suppose that an optimal server placement selects h server locations $s_1, s_2, \dots$, and s_h . For each location s_i , denote by C_i the set of clients connecting to the server at s_i under the optimal placement. It is obvious that $\bigcup_{i=1}^h C_i = C$, where C is the set of all clients.

If only one server is placed at a location s in the network, all the clients would connect to that server and thus the total interaction path length is given by

$$\sum_{c \in C} \sum_{c' \in C} (d(c, s) + d(s, s) + d(s, c')) = 2|C| \cdot \sum_{c \in C} d(c, s).$$

Suppose that s^* is the first server location selected by the GREEDY algorithm. Since the GREEDY algorithm always selects the best known server location in each iteration, the total interaction path length resulting from selecting s^* cannot be longer than those resulting from selecting $s_1, s_2, \dots$, or s_h in the first iteration, i.e.,

$$2|C| \cdot \sum_{c \in C} d(c, s^*) \leq 2|C| \cdot \sum_{c \in C} d(c, s_1),$$

$$2|C| \cdot \sum_{c \in C} d(c, s^*) \leq 2|C| \cdot \sum_{c \in C} d(c, s_2),$$

.....

and

$$2|C| \cdot \sum_{c \in C} d(c, s^*) \leq 2|C| \cdot \sum_{c \in C} d(c, s_h).$$

Taking a weighted average of the interaction path lengths on the right sides of the above inequalities, we have

$$\begin{aligned}
2|C| \cdot \sum_{c \in C} d(c, s^*) &\leq \frac{1}{|C|} \cdot \sum_{i=1}^h \left(|C_i| \cdot 2|C| \cdot \sum_{c \in C} d(c, s_i) \right) \\
&= 2 \sum_{i=1}^h \left(|C_i| \cdot \sum_{c \in C} d(c, s_i) \right) \\
&= 2 \sum_{i=1}^h \left(|C_i| \cdot \left(\sum_{c \in C_i} d(c, s_i) + \sum_{j \neq i} \sum_{c \in C_j} d(c, s_i) \right) \right). \tag{9}
\end{aligned}$$

By the α -triangle inequality, for each client $c \in C_j$, we have

$$d(c, s_i) \leq \alpha \cdot (d(c, s_j) + d(s_j, s_i))$$

$$= (1 + (\alpha - 1)) \cdot (d(c, s_j) + d(s_j, s_i)).$$

Therefore,

$$\begin{aligned} \sum_{j \neq i} \sum_{c \in C_j} d(c, s_i) &\leq \sum_{j \neq i} \sum_{c \in C_j} (d(c, s_j) + d(s_j, s_i)) \\ &\quad + (\alpha - 1) \sum_{j \neq i} \sum_{c \in C_j} (d(c, s_j) + d(s_j, s_i)) \\ &\leq \sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) - \sum_{c \in C_i} d(c, s_i) + \sum_{j \neq i} \sum_{c \in C_j} d(s_j, s_i) \\ &\quad + (\alpha - 1) \sum_{j \neq i} \sum_{c \in C_j} (d(c, s_j) + d(s_j, s_i)). \end{aligned}$$

It follows from (9) that

$$\begin{aligned} &2|C| \cdot \sum_{c \in C} d(c, s^*) \\ &\leq 2 \cdot \sum_{i=1}^h \left(|C_i| \left(\sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) + \sum_{j \neq i} \sum_{c \in C_j} d(s_j, s_i) \right) \right. \\ &\quad \left. + (\alpha - 1) \sum_{j \neq i} \sum_{c \in C_j} (d(c, s_j) + d(s_j, s_i)) \right) \\ &= 2|C| \sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) + 2 \sum_{i=1}^h \sum_{j \neq i} |C_i| |C_j| d(s_i, s_j) \\ &\quad + (2\alpha - 2) \sum_{i=1}^h \left(|C_i| \cdot \left(\sum_{j \neq i} \sum_{c \in C_j} (d(c, s_j) + d(s_i, s_j)) \right) \right) \\ &\leq 2 \left(2|C| \sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) + \sum_{i=1}^h \sum_{j \neq i} |C_i| |C_j| d(s_i, s_j) \right) \\ &\quad + (2\alpha - 2) \left(\sum_{i=1}^h |C_i| \sum_{j \neq i} \sum_{c \in C_j} d(c, s_j) \right. \\ &\quad \left. + \sum_{i=1}^h \sum_{j \neq i} |C_i| |C_j| d(s_i, s_j) \right). \end{aligned}$$

