

Benefit-Based Data Caching in Ad Hoc Networks

Bin Tang, *Member, IEEE*, Himanshu Gupta, *Member, IEEE*, and Samir R. Das, *Member, IEEE*

Abstract—Data caching can significantly improve the efficiency of information access in a wireless ad hoc network by reducing the access latency and bandwidth usage. However, designing efficient distributed caching algorithms is nontrivial when network nodes have limited memory. In this article, we consider the cache placement problem of minimizing total data access cost in ad hoc networks with multiple data items and nodes with limited memory capacity. The above optimization problem is known to be NP-hard. Defining *benefit* as the reduction in total access cost, we present a polynomial-time centralized approximation algorithm that provably delivers a solution whose benefit is at least 1/4 (1/2 for uniform-size data items) of the optimal benefit. The approximation algorithm is amenable to localized distributed implementation, which is shown via simulations to perform close to the approximation algorithm. Our distributed algorithm naturally extends to networks with mobile nodes. We simulate our distributed algorithm using a network simulator (*ns2*) and demonstrate that it significantly outperforms another existing caching technique (by Yin and Cao [33]) in all important performance metrics. The performance differential is particularly large in more challenging scenarios such as higher access frequency and smaller memory.

Index Terms—Caching placement policy, ad hoc networks, algorithm/protocol design and analysis, simulations.

1 INTRODUCTION

AD hoc networks are multihop wireless networks of small computing devices with wireless interfaces. The computing devices could be conventional computers (for example, PDA, laptop, or PC) or backbone routing platforms or even embedded processors such as sensor nodes. The problem of optimal placement of caches to reduce overall cost of accessing data is motivated by the following two defining characteristics of ad hoc networks. First, the ad hoc networks are multihop networks without a central base station. Thus, remote access of information typically occurs via multihop routing, which can greatly benefit from caching to reduce access latency. Second, the network is generally resource constrained in terms of channel bandwidth or battery power in the nodes. Caching helps in reducing communication, which results in savings in bandwidth, as well as battery energy. The problem of cache placement is particularly challenging when each network node has a limited memory to cache data items.

In this paper, our focus is on developing efficient caching techniques in ad hoc networks with memory limitations. Research into data storage, access, and dissemination techniques in ad hoc networks is not new. In particular, these mechanisms have been investigated in connection with sensor networking [14], [26], peer-to-peer networks [1], [18], mesh networks [17], world wide Web [25], and even

more general ad hoc networks [12], [33]. However, the presented approaches have so far been somewhat “ad hoc” and empirically based, without any strong analytical foundation. In contrast, the theory literature abounds in analytical studies into the optimality properties of caching and replica allocation problems (see, for example, [3]). However, distributed implementations of these techniques and their performances in complex network settings have not been investigated. It is even unclear whether these techniques are amenable to efficient distributed implementations. Our goal in this paper is to develop an approach that is both analytically tractable with a provable performance bound in a centralized setting and is also amenable to a natural distributed implementation.

In our network model, there are multiple data items; each data item has a server, and a set of clients that wish to access the data item at a given frequency. Each node carefully chooses data items to cache in its limited memory to minimize the overall access cost. Essentially, in this article, we develop efficient strategies to select data items to cache at each node. In particular, we develop two algorithms—a centralized approximation algorithm, which delivers a 4-approximation (2-approximation for uniform-size data items) solution, and a localized distributed algorithm, which is based on the approximation algorithm and can handle mobility of nodes and dynamic traffic conditions. Using simulations, we show that the distributed algorithm performs very close to the approximation algorithm. Finally, we show through extensive experiments on *ns2* [10] that our proposed distributed algorithm performs much better than a prior approach over a broad range of parameter values. Ours is the first work to present a distributed implementation based on an approximation algorithm for the general problem of cache placement of multiple data items under memory constraint.

- B. Tang is with the Computer Science Department, Wichita State University, Wichita, KS 67260. E-mail: bintang@cs.wichita.edu.
- H. Gupta and S.R. Das are with the Computer Science Department, Stony Brook University, Stony Brook, NY 11794. E-mail: {hgupta, samir}@cs.sunysb.edu.

Manuscript received 1 Oct. 2006; revised 20 June 2007; accepted 17 Sept. 2007; published online 3 Oct. 2007.

For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-0259-1006. Digital Object Identifier no. 10.1109/TMC.2007.70770.

The rest of the paper is organized as follows: In Section 2, we formally define the cache placement problem addressed in this paper and present an overview of the related work. In Section 3, we present our designed centralized approximation and distributed algorithms. Section 4 presents simulation results. We end with concluding remarks in Section 5.

2 CACHE PLACEMENT PROBLEM

In this section, we formally define the cache placement problem addressed in our article and discuss related work.

A multihop ad hoc network can be represented as an undirected graph $G(V, E)$, where the set of vertices V represents the nodes in the network, and E is the set of weighted edges in the graph. Two network nodes that can communicate directly with each other are connected by an edge in the graph. The edge weight may represent a link metric such as loss rate, delay, or transmission power. For the cache placement problem addressed in this article, there are multiple data items, and each data item is served by its server (a network node may act as a server for more than one data items). Each network node has limited memory and can cache multiple data items subject to its memory capacity limitation. The objective of our cache placement problem is to minimize the overall access cost. Below, we give a formal definition of the cache placement problem addressed in this article.

Problem formulation. Given a general ad hoc network graph $G(V, E)$ with p data items D_1, D_2, \dots, D_p , where a data item D_j is served by a server S_j , a network node may act as a server for multiple data items. For clarity of presentation, we assume uniform-size (occupying unit memory) data items for now. Our techniques easily generalize to nonuniform size data items, as discussed later. Each node i has a memory capacity of m_i units. We use a_{ij} to denote the access frequency with which a node i requests the data item D_j and d_{il} to denote the weighted distance between two network nodes i and l . The *cache placement problem* is to select a set of sets $M = \{M_1, M_2, \dots, M_p\}$, where M_j is a set of network nodes that store a copy of D_j , to minimize the total access cost:

$$\tau(G, M) = \sum_{i \in V} \sum_{j=1}^p a_{ij} \times \min_{l \in (\{S_j\} \cup M_j)} d_{il},$$

under the memory capacity constraint that

$$|\{M_j | i \in M_j\}| \leq m_i \quad \text{for all } i \in V,$$

which means that each network node i appears in at most m_i sets of M . The cache placement problem is known to be NP-hard [3].

Example 1. Fig. 1 illustrates the above described cache placement problem in a small ad hoc network. In Fig. 1, each graph edge has a unit weight. All the nodes have the same memory capacity of two pages, and the size of each data item is one memory page. Each of the nodes 1, 2, 3, 4, and 12 have one distinct data item to be served (as shown in the parenthesis with their node numbers). Each of the client nodes (9, 10, 11, 13, and 14) accesses

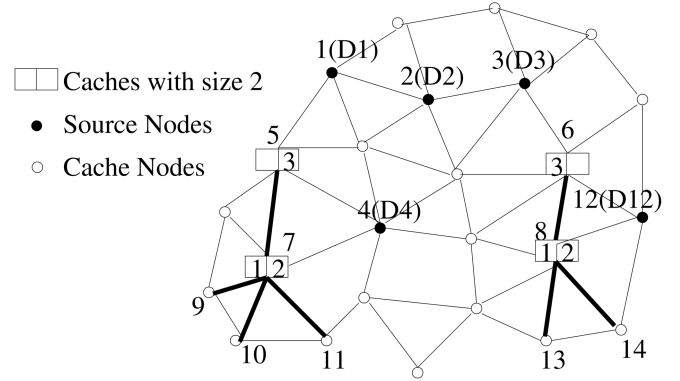


Fig. 1. Illustrating cache placement problem under memory constraint.

each of the data items D_1 , D_2 , and D_3 with unit access frequency. Fig. 1 shows that the nodes 5, 6, 7, and 8 have cached one or more data items and also shows the cache contents in those nodes. As indicated by the bold edges, the clients use the nearest cache node instead of the server to access a data item. The set of cache nodes of each data item are $M_1 = \{7, 8\}$, $M_2 = \{7, 8\}$, and $M_3 = \{5, 6\}$. One can observe that total access cost is 20 units for the given cache placement.

2.1 Related Work

Below, we categorize the prior work by the number of data items and network topology.

2.1.1 Single Data Item in General Graphs

The general problem of determining optimal cache placements in an arbitrary network topology has similarity to two problems in graph theory, namely, facility location problem and the k -median problem. Both the problems consider only a single facility type (data item) in the network. In the facility-location problem, setting up a cache at a node incurs a certain fixed cost, and the goal is to minimize the sum of total access cost and the setting-up costs of all caches without any constraint. On the other hand, the k -median problem minimizes the total access cost under the number constraint, that is, that at most k nodes can be selected as caches. Both problems are NP-hard, and a number of constant-factor approximation algorithms have been developed for each of the problems [8], [9], [15] under the assumption of triangular inequality of edge costs. Without the triangular inequality assumption, either problem is as hard as approximating the set cover [15], [21] and thus cannot be approximated better than $O(\log |V|)$ unless $\mathbf{P} = \mathbf{NP}$. Here, $|V|$ is the size of the network. In other related work, Nuggehalli et al. [22] formulate the caching problem in ad hoc networks as a special case of the connected facility location [27].

2.1.2 Single Data Item in Tree Topology

Several papers in the literature circumvent the hardness of the facility location and k -median problems by assuming that the network has a tree topology [4], [16], [19], [20], [28]. In particular, Tamir [28] and Vigneron et al. [16] design optimal dynamic programming polynomial algorithms for

the k -median problem in undirected and directed trees, respectively. In other works, Krishnan et al. [20] consider the placement of k “transparent” caches, Kalpakis et al. [19] consider a cost model involving reads, writes, and storage, and Bhattacharya et al. [4] present a distributed algorithm for sensor networks to reduce the total power expended. All of the above works consider only a single data time in a tree network topology.¹

2.1.3 Multiple Data Items

Hara [12] proposes three algorithms for cache placement of multiple data items in ad hoc networks. In the first approach, each node caches the items most frequently accessed by itself. The second approach eliminates replications among *neighboring* nodes introduced by the first approach. The third approach requires creation of “stable” groups to gather neighborhood information and determine caching placements. The first two approaches are largely localized and, hence, would fare very badly when the percentage of client nodes in the network is low or when the access frequencies are uniform. For the third approach, it is hard to find stable groups in ad hoc networks because of frequent failures and movements. All the above approaches assume the knowledge of access frequencies. In extensions of the above work, in [11] and [13], Hara generalizes the above approaches for push-based systems and updates, respectively. In other related works, Xu et al. [31] discuss placement of “transparent” caches in tree networks.

