

Cache Placement in Sensor Networks Under Update Cost Constraint

Bin Tang, Samir Das, and Himanshu Gupta

Computer Science Department,
Stony Brook University, Stony Brook NY 11790, USA
{bintang, samir, hgupta}@cs.sunysb.edu.edu

Abstract. In this paper, we address an optimization problem that arises in context of cache placement in sensor networks. In particular, we consider the cache placement problem where the goal is to determine a set of nodes in the network to cache/store the given data item, such that the overall communication cost incurred in accessing the item is minimized, under the constraint that the total communication cost in updating the selected caches is less than a given constant. In our network model, there is a single server (containing the original copy of the data item) and multiple client nodes (that wish to access the data item). For various settings of the problem, we design optimal, near-optimal, heuristic-based, and distributed algorithms, and evaluate their performance through simulations on randomly generated sensor networks.

1 Introduction

Advances in embedded processing and wireless networking have made possible creation of sensor networks [1, 9]. A sensor network consists of sensor nodes with short-range radios and limited on-board processing capability, forming a multi-hop network of irregular topology. Sensor nodes must be powered by small batteries, making energy efficiency a critical design goal. There has been a significant interest in designing algorithms, applications, and network protocols to reduce energy usage of sensors. Examples include energy-aware routing [13], energy-efficient information processing [8, 9], and energy-optimal topology construction [21]. In this article, we focus on designing techniques to conserve energy in the network by caching data items at selected sensor nodes in a sensor network. The techniques developed in this paper are orthogonal to some of the other mentioned approaches, and can be used in combination with them to conserve energy.

Existing sensor networks assume that the sensors are preprogrammed and send data to a sink node where the data is aggregated and stored for offline querying and analysis. Thus, sensor networks provide a simple sample-and-gather service, possibly with some in-network processing to minimize communication cost and energy consumption. However, this view of sensor network architecture is quite limited. With the rise in embedded processing technology, sensor networks are set to become a more general-purpose, heterogeneous, distributed databases

that generate and process time-varying data. As energy and storage limitations will always remain an issue – as much of it comes from pure physical limitations – new techniques for efficient data handling, storage, and dissemination must be developed. In this article, we take a general view of the sensor network where a subset of sensor nodes (called *servers*) generate data and another subset of nodes (called *clients*) consume this data. The data generation and consumption may not be synchronous with each other, and hence, the overall communication cost can be optimized by caching generated data at appropriately selected intermediate nodes. In particular, the data-centric sensor network applications which require efficient data dissemination [4, 6] will benefit from effective data caching strategies.

In our model of the sensor network, there is a single data item at a given server node, and many client nodes. (See Section 6 for a discussion on multiple data items and servers.) The server is essentially the data item producer and maintains the original copy of the item. All the nodes in the network cooperate to reduce the overall communication cost of accessing the data via a caching mechanism, wherein any node in the network can serve as a cache. A natural objective in the above context could be to select cache nodes such that the sum of the overall access and update cost is minimized. However, such an objective does not guarantee anything about the general distribution of energy usage across the sensor network. In particular, the updates always originate from the server node, and hence, the server node and the surrounding nodes bear most of the communication cost incurred in updating. Hence, there is a need to constrain the total update cost incurred in the network, to prolong the lifetime of the server node and the nodes around it – and hence, possibly of the sensor network. Thus, in this article, we address the cache placement problem to minimize the total access cost under an update cost constraint. More formally, we address the problem of selecting nodes in the network to serve as caches in order to minimize the total access cost (communication cost incurred in accessing the data item by all the clients), under the constraint that the total update cost (communication cost incurred in updating the cache nodes using an optimal Steiner tree over the cache nodes and the server) is less than a given constant. Note that since we are considering only a single data item, we do not need to consider memory constraints of a node.