According to (8), the total interaction path length under the optimal server placement is given by

$$OPT = 2|C| \sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) + \sum_{i=1}^h \sum_{j \neq i} |C_i| |C_j| d(s_i, s_j).$$

Thus,

$$\begin{aligned} &2|C| \cdot \sum_{c \in C} d(c, s^*) \\ &\leq 2 \cdot OPT + (2\alpha - 2) \left(|C| \sum_{j=1}^h \sum_{c \in C_j} d(c, s_j) \right. \\ &\quad \left. + \sum_{i=1}^h \sum_{j \neq i} |C_i| |C_j| d(s_i, s_j) \right) \\ &\leq 2 \cdot OPT + (2\alpha - 2) \cdot OPT \\ &= 2\alpha \cdot OPT. \end{aligned}$$

This implies that the total interaction path length on selecting the first server location in the GREEDY algorithm is

Fig. 4. The approximation ratio 2α of the GREEDY algorithm is tight.

within 2α times of that in an optimal server placement. Since the GREEDY algorithm adds a new server location in each subsequent iteration only if it reduces the total interaction path length, the total interaction path length eventually produced by GREEDY must also be within 2α times of that in an optimal placement. Hence, the theorem is proven. ■

The approximation ratio 2α of the GREEDY algorithm is tight. Figure 4 presents a tight example in which the network contains two node groups $G_A = \{c_1, c_2, c_3, s_1, s_2, s_3\}$ and $G_B = \{c_4, c_5, c_6, s_4, s_5, s_6\}$, and an additional node g . Each node c_i ($1 \leq i \leq 6$) has latency ε to node s_i , latency α to the rest nodes in the same group, and latency $\alpha(\frac{4}{3} + 3\delta)$ to all the nodes in the other groups, where $\delta, \varepsilon > 0$. Figure 4 illustrates the latencies between c_3 and the other nodes. In addition, the latency between two nodes s_i and s_j ($i \neq j$) is 1 if they are in the same group, and is $\frac{4}{3} + 3\delta$ otherwise. All the nodes in G_A and G_B have latency $\alpha(1 + \delta)$ to node g . It is easy to infer that the network latencies satisfy the α -triangle inequality when $\delta \leq \frac{2}{3}$ and $\varepsilon \leq \alpha(1 + \alpha)$. Suppose that a client is located at each node c_i ($1 \leq i \leq 6$), nodes $g, s_1, s_2, \dots, s_6$ are candidate server locations, and there is no limit on the number of server locations to select. Then, the optimal placement is to select nodes $s_1, s_2, \dots, s_6$ as server locations. Under such server placement, each client c_i connects to the server at s_i , so the total interaction path length is

$$\begin{aligned} OPT &= 12 \cdot \sum_{i=1}^6 d(c_i, s_i) + \sum_{i=1}^6 \sum_{j \neq i} d(s_i, s_j) \\ &= 12 \cdot 6\varepsilon + 12 \cdot 1 + 18 \cdot \left(\frac{4}{3} + 3\delta \right) = 36 + 54\delta + 72\varepsilon. \end{aligned}$$