Our work on cache placement problem is most closely related to the works by Yin and Cao [33] and Baev and Rajaraman [3]. Yin and Cao [33] design and evaluate three simple distributed caching techniques, namely, *CacheData*, which caches the passing-by data item, *CachePath*, which caches the path to the nearest cache of the passing-by data item, and *HybridCache*, which caches the data item if its size is small enough or, else, caches the path to the data. They use LRU policy for cache replacement. To the best of our knowledge, the work by Yin and Cao [33] is the only one that presents a distributed cache placement algorithm in a multihop ad hoc network with a memory constraint at each node. Thus, we use the algorithms in [33] as a comparison point for our study.

Baev and Rajaraman [3] design a 20.5-approximation algorithm for the cache placement problem with uniform-size data items. For the nonuniform-size data items, they show that there is no polynomial-time approximation unless $\mathbf{P} = \mathbf{NP}$. They circumvent the nonapproximability by increasing the given node memory capacities by the size of the largest data item and generalize their 20.5-approximation algorithm. However, their approach (as noted by themselves) is not amenable to an efficient distributed implementation.

2.1.4 Our Work

In this article, we circumvent the nonapproximability of the cache placement problem by choosing to maximize the benefit (*reduction* in total access cost) instead of minimizing the total access cost. In particular, we design a simple

centralized algorithm that delivers a solution whose benefit is at least $1/4$ ($1/2$ for uniform-size data items) of the optimal benefit without using any more than the given memory capacities. To the best of our knowledge, ours and that by Baev and Rajaraman [3] are the only² works that present approximation algorithms for the general placement of cache placement for *multiple* data items in networks with *memory constraint*. However, as noted before, Baev and Rajaraman’s [3] approach is not amenable to an efficient distributed implementation, whereas our approximation algorithm yields a natural distributed implementation that is localized and shown (using ns2 simulations) to be efficient even in mobile and dynamic traffic conditions. Moreover, as stated in Theorem 2, our approximation result is an improvement over that in [3] when optimal access cost is at least $(1/40)$ th of the total access cost without the caches. Finally, unlike that in [3], we do not make the assumption of the cost function satisfying the triangular inequality.

A preliminary version of this article has appeared in [29]. The main additions in this article (compared to that in [29]) are

1. the approximation result for nonuniform size data items (Theorem 3),
2. details on the maintenance of the nearest and second-nearest cache table entries (Section 3.2),
3. data expiry and cache update models (Section 3.2), and
4. more extensive simulations.

3 CACHE PLACEMENT ALGORITHMS

In this section, we first present our centralized approximation algorithm. Then, we design its localized distributed implementation that performs very closely to the approximation algorithm in our simulations.

3.1 Centralized Greedy Algorithm (CGA)

The designed centralized algorithm is essentially a greedy approach, and we refer to it as CGA. CGA starts with all network nodes having all empty memory pages and, then, iteratively caches data items into memory pages, maximizing the benefit in a greedy manner at each step. Thus, at each step, the algorithm picks a data item D_j to cache into an empty memory page r of a network node such that the benefit of caching D_j at r is the maximum among all possible choices of D_j and r at that step. The algorithm terminates when all memory pages have been cached with data items.

For a formal analysis of CGA, we first define a set of variables A_{ijk} , where the selection of a variable A_{ijk} indicates that the k th memory page of node i has been selected for storage of data item D_j and reformulate the cache placement problem in terms of selection of A_{ijk} variables. Recall that for simplicity, we have assumed that each data item is of unit size and occupies one memory page of a node.

3.1.1 Problem Formulation Using A_{ijk}

Given a network graph $G(V, E)$, where each node $i \in V$ has a memory capacity of m_i pages and p data items D_1, \dots, D_p

1. Krishnan et al. [20] formulate the problem in general graphs but designs algorithms for tree topologies with single server.

2. Awerbuch et al. [2] present a competitive online algorithm but uses polylog-factor bigger memory capacity at nodes compared to the optimal.

in the network with the respective servers S_1, \dots, S_p , select a set Γ of variables A_{ijk} , where $i \in V$, $1 \leq j \leq p$ and $1 \leq k \leq m_i$. If $A_{ijk} \in \Gamma$ and $A_{ij'k} \in \Gamma$, then $j = j'$, such that the total access cost $\tau(G, \Gamma)$ (as defined below) is minimized. Note that the memory constraint is subsumed in the restriction on Γ that if $A_{ijk} \in \Gamma$, then $A_{ij'k} \notin \Gamma$ for any $j' \neq j$. The total access cost $\tau(G, \Gamma)$ for a selected set of variables can be easily defined as

$$\tau(G, \Gamma) = \sum_{j=1}^p \sum_{i \in V} a_{ij} \times \min_{l \in (\{S_j\} \cup \{i' | A_{ij'k} \in \Gamma\})} d_{il}.$$

Note that the set of cache nodes M_j that store a particular data item D_j can be easily derived from the selected set of variables Γ .

3.1.2 Centralized Greedy Algorithm

CGA works by iteratively selecting a variable A_{ijk} that gives the highest "benefit" at that stage. The benefit of adding a variable A_{ijk} into an already selected set of variables Γ is the reduction in the total access cost if the data item D_j is cached into the empty k th memory page of the network node i . The benefit of selecting a variable is formally defined below.

Definition 1 (benefit of selecting A_{ijk}). Let Γ denote the set of variables that have been already selected by the CGA at some stage. The benefit of a variable A_{ijk} ($i \in V$, $j \leq p$, and $k \leq m_i$) with respect to Γ is denoted as $\beta(A_{ijk}, \Gamma)$ and is defined as follows:

$$\beta(A_{ijk}, \Gamma) = \begin{cases} \text{Undefined} & \text{if } A_{ij'k} \in \Gamma, j' \neq j, \\ 0 & \text{if } A_{ij'k'} \in \Gamma, \\ \tau(G, \Gamma) - \tau(G, \Gamma \cup \{A_{ijk}\}) & \text{otherwise,} \end{cases}$$

where $\tau(G, \Gamma)$ is as defined before. The first condition of the above definition stipulates that if the k th memory page of the node i is not empty (that is, has already been selected to store another data item j' due to $A_{ij'k} \in \Gamma$), then the benefit $\beta(A_{ijk}, \Gamma)$ is undefined. The second condition specifies that the benefit of a variable A_{ijk} with respect to Γ is zero if the data item D_j has already been stored at some other memory page k' of the node i .

Algorithm 1 CGA.

```

BEGIN
   $\Gamma = \emptyset$ ;
  while (there is a variable  $A_{ijk}$  with defined benefit)
    Let  $A_{ijk}$  be the variable with maximum
     $\beta(A_{ijk}, \Gamma)$ .  $\Gamma = \Gamma \cup \{A_{ijk}\}$ ;
  end while;
  RETURN  $\Gamma$ ;
END.

```

The total runtime of CGA is $O(p^2|V|^3\bar{m})$, where $|V|$ is the size in the network, \bar{m} is the average number of memory pages in a node, and p is the total number of data items. Note that the number of iterations in the above algorithm is bounded by $|V|\bar{m}$, and at each stage, we need to compute at most pV benefit values where each benefit value computation may take $O(pV)$ time.

Theorem 1. CGA (Algorithm 1) delivers a solution whose total benefit is at least half of the optimal benefit.

Proof. Let L be the total number of iterations of CGA. Note that L is equal to the total number of memory pages in the network. Let Γ_l be the set of variables selected at the end of the l th iteration, and let ζ_l be the variable added to the set Γ_{l-1} in the l th iteration. Let ζ_l be a variable A_{ijk} , signifying that, in the l th iteration, CGA decided to store the j th data item in the k th memory page of the i node. Without loss of generality, we can assume that the optimal solution also stores data items in all memory pages. Now, let λ_l be the variable $A_{ij'k'}$, where j' is the data item stored by the optimal solution in the k th memory page of node i . By the greedy choice of ζ_l , we have

$$\beta(\zeta_l, \Gamma_{l-1}) \geq \beta(\lambda_l, \Gamma_{l-1}), \quad \forall l \leq L. \quad (1)$$

Let O be the optimal benefit³ and C be the benefit of the CGA solution. Note that⁴

$$C = \sum_{l=1}^L \beta(\zeta_l, \Gamma_{l-1}). \quad (2)$$

Now, consider a modified network G' , wherein each node i has a memory capacity of $2m_i$. We construct a cache placement solution for G' by taking a union of data items selected by CGA and data items selected in an optimal solution for each node. More formally, for each variable $\lambda_l = A_{ij'k'}$ as defined above, create a variable $\lambda'_l = A_{ij'k'}$, where $k' = m_i + k$. Obviously, the benefit O' of the set of variables $\{\zeta_1, \zeta_2, \dots, \zeta_L, \lambda'_1, \lambda'_2, \dots, \lambda'_L\}$ in G' is greater than or equal to the optimal benefit O in G . Now, to compute O' , we add the variables in the order of $\zeta_1, \zeta_2, \dots, \zeta_L, \lambda'_1, \lambda'_2, \dots, \lambda'_L$ and add up the benefits of each newly added variable. Let $\Gamma'_l = \{\zeta_1, \zeta_2, \dots, \zeta_L\} \cup \{\lambda_1, \lambda_2, \dots, \lambda_l\}$ and recall that $\Gamma_l = \{\zeta_1, \zeta_2, \dots, \zeta_l\}$. Now, we have

$$\begin{aligned} O \leq O' &= \sum_{l=1}^L \beta(\zeta_l, \Gamma_{l-1}) + \sum_{l=1}^L \beta(\lambda'_l, \Gamma'_{l-1}) \\ &= C + \sum_{l=1}^L \beta(\lambda'_l, \Gamma'_{l-1}) && \text{From (2)} \\ &\leq C + \sum_{l=1}^L \beta(\lambda_l, \Gamma_{l-1}) && \text{Since } \lambda_l = \lambda'_l, \Gamma_{l-1} \subseteq \Gamma'_{l-1} \\ &\leq 2C && \text{From (1) and (2).} \end{aligned}$$

□

The following theorem follows from the above theorem and the definition of benefit and shows that our above result is an improvement of the 20.5-approximation result in [3] when the optimal access cost is at least $(1/40)$ th of the total access cost without the caches.