Paper Outline. We start with formulating the problem addressed in this article and a discussion on related work in Section 2. For the cache placement problem under an update cost constraint, we consider a tree topology and a general graph topology of the sensor network. For the tree topology, we design an optimal dynamic programming algorithm in Section 3. The optimal algorithm for the tree topology can be applied to the general graph topology by extracting an appropriate tree from the given network graph. For the general graph topology, we consider a simplified multiple-unicast update cost model, and design a constant-factor approximation algorithm in Section 4.1. In Section 4.2, we present an efficient heuristic for the general cache placement problem under an update cost constraint, i.e., for a general update cost model in general graph topology. In Sec-

tion 4.3, we present an efficient distributed implementation. Finally, we present simulation results in Section 5, and give concluding remarks in Section 6.

2 Problem Formulation and Related Work

In this section, we formulate the problem addressed in this article. We start with describing the sensor network model.

Sensor Network Model. A sensor network consists of a large number of sensor nodes distributed randomly in a geographical region. Each sensor node has a unique identifier (ID). Each sensor node has a radio interface and can communicate directly with some of the sensor nodes around it. For brevity, we sometimes just use *node* to refer to a sensor node. The sensor network can be modeled as an undirected weighted graph $G = (V, E)$, where V is the set of nodes, and E is the set of edges formed by pairs of nodes that can directly communicate with each other. The communication distance between any two nodes i and j is the number of hops d_{ij} between the two nodes. The network has a data item, which is stored at a unique node called a *server*, and is updated at a certain update frequency. Each sensor node could be a client node. A client node i requests the data item with an *access frequency* a_i . The cost of accessing a data item (*access cost*) by a node i from a node j (the server or a cache) is $a_i d_{ij}$, where d_{ij} is the number of hops from node i to node j .

Problem. Informally, our article addresses the following *cache placement problem* in sensor networks. Select a set of nodes to store copies of the data item such that the total access cost is minimized under a given update cost constraint. The total access cost is the sum of all individual access costs over all clients for accessing the data item from the nearest node (either a cache or the server) having a copy of the data item. The *update cost* incurred in updating a set of caches M is modeled as the cost of the optimal Steiner tree [10] spanning the server and the set of caches. This problem is obviously NP-hard, as even the Steiner tree problem is known to be NP-hard [3]. In this article, we look at the above problem in various stages – a tree topology, a graph topology with a simplified update cost model, a graph topology with the general update cost model – and present optimal, approximation, and heuristic-based algorithms respectively.

More formally, given a sensor network graph $G = (V, E)$, a server r with the data item, and an update cost Δ , select a set of cache nodes $M \subseteq V$ ($r \in M$) to store the data item such that the total access cost

$$\tau(G, M) = \sum_{i \in V} a_i \times \min_{j \in M} d_{ij}$$

is minimum, under the constraint that the total update cost $\mu(M)$ is less than a given constant Δ , where $\mu(M)$ is the cost of the minimum Steiner tree over the set of nodes M . Note that in the above definition *all* network nodes are considered as potential clients. If some node i is not a client, the corresponding a_i would be zero.

Related Work. The general problem of determining optimal cache placements in an arbitrary network topology has similarity to two problems widely studied in graph theory viz., 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 for all the caches, without any constraint. On the other hand, the k -median problem minimizes the total access cost under the number constraint, i.e., 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 [7, 15], under the assumption that the edge costs in the graph satisfy the triangular inequality. Without the triangular inequality assumption, either problem is as hard as approximating the set cover [14], and therefore cannot be approximated better than $O(\log |V|)$ unless $\mathbf{NP} \subseteq \tilde{\mathbf{P}}$. Here, $|V|$ is the size of the network.

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 [19, 22]. In particular, Li et al. [19] address the optimal placement of web proxies in a tree topology, essentially designing an $O(n^3 k^2)$ time dynamic programming algorithm to solve the k -median problem optimally in a tree of n nodes. In other related works on cache placement in trees, Xu et al. [22] discuss placement of “transparent” caches to minimize the sum of reads and writes, Krishnan et al. [18] consider a cost model based on cache misses, and Kalpakis et al. [16] consider a cost model involving reads, writes, and storage. In sensor networks, which consist of a large number of energy-constrained nodes, the constraint on the number of cache nodes is of little relevance.