On the other hand, the first server location selected by the GREEDY algorithm is g . This is because placing a server at g leads to the total interaction path length of $D_1 = 12 \cdot \sum_{i=1}^6 d(c_i, g) = 12 \cdot 6 \cdot \alpha(1 + \delta) = 72\alpha + 72\alpha\delta$, whereas if a server is placed at any node s_i , the total interaction path length is $12 \cdot (\varepsilon + 2 \cdot \alpha + 3 \cdot \alpha(\frac{4}{3} + 3\delta)) = 72\alpha + 108\alpha\delta + 12\varepsilon > D_1$. In the second iteration of the GREEDY algorithm, if a new server location s is further selected from $\{s_1, s_2, \dots, s_6\}$, all the three clients in s 's group would connect to the server at s and the other three clients remain connecting to the server at g . Therefore, the total interaction path length becomes $D_2 = 12 \cdot (\varepsilon + 2 \cdot \alpha + 3 \cdot \alpha(1 + \delta)) + 18 \cdot \alpha(1 + \delta) = 78\alpha + 54\alpha\delta + 12\varepsilon$. When $\delta < \frac{1}{3} + \frac{2\alpha}{3}\varepsilon$, we have $D_2 > D_1$. Thus, the GREEDY

algorithm does not add any new server location and terminates at the second iteration. As a result, the final set of server locations selected is $\{g\}$. When both δ and ε approach 0, the ratio between the total interaction path lengths produced by the GREEDY algorithm and the optimal server placement is

$$\lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \frac{D_1}{OPT} = \lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \frac{72\alpha + 72\alpha\delta}{36 + 54\delta + 72\varepsilon} = 2\alpha.$$

For any $\alpha \geq 1$, we have

$$2\alpha \leq \frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 \leq \max\left\{\frac{1}{2}\alpha^3 + \frac{3}{2}\alpha^2 + 1, \frac{2}{3}\alpha^3 + \alpha^2 + \frac{2}{3}\right\}.$$

Therefore, the approximation ratio of the GREEDY algorithm (Theorem 3) is much lower than that of the k -median server placement (Theorem 2). Asymptotically, the approximation ratio of GREEDY grows just linearly with α , whereas the approximation ratio of the k -median placement grows cubically with α . This indicates that the GREEDY algorithm is far more resilient to the extent of triangle inequality violations than the k -median placement. Moreover, due to its incremental nature, the GREEDY algorithm can also be used directly to find substitute server locations in the presence of server failures.

VI. EXPERIMENTAL EVALUATION

Besides theoretical approximability analysis, we have also evaluated the performance of the proposed GREEDY algorithm using different latency datasets collected from the real Internet and synthetic networks, including Meridian [25], PeerWise [26] and p2psim-kingdata [27]. The experimental results using these datasets show similar performance trends. Due to space limitations, we present in this section the results using the Meridian dataset, which is the largest publicly available real Internet latency dataset to our knowledge. The Meridian dataset contains the pairwise latency measurements between 2500 nodes. The measurements for some node pairs are not available in the dataset. We discard the nodes involved in the missing measurements, and use the complete pairwise latency matrix among the remaining 1796 nodes as the simulated network. The simulated network does not satisfy the triangle inequality. To examine the extent of triangle inequality violations, we check all the triangles among the nodes. For each triangle (u, v, w) , we compute the ratio of $\frac{d(u, v)}{d(u, w) + d(w, v)}$, where $d(u, v)$ is the longest edge among the three. If the ratio does not exceed 1, the triangle satisfies the triangle inequality. Figure 5 shows the cumulative distribution of the ratio. It can be seen that 78% of the triangles satisfy the triangle inequality and for the remaining triangles, their extent of triangle inequality violations has a heavy-tailed distribution.

We compare our proposed algorithm with the k -median and other baseline server placements. To quantify the relative performance of the algorithms, we normalize the total interaction path lengths produced by different algorithms by a theoretical lower bound, which is derived as follows. Given a set Z of candidate server locations, the shortest possible interaction path length between two clients c_i and c_j is given by $\min_{s_a, s_b \in Z} (d(c_i, s_a) + d(s_a, s_b) + d(s_b, c_j))$. Therefore, the total length of interaction paths between all client pairs has a

Fig. 5. Cumulative distribution of the extent of triangle inequality violations for all triangles in the Meridian dataset

lower bound of

$$\sum_{c_i \in C} \sum_{c_j \in C} \left(\min_{s_a, s_b \in Z} (d(c_i, s_a) + d(s_a, s_b) + d(s_b, c_j)) \right). \quad (10)$$

Note that this bound is a super-optimum, which may not be achievable by any real server placement, because it does not enforce each client to connect to its nearest server or even a single server through which it interacts with all the other clients. The total interaction path length normalized by the above bound shall be called the *normalized interactivity*.