Theorem 2. If the access cost without the caches is less than 40 times the optimal access cost using optimal cache

3. Note that a solution with optimal benefit also has optimal access cost.

4. Note that $O \neq \sum_{l=1}^L \beta(\lambda_l, \Gamma_{l-1})$. Also, in spite of (2), the benefit value C is actually independent of the order in which ζ_l is selected.

placement, then the total access cost of the CGA solution is less than 20.5 times the optimal access cost.

Proof. Let the total access cost without the caches be W and the optimal access cost (using optimal cache placement) be O . Thus, the optimal benefit is $W - O$. Since the benefit of the CGA solution is at least half of the optimal benefit, the total access time of the CGA solution is at most $W - (W - O)/2$, which is at most $20.5O$. \square

3.1.3 Nonuniform Size Data Items

To handle nonuniform size data items, at each stage, CGA selects a data item to cache at a node such that the (data item, node) pair has the maximum benefit per page at that stage. CGA continues to cache data items at nodes in the manner stated above until each node's memory is exceeded by the last data item cached. Let S be the solution obtained at the end of the above process. Now, CGA picks the better of the following two feasible solutions: (S_1) Each node caches only its last data item. (S_2) Each node caches all the selected data items except the last. For the above solutions to be feasible, we assume that size of the largest data item in the system is less than the memory capacity of any node. Below, we show that the better one of the above two solutions has a benefit of at least 1/4 of the optimal benefit.

Theorem 3. For nonuniform size data items, the above-described modified CGA algorithm delivers a solution whose benefit is at least 1/4 of the optimal benefit. We assume that the size of the largest data item is at most the size of any node's memory capacity.

Proof. First, note that since the solution S is a union of solutions S_1 and S_2 , the benefit of either S_1 or S_2 is at least half of the benefit of S . We prove the theorem by showing below that the benefit of S is at least half of the optimal benefit. The proof below is similar to that of Theorem 1.

As before, we define variables A_{ijk} signifying that (part of) the data item D_j was stored in the k th memory page of node i . However, note that A_{ijk} only corresponds to a *unit* memory page, whereas a data item may occupy more than one memory page. Thus, a cache placement solution may select one or more variables A_{ijk} with the same i and j . Moreover, in this context, the benefit $\beta(A_{ijk}, \Gamma)$ is defined as the benefit *per unit space* of caching D_j at node i , when the certain data items have already been cached at nodes (as determined by the variables in $\Gamma - \{A_{ijk}\}$).

Now, let L' be the total amount of memory used by the solution S . Note that L' may be more than the total number of memory pages in the network. As in Theorem 1, let Γ_l be the set of variables selected at the end of l th iteration, and let ζ_l be the variable added to the set Γ_{l-1} in the l th iteration. Thus, the solution S is the set of variables $\{\zeta_1, \zeta_2, \dots, \zeta_{L'}\}$. Similarly, let the optimal solution be $\{\lambda_1, \lambda_2, \dots, \lambda_{L'}\}$, where λ_l corresponds to the same node and memory page as ζ_l . Without loss of generality, we assume that the optimal solution is "completely" different than S . That is, there is no data item that is cached by S , as well as the

optimal solution at the same node.⁵ As in Theorem 1, by the greedy choice of ζ_l and the above assumption of completely different solutions, we have

$$\beta(\zeta_l, \Gamma_{l-1}) \geq \beta(\lambda_l, \Gamma_{l-1}), \quad \forall l \leq L'. \quad (3)$$

Now, consider a modified network G' , wherein each node i has a memory capacity of $2m_i$. We construct a cache placement solution for G' by taking a union of the solution S and the optimal solution at each node. More formally, for each variable $\lambda_l = A_{ij'k}$ as defined above, create a variable $\lambda'_l = A_{ij'k'}$, where $k' = m'_i + k$, where m'_i is the memory used by the solution S at node i . Obviously, the benefit O' of the set of variables $\{\zeta_1, \zeta_2, \dots, \zeta_{L'}, \lambda'_1, \lambda'_2, \dots, \lambda'_{L'}\}$ in G' is greater than or equal to the optimal benefit. Now, to compute O' , we add the variables in the order of $\zeta_1, \zeta_2, \dots, \zeta_{L'}, \lambda'_1, \lambda'_2, \dots, \lambda'_{L'}$, and add up the benefits of each newly added variable. Let $\Gamma'_l = \{\zeta_1, \zeta_2, \dots, \zeta_{L'}\} \cup \{\lambda_1, \lambda_2, \dots, \lambda_l\}$ and recall that $\Gamma_l = \{\zeta_1, \zeta_2, \dots, \zeta_l\}$. Let C be the benefit of solution S . Now, we have

$$\begin{aligned} O \leq O' &= \sum_{l=1}^{L'} \beta(\zeta_l, \Gamma_{l-1}) + \sum_{l=1}^{L'} \beta(\lambda'_l, \Gamma'_{l-1}) \\ &= C + \sum_{l=1}^{L'} \beta(\lambda'_l, \Gamma'_{l-1}) \\ &\leq C + \sum_{l=1}^{L'} \beta(\lambda_l, \Gamma_{l-1}) \quad \text{Since } \lambda_l = \lambda'_l, \Gamma_{l-1} \subseteq \Gamma'_{l-1} \\ &\leq 2C \quad \text{From (1) and (2).} \end{aligned}$$

\square

3.2 Distributed Greedy Algorithm (DGA)

In this section, we describe a localized distributed implementation of CGA. We refer to the designed distributed implementation as DGA. The advantage of DGA is that it adapts to dynamic traffic conditions and can be easily implemented in an operational (possibly, mobile) network with low communication overheads. Although we cannot prove any bound on the quality of the solution produced by DGA, we show through extensive simulations that the performance (in terms of the quality of the solution delivered) of the DGA is very close to that of the CGA. The DGA is composed of two important components—nearest cache tables and localized caching policy—as described below.

3.2.1 Nearest Cache Tables

For each network node, we maintain the *nearest* node (including itself) that has a copy of the data item D_j for each data item D_j in the network. More specifically, each node i in the network maintains a *nearest cache table*, where an entry in the nearest cache table is of the form (D_j, N_j) , where N_j is the closest node that has a copy of D_j . Note that if i is the server of D_j or has cached D_j , then N_j is i . In addition, if a node i has cached D_j , then it also maintains an entry (D_j, N_j^2) , where N_j^2 is the *second-nearest cache*, that is, the closest node (other than i itself)

⁵ Else, we could remove the common data items from both solutions and prove the claim below about the remaining solutions.

that has a copy of D_j . The second-nearest cache information is helpful when node i decides to remove the cached item D_j . Note that if i is the server of D_j , then N_j is i . The above information is in addition to any information (such as routing tables) maintained by the underlying routing protocol. The nearest cache tables at network nodes in the network are maintained as follows in response to cache placement changes:

Addition of a cache. When a node i caches a data item D_j , N_j^2 (the second-nearest cache) is set to its current N_j (nearest cache node), and N_j is updated to i itself. In addition, the node i broadcasts an AddCache message to all of its neighbors. The AddCache message contains the tuple (i, D_j) signifying the ID of the originating node and the ID of the newly cached data item. Consider a node l that receives the AddCache message (i, D_j) . Let (D_j, N_j) be the nearest cache table entry at node l signifying that N_j is the cache node currently closest to l that has the data item D_j . If $d_{li} < d_{lN_j}$,⁶ then the nearest cache table entry (D_j, N_j) is updated to (D_j, i) , and the AddCache message is forwarded by l to all of its neighbors. Otherwise, the node l sends the AddCache message to the single node N_j (which could be itself) so that N_j can possibly update information about its second-nearest cache. The above process maintains the correctness of the nearest cache entries and second-nearest cache entries in a static network with bounded communication delays, because of the following observations:

- O1. Consider a node k whose nearest cache table entry *needs* to change (to i) in response to the addition of a cache at node i . Then, every intermediate node on the shortest path connecting k to i also needs to change its nearest cache table entry (and hence, forward the AddCache message).
- O2. Consider a cache node k such that its second-nearest cache node *should* be changed to i in response to the addition of a cache at node i . Then, there exists two distinct *neighboring* nodes i_1 and i_2 (not necessarily different from k or i) on the shortest path from k to i such that the nearest cache node of i_1 is k , and the nearest cache node of i_2 is i .

The first observation ensures the correctness of the nearest cache entries since any node k that needs to receive the AddCache message receives it through the intermediate nodes on the shortest path connecting k to i . The second observation ensures the correctness of the second-nearest cache entries, since for any cache node k whose second-nearest cache entry must change to the newly added cache i , there exists a node i_1 that sends the AddCache message (received from the forwarding neighboring node i_2) to k (i_1 's nearest cache node). We now prove the above two observations.

We prove the first observation O1 by contradiction. Consider a node k whose nearest cache table entry *needs* to change (to i) in response to addition of a cache at node i . Assume that there is an intermediate node j on the shortest path \mathcal{P} connecting k to i such that j does not need to change its nearest cache entry. Thus, there is another cache node l

that is nearer to j than i . Then, the cache node l is also closer to k than i , and thus, k does not need to change its nearest cache entry due to the addition of cache at node i —which is a contradiction. The second observation O2 is true because the nearest cache node of each intermediate node on the shortest path connecting the *cache nodes* k and i (such that i is the second-nearest cache node of k) is either k or i .

Deletion of a cache. To efficiently maintain the nearest cache tables in response to deletion of a cache, we maintain a *cache list* C_j for each data item D_j at its server S_j . The cache list C_j contains the set of nodes (including S_j) that have cached D_j . To keep the cache list C_j up to date, the server S_j is informed whenever the data time D_j is cached at or removed from a node. Note that the cache list C_j is almost essential for the server S_j to efficiently update D_j at the cache nodes. Now, when a node i removes a data item D_j from its local cache, it updates its nearest cache node (N_j) to its second-nearest cache (N_j^2) and deletes the second-nearest cache entry. In addition, the node i requests C_j from the server S_j and, then, broadcasts a DeleteCache message with the information (i, D_j, C_j) to all of its neighbors. Consider a node l that receives the DeleteCache message, and let (D_j, N_j) be its nearest cache table entry. If $N_j = i$, then the node l updates its nearest cache entry using C_j and forwards the DeleteCache message to all its neighbors. Otherwise, the node l sends the DeleteCache message to the node N_j . The above process ensures the correctness of the nearest cache and the second-nearest cache entries due to the above two observations (O1 and O2). If the maintenance of a complete cache list at the server is not feasible, then we can either broadcast the DeleteCache message with $\{S_j\}$ as the cache list or not use any DeleteCache messages at all. In the latter case, when a data request for the deleted cache is received, the data request can be redirected to the server.

Integrated cache-routing tables. Nearest caching tables can be used in conjunction with any underlying routing protocol to reach the nearest cache node, as long as the distances to other nodes are maintained by the routing protocol (or available otherwise). If the underlying routing protocol maintains routing tables [24], then the nearest cache tables can be integrated with the routing tables as follows: For a data item D_j , let H_j be the next node on the shortest path to N_j , the closest node storing D_j . Now, if we maintain a *cache-routing* table having entries of the form (D_j, H_j, δ_j) , where δ_j is the distance to N_j , then there is no need for routing tables. However, note that maintaining the cache-routing tables instead of the nearest cache tables *and* the routing tables does not offer any clear advantage in terms of the number of message transmissions.

3.2.2 Mobile Networks

To handle mobility of nodes, we could maintain the integrated cache-routing tables in a similar vein as routing tables [24] are maintained in mobile ad hoc networks. Alternatively, we could have the servers periodically broadcast the latest cache lists. In our simulations, we adopted the latter strategy, since it precludes the need to broadcast AddCache and DeleteCache messages to some extent.

6. The distance values are assumed to be available from the underlying routing protocol.

3.2.3 Localized Caching Policy

The caching policy of DGA is given as follows: Each node computes the benefit of data items based on its “local traffic” observed for a sufficiently long time. The *local traffic* of a node i includes its own local data requests, nonlocal data requests to data items cached at i , and the traffic that the node i is forwarding to other nodes in the network.

Local benefit. We refer to the benefit computed based on the node’s local traffic as the *local benefit*. For each data item D_j not cached at node i , the node i calculates the local benefit gained by caching the item D_j , whereas for each data item D_j cached at node i , the node i computes the local benefit lost by removing the item. In particular, the local benefit B_{ij} of caching (or removing) D_j at node i is the reduction (or increase) in access cost given by

$$B_{ij} = t_{ij}\delta_j,$$

where t_{ij} is the access frequency observed by node i for the item D_j in its local traffic, and δ_j is the distance from i to N_j or N_j^2 —the nearest node other than i that has the copy of the data item D_j . Using the nearest cache tables, each node can compute the local benefits of data items in a localized manner using only local information. Since the traffic changes dynamically (due to new cache placements), each node needs to continually recompute local benefits based on most recently observed local traffic.

Caching policy. A node decides to cache the most beneficial (in terms of local benefit per unit size of data item) data items that can fit in its local memory. When the local cache memory of a node is full, the following cache replacement policy is used. Let $|D|$ denote the size of a data item (or a set of data items) D . If the local benefit of a newly available data item D_j is higher than the total local benefit of some set D of cached data items, where $|D| > |D_j|$, then the set D is replaced by D_j . Since, adding or replacing a cache entails communication overhead (due to AddCache or DeleteCache messages), we employ the concept of *benefit threshold*. In particular, a data item is newly cached only if its local benefit is higher than the benefit threshold, and a data item replaces a set of cached data items only if the difference in their local benefits is greater than the benefit threshold.

3.2.4 Distributed Greedy Algorithm (DGA)

The above components of the nearest cache table and the cache replacement policy are combined to yield our DGA for cache placement problem. In addition, the server uses the cache list to periodically update the caches in response to changes to the data at the server. The departure of DGA from CGA is primarily in its inability to gather information about all traffic (access frequencies). In addition, the inaccuracies and staleness of the nearest cache table entries (due to message losses or arbitrary communication delays) may result in approximate local benefit values. Finally, in DGA, the placement of caches happens simultaneously at all nodes in a distributed manner, which is in contrast to the sequential manner in which the caches are selected by the CGA. However, DGA is able to cope with dynamically changing access frequencies and cache placements. As noted before, any changes in cache placements trigger updates in the nearest cache table, which in turn affect the

local benefit values. Below is a summarized description of the DGA.

Algorithm 2 DGA.

Setting

A network graph $G(V, E)$ with p data items. Each node i has a memory capacity of m_i pages. Let Θ be the benefit threshold.

Program of Node i

BEGIN

When a data item D_j passes by

if local memory has available space and $(B_{ij} > \Theta)$

then cache D_j

else if there is a set D of cached data items such

that (local benefit of $D < B_{ij} - \Theta$) and

$(|D| \geq |D_j|)$, then replace D with D_j

When a data item D_j is added to a local cache

Notify the server of D_j

Broadcast an AddCache message with (i, D_j)

When a data item D_j is deleted from a local cache

Get the cache list C_j from the server of D_j

Broadcast a DeleteCache message with

(i, D_j, C_j)

On receiving an AddCache message (i', D_j)

if i' is nearer than the current nearest cache for D_j

then update the nearest cache table entry and broadcast the AddCache message to neighbors

else send the message to the nearest cache of i

On receiving a DeleteCache message (i', D_j, C_j)

if i' is the current nearest cache for D_j

then update the nearest case of D_j using C_j , and

broadcast the DeleteCache message

else send the message to the nearest cache of i

For **mobile networks**, instead of AddCache and

DeleteCache messages, for each data item, its server periodically broadcasts (to the entire network) the latest cache list.

END.

3.2.5 Performance Analysis

Note that the performance guarantee of CGA (that is, proof of Theorem 1) holds even if the CGA were to consider the memory pages in some arbitrary order and select the most beneficial caches for each one of them. Now, based on the above observation, if we assume that the local benefit is reflective of the accurate benefit (that is, if the local traffic seen by a node i is the only traffic that is affected by caching a data item at node i), then DGA also yields a solution whose benefit is 1/4 of the optimal benefit. Our simulation results in Section 4.1 show that DGA and CGA indeed perform very closely.

3.2.6 Data Expiry and Cache Updates

We incorporate the concepts of data expiry and cache updates in our overall framework as follows: For data expiry, we use the concept of *Time to Live (TTL)* [33], which is the time until the given copy of the data item is considered valid/fresh. The data item or its copy is considered *expired* at the end of the TTL time value. We

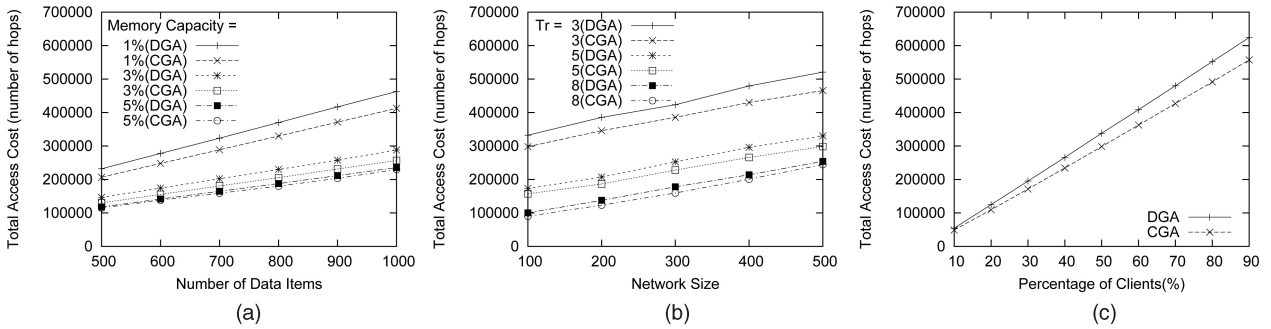


Fig. 2. Performance comparison of CGA and DGA. (a) Varying number of data items and memory capacity. (b) Varying number of nodes and transmission radius (T_r). (c) Varying client percentage. Unless being varied—the number of nodes is 500, transmission radius is 5, number of data items is 1,000, number of clients (for each data item) is 250, and each node can store 20 data items in its memory.

consider two data expiry models, namely, *TTL per request* and *TTL per item*. In the *TTL-per-request* data expiry model [33], the server responds to any data item request with the requested data item and an appropriately generated TTL value. Thus, *each* copy of the data item in the network is associated with an independent TTL value. In the *TTL-per-item* data expiry model, the server associates a TTL value with each *data item* (rather than each request), and all requests for the same data item are associated with the same TTL (until the data item expires). Thus, in the *TTL-per-item* data expiry model, all the fresh copies of a data item in the network are associated with the same TTL value. On the expiration of the data item, the server generates a new TTL value for the data item.

For updating the cached data items, we consider two mechanisms. For the case of the *TTL-per-request* data expiry model, we use the *cache deletion* update model, where each cache node independently deletes its copy of the expired data item. Such deletions are handled in a similar way as describe before, that is, by broadcasting a DeleteCache request. In the case of *TTL-per-item* data expiry model, all the copies of a particular data item expire simultaneously. Thus, we use the *server multicast* cache update model, wherein the server multicasts the fresh copy of the data item to all the cache nodes on expiration of the data item (at the server). If the cache list is not maintained at the server, then the above update is implemented using a network-wide broadcast.

4 PERFORMANCE EVALUATION

We demonstrate through simulations the performance of our designed cache placement algorithms over randomly generated network topologies. We first compare the relative quality of the solutions returned by CGA and its distributed algorithm DGA. Then, we turn our attention to application-level performance in complex network settings and evaluate our designed DGA with respect to a naive distributed algorithm and the HybridCache algorithm [33] using the ns2 simulator [10].

4.1 CGA versus DGA

In this section, we evaluate the relative performance of CGA and DGA by comparing the benefits of the solutions delivered.