The k -median and k -center placements have been considered to be closely related to the server placement for DIAs [11], [12]. For the purpose of comparison, we implement a k -median heuristic [8] and k -center heuristic [9] in our experiments. Both heuristics work for k iterations. In each iteration, the k -median heuristic adds a new server location that results in the largest reduction in the total latency from the clients to their nearest servers. Similarly, the k -center heuristic adds a new server location that results in the lowest value of the maximum latency from the clients to their nearest servers. In addition, we also implement a k -favourable heuristic that considers not only the client-to-server latencies but also the inter-server latencies. The k -favourable heuristic evolves from the above derivation of the lower bound. Specifically, for each pair of clients c_i and c_j , the candidate server locations giving rise to the shortest possible interaction path between c_i and c_j are recorded. Then, the k -favourable heuristic selects the top k candidate server locations that are most frequently involved in the shortest possible interaction paths between all client pairs. Intuitively, these locations have high potential to optimize the interactivity performance of DIAs.

A. Performance Comparison for Different Algorithms

We start by not placing any restriction on the available choices of server locations. Specifically, we assume that all 1796 nodes are candidate server locations, and a client is located at each of the nodes in the network. Figure 6 shows the performance of the four algorithms for different caps on the number of server locations to select. As can be seen, server placement has great impacts on the interactivity performance.

When only one server location is selected, the GREEDY algorithm degenerates to the 1-median placement. Figure 6 shows that by selecting additional server locations, GREEDY substantially reduces the network latency involved in the interaction between clients over the 1-median placement. This demonstrates the great potential of distributed server

Fig. 6. Normalized interactivity of different algorithms when all nodes are candidate server locations

Fig. 7. Normalized interactivity of different algorithms for different numbers of candidate server locations

deployment to enhance the interactivity performance of DIAs. However, deploying more servers does not always result in better interactivity between clients—the server locations must be chosen carefully. For example, in the k -median and k -center heuristics, the latencies from the clients to their connected servers normally decrease with increasing number of server locations selected since each client connects to its nearest server. Thus, the k -median and k -center heuristics always place servers up to the cap number. However, aggressively reducing client-to-server latencies alone may result in longer latencies between the servers of different clients and consequently increase the interaction path length between clients. As seen from Figure 6, the interaction path lengths of the k -median and k -center heuristics do not always reduce with increasing number of server locations selected because they do not consider inter-server latencies in server provisioning. By taking inter-server latencies into account, the k -favourable heuristic performs much better than the k -median and k -center heuristics when more servers are allowed to be selected. However, its interactivity performance is still considerably worse than our GREEDY algorithm. GREEDY stops selecting new server locations when the total interaction path length cannot be further reduced. Therefore, its interactivity performance does not deteriorate with a higher cap on the number of server locations to select. Figure 6 shows that the GREEDY algorithm significantly outperforms the other three placements.

Next, we restrict the choices of server locations in the network. We assume that a client is located at each of 896 randomly selected nodes. Then, we randomly choose subsets of 75, 150, 300, 600 and 900 nodes from the remaining nodes in the network as the candidate server locations. In this experiment, no limit is set on the number of server locations

for our GREEDY algorithm to select. We run the GREEDY algorithm until it terminates and record the number of server locations actually selected by GREEDY. This number is then used as the number of server locations to select in executing the other three heuristics. In this way, all the four algorithms select the same number of server locations in the network to allow for fair comparison. For each set size, we perform 1000 simulation runs using 1000 different sets of candidate server locations chosen at random. Figure 7 shows the average performance of the algorithms together with the 90th and 10th percentile results for different numbers of candidate server locations. Here, the total interaction path lengths are normalized by the lower bound (10) derived using the respective sets of candidate server locations chosen. Thus, the results of various simulation runs are normalized by different lower bound values. It can be seen from Figure 7 that our GREEDY algorithm consistently results in much better interactivity than the other three placements. It can also be observed that the normalized interactivity of the algorithms generally improves when the number of candidate server locations reduces. This is because fewer candidate server locations imply smaller search space, so that it is more likely to obtain a good solution within certain number of attempts.