For the purposes of implementing a centralized strategy, we use our own simulator for implementation and comparison of our designed algorithms. In our simulator, DGA is implemented as a dynamically evolving process wherein, initially, all the memory pages are free, and the nearest cache table entries point to the corresponding servers. This initialization of the nearest cache table entries results in traffic being directed to servers, which triggers caching of data items at nodes, which in turn causes changes in the nearest cache tables and further changes in cache placements. The process continues until convergence. To provide a semblance of an asynchronous distributed protocol, our simulation model updates routing and the nearest cache table entries in an arbitrary order across nodes.

Simulation parameters. In our cache placement problem, the relevant parameters are

1. number of nodes in the network,
2. transmission radius T_r (two nodes can directly transmit with each other if and only if they are within T_r distance from each other),
3. number of data items,
4. number of clients accessing each data item, and
5. memory capacity on each node.

The first two parameters are related to network topology, the next two parameters are application dependent, and the last parameter is the problem constraint (property of the nodes). Here, we assume each data item to be of unit size (one memory page). Below, we present a set of plots wherein we vary some of the above parameters while keeping the others constant.

4.1.1 Varying Number of Data Items and Memory Capacity

Fig. 2a plots the access costs for CGA and DGA against the number of data items in the network for different local memory capacities. Here, the network size is 500 nodes in a 30×30 area.⁷ We use a transmission radius (T_r) of five units. The memory capacity in each node is expressed as the percentage of the number of data items in the network. We vary the number of data items from 500 to 1,000 and the memory capacity of each node from 1 percent to 5 percent of

7. Since the complexity of CGA is a high-order polynomial, the runtime is quite slow. Thus, we have not been able to evaluate the performance on very large networks.

the number of data items. The number of clients accessing each data item is fixed at 50 percent of the number of nodes in the network.

We observe that the access cost increases with the number of data items as expected. Also, as expected, we see that CGA performs slightly better since it exploits global information, but DGA performs quite close to CGA. The performance difference between the algorithms decreases with increasing memory capacity, since, with increasing memory capacity, both the algorithms must converge to the same solution (access cost zero) as all client nodes will eventually be able to cache all the data items they wish to access. Although this degenerate situation is not reached, the trend is indeed observed.

4.1.2 Varying Network Size and Transmission Radius

In the next plot (Fig. 2b), we fix the number of data items in the network to 1,000 and the memory capacity of each node to 2 percent of data items. As before, 50 percent of the network nodes act as clients for each of the data item. In this plot, we vary the network size from 100 nodes to 500 nodes and transmission radius (T_r) from 3 to 8. Essentially, Fig. 2b shows the access cost as a function of network size and transmission radius for the two algorithms. Once again, as expected, CGA slightly outperforms DGA, but DGA performs very close to that of CGA.

4.1.3 Varying Client Percentage

We also investigated the effect of the number of clients on the access cost. See Fig. 2c. We note a similar behavior. The performances are more similar as the independent parameter is varied toward the degenerate case. Here, the degenerate case represents a single client, where both algorithms must perform similarly.

4.2 DGA versus HybridCache

In this section, we compare DGA with the HybridCache approach proposed in [33] by simulating both approaches in *ns2* [10] (version 2.27). The *ns2* simulator contains models for common ad hoc network routing protocols, IEEE Standard 802.11 MAC layer protocol, and two-ray ground reflection propagation models [6]. Due to the lack of scalability of the *ns2* simulator, we use a smaller network size (compared to Section 4.1) and other appropriate parameter values. The DSDV routing protocol [24] is used to provide routing services. For comparison, we also implemented a *Naive* approach, wherein each node caches any passing-by data item if there is a free memory space and uses the LRU (least recently used) policy for replacement of caches. We start with presenting the simulation setup and then present the simulation results in Section 4.2.1. In all plots, each data point represents an average of 5 to 10 runs. In some plots, we show error bars indicating the 95 percent confidence interval; for sake to clarity, we show confidence intervals in only those graphs that are relevant to our claims.

4.2.1 Simulation Setup

In this section, we briefly discuss the network set up, client query model, data access pattern model, and performance metrics used for our simulations.

Network setup. We simulated our algorithms on a network of randomly placed 100 nodes in an area⁸ of $2,000 \times 500$ m². Note that the nominal radio range for two directly communicating nodes in the *ns2* simulator is about 250 meters. We assume that the wireless bandwidth is 2 megabits per second. Other network parameters (such as propagation delay, etc.) are the same as the *ns2* default. When a node relays a message, we introduce a random delay between 0.001-0.01 second to reduce collisions. There is no other routing delay except queueing delays that occur naturally at the network interface. In our simulations, we assume 1,000 data items of varying sizes, two randomly placed servers S_0 and S_1 , where S_0 stores the data items with even IDs, and S_1 stores the data items with odd IDs. We choose the size of a data item randomly between 100 and 1,500 bytes.⁹

Client query model. In our simulations, each network node is a client node. Each client node in the network sends out a single stream of read-only queries. Each query is essentially a request for a data item. In our DGA scheme, the query is forwarded to the nearest cache (based on the nearest cache table entry). In the Naive scheme, the query is forwarded to the server unless the data item is available locally; if the query encounters a node with the requested data item cached, then the query is answered by the encountered node itself. The time interval between two consecutive queries is known as the *query generate time* and follows exponential distribution with mean value T_{query} , which we vary from 3 to 40 seconds. We do not consider values of T_{query} less than 3 seconds, since they result in a query success ratio of much less than 80 percent for Naive and HybridCache approaches. Here, the *query success ratio* is defined as the percentage of the queries that receive the requested data item within the *query success timeout* period. In our simulations, we use a query success timeout of 40 seconds.

The above client query model is similar to the model used in previous studies [7], [33]. However, query generation process differs slightly from the one used in [33] in how the queries are generated. In [33], if the query response is not received within the query success timeout period, then the same query is sent repeatedly until it succeeds, whereas on the success of a query, a new query is generated (as in our model) after some random interval.¹⁰ Our querying model is better suited (due to the exact periodicity of querying) for comparative performance evaluation of various caching strategies, whereas the querying model in [33] depicts a more realistic model of a typical application (due to repeated querying until success).

Data access models. For our simulations, we use the following patterns for modeling data access frequencies at nodes:

1. *Spatial pattern.* In this pattern of data access, the data access frequencies at a node depends on its geographic location in the network area such that nodes

8. This is similar to the configuration in [32] and [33], wherein they distribute 50 nodes in an area of $1,500 \times 320$ m².

9. The maximum data size used in [33] is 10 Kbytes, which is not a practical choice due to lack of MAC layer fragmentation/reassembly mechanism in the 2.27 version of *ns2* that we used.

10. In the original simulation code of HybridCache [33], the time interval between two queries is actually 4 seconds plus the query generate time (which follows an exponential distribution with mean value T_{query}).




Fig. 3. Varying mean query generate time on spatial data access pattern. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

that are closely located have similar data access frequencies.¹¹ More specifically, we start with laying the given 1,000 data items uniformly over the network area in a gridlike manner resulting in a virtual coordinate for each data item. Then, each network node accesses the 1,000 data items in a *Zipf-like* distribution [5], [34], with the access frequencies of the data items ordered by the distance of the data item's virtual coordinates from the network node. More specifically, the probability of accessing (which can be mapped to access frequency) the j th ($1 \leq j \leq 1,000$) closest data item is represented by $P_j = \frac{1}{j^\theta \sum_{h=1}^{1,000} 1/h^\theta}$,

where $0 \leq \theta \leq 1$. Here, we have assumed the number of data items to be 1,000. When $\theta = 1$, the above distribution follows the strict Zipf distribution, whereas for $\theta = 0$, it follows the uniform distribution. As in [33], we choose θ to be 0.8 based on real Web trace studies [5].

2. *Random pattern.* In this pattern of data access, each node uniformly accesses a predefined set of 200 data items chosen randomly from the given 1,000 data items.

Performance metrics. We measure three performance metrics for comparing various caching strategies, namely, average query delay, total number of messages, and query success ratio. Query delay is defined as the time elapsed between query request and query response, and the average query delay is the average of query delays over all queries. Total number of messages includes *all* message transmissions between neighboring nodes, including messages due to queries, maintenance of the nearest cache tables and cache lists, and periodic broadcast of cache lists in mobile networks. Messages to implement routing protocol are not counted, as they are the same in all three approaches compared. Query success ratio has been defined before. Each data point in our simulation results is an average of over five different random network topologies, and to achieve stability in performance metrics, each of our experiments is run for a sufficiently long time (20,000 seconds for our experiments).

DGA parameter values. We now present a brief discussion on the choice of values of the benefit threshold

and local traffic window size for DGA. For static networks, we compute local benefits based on the most recent 1,000 queries. Since, the data access frequencies remains static in our experiment setting, computing local benefits based on as large a number of queries as possible is a good idea. However, we observed that the most recent 1,000 queries are sufficient to derive a complete knowledge of local traffic. For mobile networks with spatial data access pattern, the access frequencies at a client node change with the node's location. Thus, we compute local benefits using only 50 recent queries.

Also, we chose a benefit threshold value of 0.008 when the cache size is default at 75 Kbytes (capable of storing 100 average-sized data items), based on the typical benefit value of the 100th most beneficial data item at a node. We use a similar methodology for choosing benefit threshold values for other values of cache sizes. In general, the chosen benefit threshold value should be higher than the communication overhead that incurred (in terms of maintenance of the nearest cache tables and the cache list) due to caching of a data item.

4.2.2 Simulation Results

We now present simulation results comparing the three caching strategies, namely, Naive Approach, HybridCache approach in [33], and our DGA, under the random and spatial data access patterns (as defined above) and study the effect of various parameter values on the performance metrics.

Varying mean query generate time. In Fig. 3, we vary the mean query generate time T_{query} in the spatial data access pattern while keeping the cache size as a constant and all network nodes as client nodes. We choose the cache size to be big enough to fit about 100 average-sized data items (that is, 75 Kbytes). We observe that our DGA outperforms the other two approaches in terms of all three performance metrics of query average delay, query success ratio, and total number of messages. In comparison with HybridCache strategy, our DGA has an average query delay of less than half for all parameter values (corroborated by confidence intervals of 95 percent), always has better query success ratio and lower message overhead. For the mean query generate time of 3 seconds, the average query delay in all approaches is high, but our DGA outperforms HybridCache by more than a factor of 10. Also, for very low mean query generate times, our DGA has a significantly better query

11. This data access pattern is similar to the spatial gridlike access pattern used in [33].

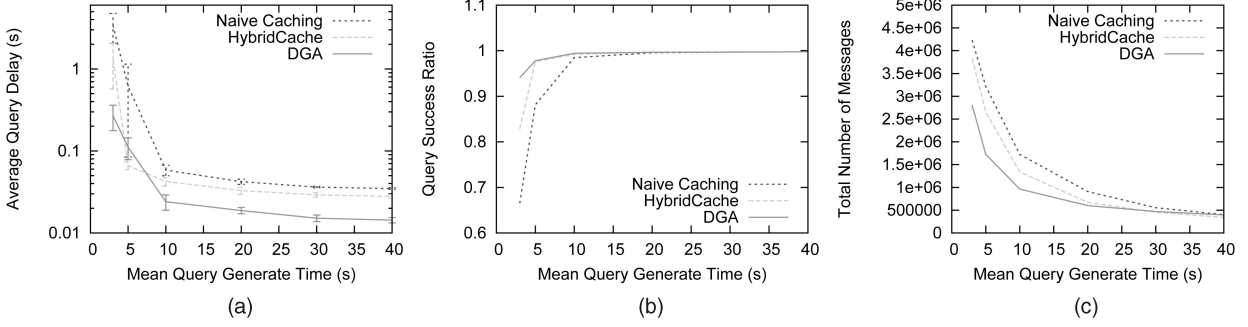


Fig. 4. Varying mean of query generate time on random data access pattern. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

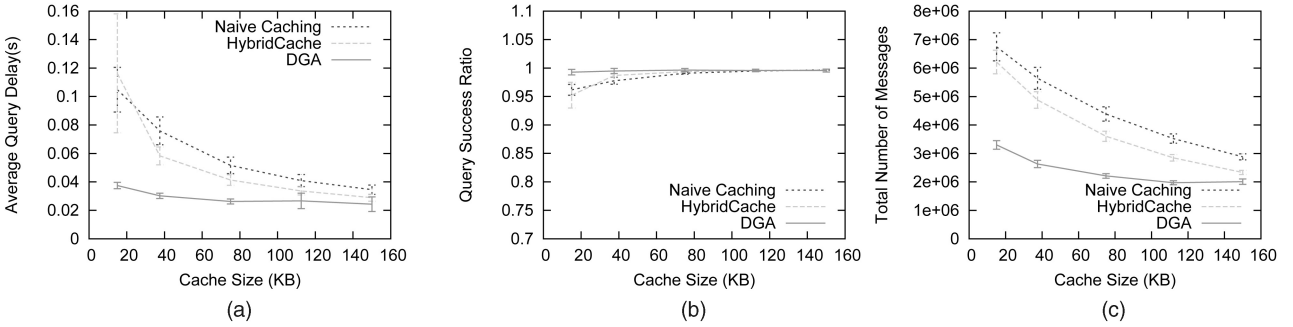


Fig. 5. Varying cache size on spatial data access pattern. Here, $T_{query} = 10$ seconds. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

success ratio. Fig. 4 depicts similar observations for the random access data patterns, except that for a mean query generate time of 5 seconds, we have a slightly worse average query delay than that of HybridCache (but a significantly better query success ratio).

Varying cache memory size. In Fig. 5, we vary the local cache size of each node in the spatial data access pattern while keeping the mean query generate time T_{query} constant at 10 seconds. We vary the local cache size from 15 Kbytes (capable of storing 20 data items of average size) to 50 Kbytes. We observe in Fig. 5 that our DGA outperforms the HybridCache approach consistently for all cache sizes and in terms of all three performance metrics. The difference in the average query delay is much more significant for lower cache size—which suggests that our DGA is very judicious in the choice of data items to cache. Note that HybridCache performs

even worse than the Naive Approach when each node’s memory is 15 Kbytes.

Mobile networks. Until now, we have restricted our discussion and simulations to ad hoc networks with static nodes. Now, we present the performance comparison of various caching strategies for mobile ad hoc networks, wherein the mobile nodes move based on the “random waypoint” movement model [6]. In the random waypoint movement model, initially, nodes are placed randomly in the area. Each node selects a random destination and moves toward the destination with a speed selected randomly from $(0 \text{ m/s}, v_{max} \text{ m/s})$. After the node reaches its destination, it pauses for a period of time (chosen to be 300 seconds in our simulations, as in [33]) and repeats the movement pattern. In our simulations, the server broadcasting the cache lists every 100 seconds; this time interval is sufficient for notifying the nodes in a timely manner without incurring too much overhead.

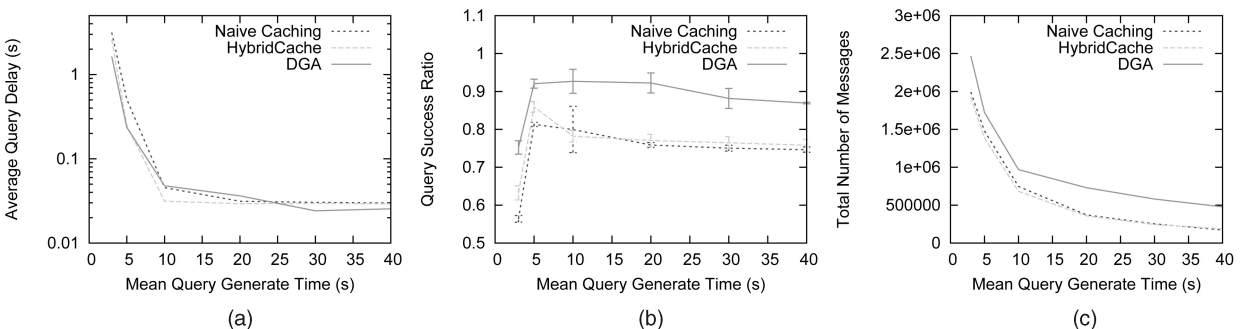


Fig. 6. Varying mean query generate time in spatial data access pattern with $v_{max} = 10 \text{ m/s}$. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

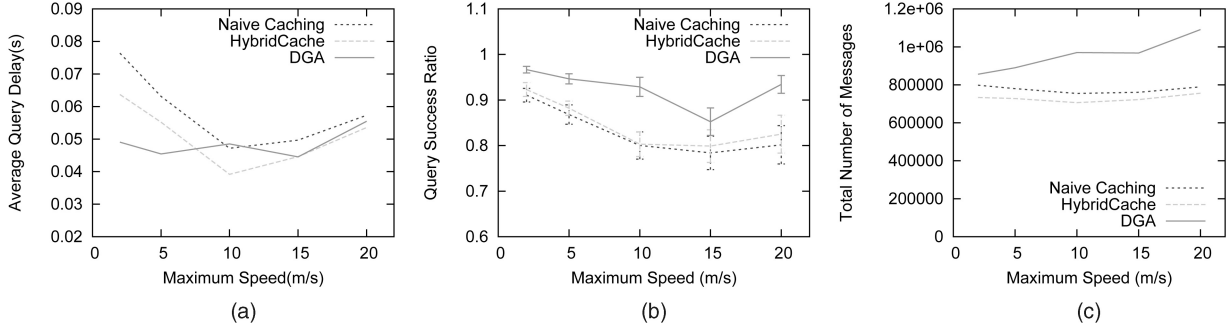


Fig. 7. Varying v_{max} in spatial data access pattern. Here, $T_{query} = 10$ seconds. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

In Fig. 6, we compare various cache placement algorithms under the spatial data access pattern for varying mean query generate time while keeping other parameters constant ($v_{max} = 10$ m/s and local cache size = 75 Kbytes). We observe that all our schemes perform similarly in terms of query delay, but DGA outperforms the other schemes by a significant margin in terms of query success ratio (again, corroborated by confidence intervals). Note that a significantly better query success ratio is much more desirable than a slightly better average query delay. In Fig. 7, we compare various schemes under the spatial data access pattern for varying v_{max} value while keeping other parameters constant ($T_{query} = 10$ seconds and local cache size = 75 Kbytes). In terms of query delay, DGA outperforms other schemes for low mobilities but has a slightly worse query delay for higher mobilities. However, more importantly, DGA has a *significantly* better query success ratio than all schemes for all mobilities. As noted before, a much better query success ratio is more desirable than slightly better query delay. We have the following explanation for the unusual (nonmonotonic) pattern of the graphs in Fig. 7, which is the only figure in this article where we have varied mobilities. First, Naive and Hybrid schemes display similar patterns—the query success ratio initially decreases with increase in mobility (as expected) and, then, stabilizes. The initial decrease in query delay is largely due to the effect of decrease in query success ratio (due to loss of longer delay queries). The later increase in query delay with an increase in mobility is as expected, when the query success ratio remains largely unchanged. In contrast, we notice that the

DGA scheme has a relatively unchanged query delay and query success ratio, suggesting that higher mobility does not deteriorate much the performance of DGA due to the presence of the nearest cache table structure.

Varying client percentage. In all previous experiments in this section, we have assumed that each network node is a client node. In Fig. 8, we vary the percentage of client nodes in the static network for the spatial data access pattern while keeping $T_{query} = 10$ seconds and cache size as 75 Kbytes. We can see that DGA outperforms HybridCache for all client percentage values. The performance difference is seen to be very less at a very low percentage of client nodes because of minimal traffic. Fig. 9 shows similar trend and results for mobile networks with mobility ($v_{max} = 10$ m/s).

Incorporating data expiry and cache updates. In all of our previous experiments, we have not considered data expiration or cache updates. We now incorporate data expiry and cache updates into our simulations. We run our simulations for a total runtime of 200,000 seconds and generate TTL values as *current time* plus a random number in [10,000, 20,000]. We use both data expiry models, namely, TTL per request and TTL per item. As mentioned before, for the TTL-per-request data expiry model, we use the cache deletion update mechanism, whereas for the TTL-per-time model, we use the server multicast update mechanism. For all the three caching algorithms (Naive, Hybrid, and DGA), a data item request destined to a node with expired data item is redirected to the server, and the TTL value of a cached expired data item is updated using the TTL values of a passing by fresh copy of the data item.