We also record the running times of different algorithms. The results show that our GREEDY algorithm has very acceptable running time. Please refer to Appendix B in the supplementary material for details.

B. Impact of Server Capacity

All the above experiments have assumed that there is no capacity limitation of the servers. This is the case when we have sufficient server capacities available to be allocated on demand at the candidate server locations. In this section, we further study the impact of server capacity on the performance of the server placement algorithms. Suppose that we have a certain number of servers to allocate, each of which has a given capacity in terms of the maximum number of clients that can connect to it. To place the servers, we first execute each server placement algorithm to select the server locations. Then, we allocate the servers to each selected server location approximately in proportion to the number of clients having it as their nearest server locations. The total capacity of the servers allocated to a server location forms the capacity of this server location. When a server location has been connected by enough clients and filled to its capacity, it would not accept any new connections from other clients, and the affected clients have to turn to their next nearest server locations with spare capacities.

We assume 896 clients and 300 candidate server locations in each simulation run. First, we study the impact of the individual server capacity by testing 36, 18, 12 and 9 servers each with the capacity of 25, 50, 75 and 100 respectively. Note that under these settings, the total server capacity is 900 and just suffices to accommodate all the 896 clients. Figure 8 shows the performance of different server placement algorithms. It can be seen that the GREEDY algorithm achieves the best interactivity for all the cases tested. Figure 8 also

Fig. 8. Normalized interactivity of different algorithms for different individual server capacities

Fig. 9. Normalized interactivity of different algorithms for different total server capacities

shows the trend that the performance of each algorithm slightly deteriorates when the individual server capacity increases. This is because the number of servers needed to support all the clients decreases with increasing individual server capacity. As a result, the number of locations that are allocated servers decreases accordingly, raising the number of clients that cannot connect to their preferred nearest server locations.

Next, we study the impact of the total available server capacity. In this experiment, the capacity of each server is fixed at 50. Figure 9 shows the performance results for different numbers of servers available for allocation. Again, the GREEDY algorithm outperforms the other server placements. We can also observe from the figure that the interactivity performance of each server placement generally improves with increasing number of servers available. This is because a larger number of servers increase the maximum number of connections that each server location can accept. Thus, more clients can connect to their preferred nearest server locations.

C. Impact of Dynamics in Network Conditions

The network latency in real networks may vary over time due to dynamics in network conditions. In this section, we study the impact of the changes in network latency on the performance of the server provisioning algorithms. Since the Meridian dataset does not contain the measurements of network latencies at different times, we use the PlanetLab All Pairs Pings dataset [28] collected on 6 June 2005 in this experiment. This dataset contains the periodic pairwise latency measurements between PlanetLab nodes at intervals of 15 minutes over a one-day time span. There are 95 measurements in the dataset and 193 nodes consistently involved in all the measurements. Thus, we simulate a network of 193 nodes, using these data to emulate the network latencies among the nodes. The latency between each node pair in the simulated network changes according to the consecutive measurements in the dataset.

We randomly select 80 nodes as candidate server locations, and assume that a client is located at each of the remaining 113 nodes. We decide the server placement based on the latency data of the first measurement and keep the server placement unchanged throughout the simulation run. Two scenarios of client connections are then simulated. In the first scenario, each client connects to a fixed server (the nearest server as indicated by the latency data of the first measurement) throughout the simulation run. That is, the connected servers of the clients do not change with the network latency. In the second scenario, each client always connects to the nearest server based on the latest latency measurement. As a result, the connected servers of the clients change over time along with the variation of network latency. The theoretical lower bound on the total interaction path length is recalculated at each latency measurement for computing the normalized interactivity. We have performed simulation with many different sets of candidate server locations chosen at random and observed similar performance trends. Figure 10 shows the performance results for a sample set of candidate server locations. As can be seen, our GREEDY algorithm achieves the best interactivity in both scenarios despite the changes in network latency. In addition, the GREEDY algorithm also has the minimum fluctuation in the normalized interactivity among the four server provisioning algorithms. This implies that the GREEDY algorithm is more resilient to the variation of network latency.

Comparing Figures 10(a) and 10(b), it can be observed that the performance of each server provisioning algorithm is slightly improved when the clients adjust their connections according to the up-to-date latency measurements. These adjustments reduce the latency between each client and its nearest server.

In summary, the results of this experiment show that good server placement for DIAs is relatively stable over time and adaptive client connections can assist to compensate for the interactivity loss due to short-term variation of network latency.

VII. CONCLUSION

In this paper, we have investigated the server provisioning problem for DIAs with the objective of reducing the network latency involved in the interaction between clients. From the computability perspective, the mutual interaction feature makes server provisioning for DIAs much more challenging than that for web content delivery. With the client-to-server latency as the sole optimization objective, placing more servers in the network can only improve the web access performance. Nevertheless, deploying more servers in DIAs does not necessarily imply shorter network latency in the interaction between clients. We have shown that the server provisioning problem for DIAs is NP-hard even if there is no limit on the number of servers that can be placed in the network. We have also proposed a GREEDY algorithm that achieves a much lower approximation ratio than the classical k -median server placement. Experimental evaluations using real Internet latency data show that aggressively reducing

(a) The connected servers of clients are fixed.

(b) The connected servers of clients change along with the variation of network latency.

Fig. 10. Performance of different algorithms with dynamic network latency

the client-to-server latency alone may considerably increase the latency between servers and thus make the interactivity between clients far worse than optimum. Our proposed GREEDY algorithm substantially outperforms the k -median and other baseline server placements and is also more resilient to the dynamics in network latency.

ACKNOWLEDGEMENTS

This work is supported by Singapore Ministry of Education Academic Research Fund Tier 2 under Grant MOE2013-T2-2-067.