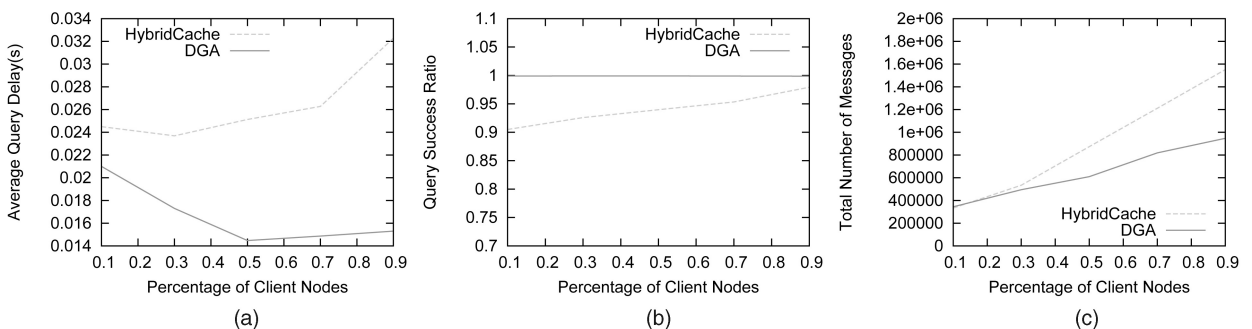


Fig. 8. Varying client percentage on spatial data access pattern in static networks. Here, $T_{query} = 10$ seconds. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

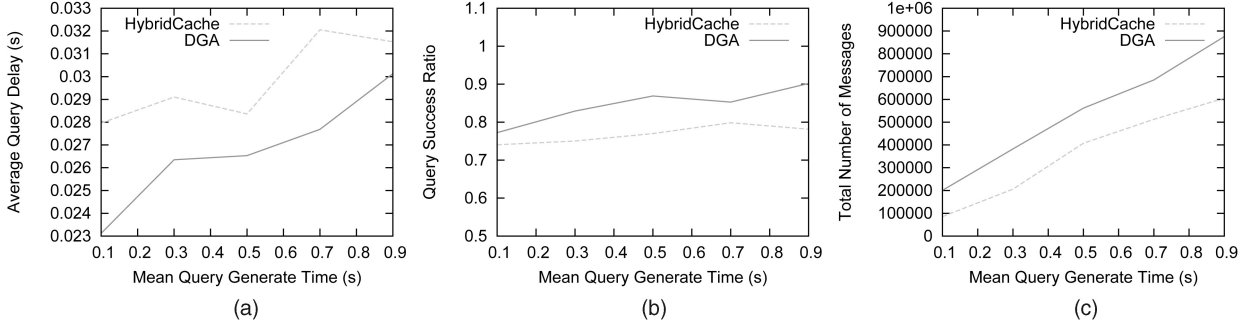


Fig. 9. Varying client percentage on spatial data access pattern for mobile networks with $v_{max} = 10$ m/s. Here, $T_{query} = 10$ seconds. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.




Fig. 10. Varying mean query generate time on spatial data access pattern with cache update in static networks. Here, the data expiry model is TTL-per-item, and the cache update model is server multicast. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

Figs. 10 and 11 show the comparison of the three caching techniques for TTL-per-item and TTL-per-request data expiry models, respectively. In Fig. 10, we see that our DGA technique outperforms HybridCache and Naive Caching in all three performance metrics; the relative performance is similar to that in Fig. 3. However, for the case of TTL-per-request data expiry model (Fig. 11), our DGA has a lower query success ratio (95 percent) due to an increase in the number of DeleteCache messages; our DGA still outperforms the other two techniques in terms of average query delay by a significant margin. Figs. 12 and 13 show the comparison of the DGA and HybridCache in mobile networks with $v_{max} = 10$ m/s for TTL-per-item and TTL-per-request data expiry models, respectively. The total runtime for these experiments is 100,000 seconds, at which the average query delay and the query success ratio values had stabilized. In this very general setting of mobility, data

expiration, and cache updates, we continue to see that our DGA technique outperforms HybridCache in terms of average query delay and query success ratio.

Comparison with random caching/nearest cache table. To demonstrate that the better performance of DGA is not only due to the presence of the nearest cache table but also due to the way the caches are placed, we compare our DGA scheme with a *Random Caching* scheme aided with the nearest cache table. In the Random Caching scheme, we cache/place data items randomly in each node's cache memory and appropriately initialize the nearest cache table. The placement of caches and the initialization of the nearest cache tables is done in a centralized way without any communication overhead, which only favors the Random Caching scheme. In Fig. 14, we compare the DGA, HybridCache, and Random Caching schemes in highly mobile networks (that is, with $v_{max} = 10$ m/s). We observe that due to high mobility, all

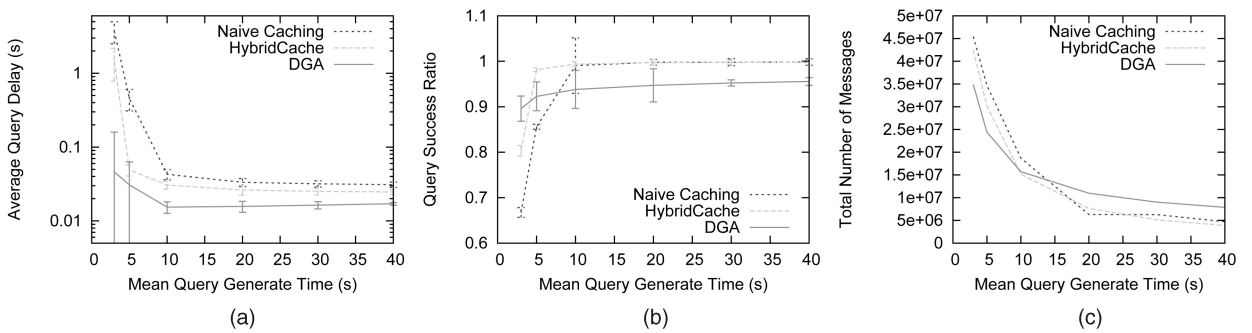


Fig. 11. Varying mean query generate time on spatial data access pattern with cache update in static networks. Here, the data expiry model is TTL-per-request, and the cache update model is cache deletion. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

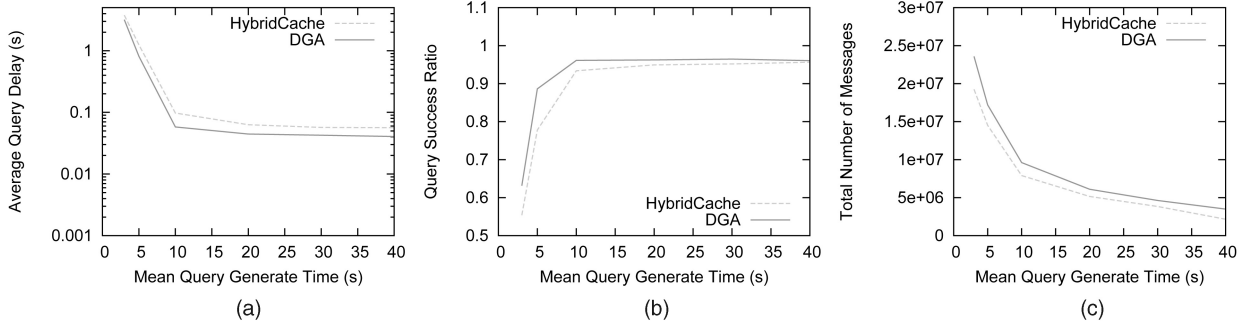


Fig. 12. Varying mean query generate time on spatial data access pattern with cache update in mobile networks with $v_{max} = 10$ m/s. Here, the data expiry model is TTL-per-item, and the cache update model is server multicast. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.




Fig. 13. Varying mean query generate time on spatial data access pattern with cache update in mobile networks with $v_{max} = 10$ m/s. Here, the data expiry model is TTL-per-request, and the cache update model is cache deletion. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

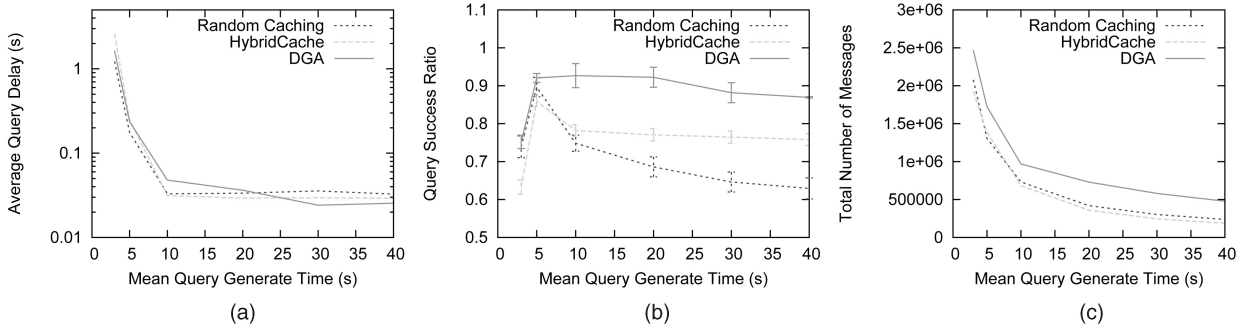


Fig. 14. Varying mean query generate time on spatial data access pattern by comparing Random Caching, HybridCache and DGA, with $v_{max} = 10$ m/s. (a) Average query delay. (b) Query success ratio. (c) Total number of messages.

three different schemes have similar average query delays. However, DGA has significantly better query success ratio than the other schemes. These results demonstrate that the superior performance of our DGA scheme is not just due to the nearest cache table structure.