REFERENCES

- [1] H. Zheng and X. Tang. On Server Provisioning for Distributed Interactive Applications. In *Proc. IEEE ICDCS*, pages 500–509, 2013.
- [2] M. Claypool and K. Claypool. Latency Can Kill: Precision and Deadline in Online Games. In *Proc. ACM MMSys*, pages 215–222, 2010.
- [3] R. B. Jennings, E. M. Nahum, D. P. Olshefski, D. Saha, Z.-Y. Shae and C. Waters. A Study of Internet Instant Messaging and Chat Protocols. *IEEE Network*, 20(4):16–21, 2006.
- [4] Agustina, F. Liu, S. Xia, H. Shen and C. Sun. CoMaya: Incorporating Advanced Collaboration Capabilities into 3D Digital Media Design Tools. In *Proc. ACM CSCW*, pages 5–8, 2008.
- [5] L. Ahmad, A. Boukerche, A. Al Hamidi, A. Shadid and R. Pazzi. Web-Based e-Learning in 3D Large Scale Distributed Interactive Simulations using HLA/RTI. In *Proc. IEEE IPDPS*, pages 1–4, 2008.
- [6] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica and M. Zaharia. A View of Cloud Computing. *Communications of the ACM*, 53(4):50–58, 2010.
- [7] F. Safaei, P. Boustead, C. D. Nguyen, J. Brun and M. Dowlatshahi. Latency-Driven Distribution: Infrastructure Needs of Participatory Entertainment Applications. *IEEE Communications Magazine*, 43(5):106–112, 2005.
- [8] L. Qiu, V. N. Padmanabhan and G. M. Voelker. On the Placement of Web Server Replicas. In *Proc. IEEE INFOCOM*, pages 1587–1596, 2001.
- [9] E. Cronin, S. Jamin, J. Cheng, A. R. Kurc, D. Raz and Y. Shavitt. Constrained Mirror Placement on the Internet. *IEEE Journal on Selected Areas in Communications*, 20(7):1369–1382, 2002.
- [10] N. Laoutaris, G. Smaragdakis, K. Oikonomou, I. Stavrakakis and A. Bestavros. Distributed Placement of Service Facilities in Large-Scale Networks. In *Proc. IEEE INFOCOM*, pages 2144–2152, 2007.
- [11] K.-W. Lee, B.-J. Ko and S. Calo. Adaptive Server Selection for Large Scale Interactive Online Games. *Computer Networks*, 49(1):84–102, 2005.
- [12] M. Kwok. *Performance Analysis of Distributed Virtual Environments*. PhD thesis, University of Waterloo, 2006.
- [13] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- [14] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala and V. Pandit. Local Search Heuristic for k -Median and Facility Location Problems. In *Proc. ACM STOC*, pages 21–29, 2001.
- [15] M. Chrobak, C. Kenyon and N. Young. The Reverse Greedy Algorithm for the Metric k -Median Problem. *Information Processing Letters*, 97(2):68–72, 2006.
- [16] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.
- [17] L. Zhang and X. Tang. Optimizing Client Assignment for Enhancing Interactivity in Distributed Interactive Applications. *IEEE/ACM Transactions on Networking*, 20(6):1707–1720, 2012.
- [18] S. D. Webb, S. Soh and W. Lau. Enhanced Mirrored Servers for Network Games. In *Proc. ACM NetGames*, pages 117–122, 2007.
- [19] C. Ding, Y. Chen, T. Xu and X. Fu. CloudGPS: A Scalable and ISP-Friendly Server Selection Scheme in Cloud Computing Environments. In *Proc. IEEE/ACM IWQoS*, 2012.
- [20] C. Lumezanu, R. Baden, N. Spring and B. Bhattacharjee. Triangle Inequality and Routing Policy Violations in the Internet. In *Proc. Passive and Active Measurement Conference*, pages 45–54, 2009.
- [21] F. Dabek, R. Cox, F. Kaashoek and R. Morris. Vivaldi: A Decentralized Network Coordinate System. In *Proc. ACM SIGCOMM*, pages 15–26, 2004.
- [22] M. Costa, M. Castro, A. Rowstron and P. Key. PIC: Practical Internet Coordinates for Distance Estimation. In *Proc. IEEE ICDCS*, pages 178–187, 2004.
- [23] H. Zheng, E. K. Lua, M. Pias and T. G. Griffin. Internet Routing Policies and Round-Trip-Times. In *Proc. Passive and Active Measurement Conference*, pages 236–250, 2005.
- [24] K. P. Gummadi, S. Saroiu and S. D. Gribble. King: Estimating Latency between Arbitrary Internet End Hosts. In *Proc. ACM SIGCOMM Workshop on Internet Measurement*, pages 5–18, 2002.
- [25] The Meridian Latency Data Set. <http://www.cs.cornell.edu/People/egs/meridian/>.
- [26] The PeerWise Data Set. <http://www.cs.umd.edu/projects/peerwise/>.
- [27] The p2psim King Data Set. <http://pdos.csail.mit.edu/p2psim/kingdata/>.
- [28] The PlanetLab All Pairs Pings Data Set. <http://pdos.csail.mit.edu/~trib/projects.html>.

Hanying Zheng received the BSc degree in computer science and technology from Sun Yat-Sen University, China. He is currently a PhD candidate in computer science at Nanyang Technological University, Singapore. His research interests include computer and communication networks, distributed systems, mobile computing and cloud computing.

Xueyan Tang is currently an associate professor in the School of Computer Engineering at Nanyang Technological University, Singapore. He has served as an associate editor of IEEE Transactions on Parallel and Distributed Systems. His research interests include distributed systems, mobile and pervasive computing, and wireless sensor networks. He is a senior member of the IEEE.