Summary of simulation results. Our simulation results can be summarized as follows: Both the HybridCache and DGA approaches outperform the Naive approach in terms of all three performance metrics, namely, average query delay, query success ratio, and total number of messages. Our designed DGA almost always outperforms the Hybrid approach in terms of *all* performance metrics for a wide range of parameters of mean query generate time, local cache size, and mobility speed. In particular, for frequent queries or smaller cache size, the DGA approach has a significantly better average query delay and query success ratio. For very high mobility speeds, sometimes, the DGA approach has a

slight worse average query delay than Hybrid but with significantly better query success ratio, which is certainly the more desirable performance metric. We show that the success of DGA comes not only from maintenance of the nearest cache tables but also from the near-optimal placement of caches. The optimized placement of caches not only reduces query delay but also message transmissions, which in turn leads to less congestion and, hence, fewer lost messages due to collisions or buffer overflows at the network interfaces. This, in turn, provides a better success ratio. This “snowballing” effect is very apparent in challenging cases such as frequent queries and small cache sizes.

5 CONCLUSIONS

We have developed a paradigm of data caching techniques to support effective data access in ad hoc networks. In

particular, we have considered the memory capacity constraint of the network nodes and developed efficient algorithms to determine near-optimal cache placements to maximize reduction in overall access cost. Reduction in access cost leads to communication cost savings and, hence, better bandwidth usage and energy savings. Our later simulation experience with *ns2* also shows that better bandwidth usage also in turn leads to less message losses and, thus, better query success ratio.

The novel contribution in our work is the development of a 4-approximation centralized algorithm, which is naturally amenable to a localized distributed implementation. The distributed implementation uses only local knowledge of traffic. However, our simulations over a wide range of network and application parameters show that the performance of the two algorithms is quite close. We note that ours is the first work that presents a distributed implementation based on an approximation algorithm for the problem of cache placement of multiple data items under memory constraint.


We further compare our distributed algorithm with a competitive algorithm (HybridCache) presented in the literature that has a similar goal. This comparison uses the *ns2* simulator with a complete wireless networking protocol stack including dynamic routing. We consider a broad range of application parameters and both stationary and mobile networks. These evaluations show that our algorithm significantly outperforms HybridCache, particularly in more challenging scenarios, such as higher query frequency and smaller memory.

ACKNOWLEDGMENTS

The authors would like to thank the associate editor and anonymous referees for their time and valuable comments. They are very grateful to Dr. Liangzhong Yin and Dr. Guohong Cao for being patient and for offering them a helpful discussion and for the simulation code of HybridCache, with which their work is compared.

REFERENCES

- [1] A. Aazami, S. Ghandeharizadeh, and T. Helmi, "Near Optimal Number of Replicas for Continuous Media in Ad-Hoc Networks of Wireless Devices," *Proc. Int'l Workshop Multimedia Information Systems*, 2004.
- [2] B. Awerbuch, Y. Bartal, and A. Fiat, "Heat and Dump: Competitive Distributed Paging," *Proc. IEEE Symp. Foundations of Computer Science (FOCS '93)*, 1993.
- [3] I. Baev and R. Rajaraman, "Approximation Algorithms for Data Placement in Arbitrary Networks," *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA '01)*, 2001.
- [4] S. Bhattacharya, H. Kim, S. Prabh, and T. Abdelzaher, "Energy-Conserving Data Placement and Asynchronous Multicast in Wireless Sensor Networks," *Proc. ACM Int'l Conf. Mobile Systems, Applications, and Services (MobiSys '03)*, 2003.
- [5] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web Caching and Zipf-Like Distributions: Evidence and Implications," *Proc. IEEE INFOCOM '99*, 1999.
- [6] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," *Proc. ACM MobiCom '98*, 1998.
- [7] G. Cao, "Proactive Power-Aware Cache Management for Mobile Computing Systems," *IEEE Trans. Computer*, vol. 51, no. 6, 2002.
- [8] M. Charikar and S. Guha, "Improved Combinatorial Algorithms for the Facility Location and k -Median Problems," *Proc. IEEE Symp. Foundations of Computer Science (FOCS '99)*, 1999.
- [9] F.A. Chudak and D. Shmoys, "Improved Approximation Algorithms for a Capacitated Facility Location Problem," *LNCS*, vol. 1610, 1999.
- [10] *The NS Manual*, K. Fall and K. Varadhan, eds., USC/ISI, LBL, UC Berkeley and Xerox Parc, <http://www-mash.cs.berkeley.edu/ns/>, 1995.
- [11] T. Hara, "Cooperative Caching by Mobile Clients in Push-Based Information Systems," *Proc. ACM Int'l Conf. Information and Knowledge Management (CIKM '02)*, 2002.
- [12] T. Hara, "Effective Replica Allocation in Ad Hoc Networks for Improving Data Accessibility," *Proc. IEEE INFOCOM '01*, 2001.
- [13] T. Hara, "Replica Allocation in Ad Hoc Networks with Periodic Data Update," *Proc. Int'l Conf. Mobile Data Management (MDM '02)*, 2002.
- [14] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *Proc. ACM MobiCom '00*, 2000.
- [15] K. Jain and V.V. Vazirani, "Approximation Algorithms for Metric Facility Location and k -Median Problems Using the Primal-Dual Schema and Lagrangian Relaxation," *J. ACM*, vol. 48, no. 2, 2001.
- [16] A. Vigneron, L. Gao, M.J. Golin, G.F. Italiano, and B. Li, "An Algorithm for Finding a k -Median in a Directed Tree," *Information Processing Letters*, vol. 74, 2000.
- [17] S. Jin and L. Wang, "Content and Service Replication Strategies in Multi-Hop Wireless Mesh Networks," *Proc. IEEE ACM Int'l Workshop Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '05)*, 2005.
- [18] S. Jin, "Replication of Partitioned Media Streams in Wireless Ad Hoc Networks," *Proc. ACM Int'l Conf. Multimedia (MULTIMEDIA '04)*, 2004.
- [19] K. Kalpakis, K. Dasgupta, and O. Wolfson, "Steiner-Optimal Data Replication in Tree Networks with Storage Costs," *Proc. Int'l Symp. Database Eng. and Applications (IDEAS '01)*, 2001.
- [20] P. Krishnan, D. Raz, and Y. Shavitt, "The Cache Location Problem," *IEEE/ACM Trans. Networking*, vol. 8, 2000.
- [21] J.-H. Lin and J. Vitter, "Approximation Algorithms for Geometric Median Problems," *Information Processing Letters*, vol. 44, no. 5, 1992.
- [22] P. Nuggehalli, V. Srinivasan, and C. Chiasserini, "Energy-Efficient Caching Strategies in Ad Hoc Wireless Networks," *Proc. ACM MobiHoc '03*, 2003.
- [23] C. Perkins, *Ad Hoc Networking*. Addison-Wesley, 2001.
- [24] C. Perkins and P. Bhagwat, "Highly Dynamic DSDV Routing for Mobile Computers," *Proc. ACM SIGCOMM '94*, 1994.
- [25] L. Qiu, V.N. Padmanabhan, and G.M. Voelker, "On the Placement of Web Server Replicas," *Proc. IEEE INFOCOM '01*, 2001.
- [26] S. Ratnasamy, B. Karp, S. Shenker, D. Estrin, R. Govindan, L. Yin, and F. Yu, "Data-Centric Storage in Sensornets with GHT, A Geographic Hash Table," *Mobile Networks and Applications*, vol. 8, no. 4, 2003.
- [27] C. Swamy and A. Kumar, "Primal-Dual Algorithms for Connected Facility Location Problems," *Proc. Int'l Workshop Approximation Algorithms for Combinatorial Optimization (APPROX '02)*, 2002.
- [28] A. Tamir, "An $O(pn^2)$ Algorithm for p -Median and Related Problems on Tree Graphs," *Operations Research Letters*, vol. 19, 1996.
- [29] B. Tang, H. Gupta, and S.R. Das, "Benefit-Based Data Caching in Ad Hoc Networks," *Proc. IEEE Int'l Conf. Network Protocols (ICNP '06)*, 2006.
- [30] B. Li, M.J. Golin, G.F. Italiano, and X. Deng, "On the Optimal Placement of Web Proxies in the Internet," *Proc. IEEE INFOCOM '99*, 1999.
- [31] J. Xu, B. Li, and D.L. Lee, "Placement Problems for Transparent Data Replication Proxy Services," *Proc. IEEE J. Selected Areas in Comm.*, vol. 20, no. 7, 2002.
- [32] Y. Xu, J. Heidemann, and D. Estrin, "Geography-Informed Energy Conservation for Sensor Networks," *Proc. ACM MobiCom '01*, 2001.
- [33] L. Yin and G. Cao, "Supporting Cooperative Caching in Ad Hoc Networks," *IEEE Trans. Mobile Computing*, vol. 5, no. 1, pp. 77-89, Jan. 2006.
- [34] G.K. Zipf, *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Addison-Wesley, 1949.



Bin Tang received the BS degree in physics from Peking University, China, in 1997, the MS degrees in materials science and computer science from Stony Brook University in 2000 and 2002, respectively, and the PhD degree in computer science from Stony Brook University in 2007. He is an assistant professor of computer science at Wichita State University, Wichita, Kansas. His research interests include wireless ad hoc and sensor networks. He is a

member of the IEEE.



Himanshu Gupta received the BTech degree in computer science and engineering from IIT, Bombay, in 1992 and the MS and PhD degrees in computer science from Stanford University in 1999. He joined the faculty of the Department of Computer Science, Stony Brook University in Fall 2002. His recent research activities focus on theoretical issues in wireless networking, particularly sensor networks and sensor databases. His other

research interests include database systems and theory, specifically materialized views, (multiple) query optimization, and data analysis. He is a member of the IEEE.



Samir R. Das received the PhD degree in computer science from Georgia Institute of Technology, Atlanta, in 1994. He is currently an associate professor in the Computer Science Department, State University of New York, Stony Brook. His research interests are in wireless networking and mobile computing, focusing on protocols, systems, and performance evaluation. He received the US National Science Foundation's CAREER Award in 1998 and the

Best Paper Award at the ACM MobiSys Conference in 2007. He has been a speaker in the Distinguished Visitor Program of the IEEE Computer Society from 2001 to 2003. He cochaired the technical program committee for the ACM MobiHoc Symposium in 2001 and the ACM MobiCom Conference in 2004. He currently serves or has served on the editorial board of *IEEE/ACM Transactions on Networking*, *IEEE Transactions on Mobile Computing*, *ACM/Kluwer Wireless Networks Journal*, and *Ad Hoc Networks*. More information about him and his research can be found at <http://www.cs.sunysb.edu/~samir>. He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**