Relatively less work has been done for caching in sensor networks. Intanagonwatt et al. [6] propose directed diffusion, a data dissemination paradigm for sensor networks, which adopts a data centric approach and enables diffusion to achieve energy savings by selecting empirically good paths and by caching/processing data in-network. Bhattacharya et al. [4] develop a distributed framework that improve energy consumption by application layer data caching and asynchronous update multicast. In this article, we consider cache placement in sensor network under update cost constraint. As mentioned before, the update cost is typically mostly borne by the server and the surrounding nodes, and hence, is a critical constraint. To the best of our knowledge, we are not aware of any prior work that considers the cache placement problem under an update cost constraint.

3 Tree Topology

In this subsection, we address the cache placement problem under the update cost constraint in a tree network. The motivation of considering a tree topology (as opposed to a general graph model which we consider in the next section) is two fold. Firstly, data dissemination or gathering in sensor networks is typically done over an appropriately constructed network tree. Secondly, for the tree topol-

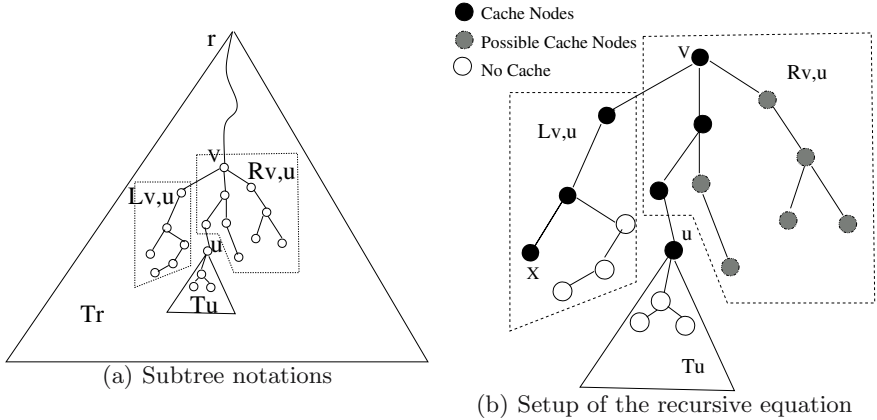


Fig. 1. Dynamic Programming algorithm for the tree topology

ogy, we can actually design polynomial time *optimal* algorithms. Thus, we can apply such optimal algorithms for the tree topology to the general graph topology by extracting an appropriate tree (e.g., shortest-path tree or near-optimal Steiner tree connecting the clients) from the general graph. In Section 5, we show through extensive simulations that such a strategy of applying an optimal tree algorithm to a general graph topology yields heuristics that deliver near-optimal cache placement solutions.

Consider an ad hoc network tree T rooted at the node r . Since the communication edges are bidirectional, any node in the network could be designated as the root; thus, we assume that the root node r is also the server for the data item. The cache placement problem under update cost constraint in a tree topology can be formally defined as follows.

Given the tree network T rooted at r , a data item whose server is r , and an update cost constraint Δ , find a set of cache nodes $M \subseteq T$ ($r \in M$) for storing copies of the data item, such that the total access cost $\tau(T, M) = \sum_{i \in T} a_i \times \min_{j \in M} d_{ij}$ is minimized under the constraint that the total update cost $\mu(M)$ is less than Δ , where $\mu(M)$ is the cost of minimum cost Steiner tree over M . Note that the minimum cost Steiner tree spanning over a set of nodes M is simply the smallest subtree connecting the root r to all the nodes in M .

3.1 Dynamic Programming Algorithm

In this subsection, we present an optimal dynamic programming algorithm for the above described cache placement problem under the update cost constraint in a tree topology. We first start with some subtree notations [19] that are needed to describe our dynamic programming algorithm.

Subtree Notations. Consider the network tree T rooted at r . We use T_u to denote the subtree rooted at u in the tree T with respect to the root r (i.e.,

the subtree rooted at u not containing r); the tree T_r represents the entire tree T . For ease of presentation, we use T_u to also represent the set of nodes in the subtree T_u . We use $p(i)$ to denote the parent node of a node i in the tree T_r . Let $\pi(i, j)$ denote the unique path from node i to node j in T_r , and $d_{k, \pi(i, j)}$ denote the distance of a node k to the closest node on $\pi(i, j)$.

Consider two nodes v and u in the network tree, where v is an ancestor of u in T_r . See Figure 1(a). Let $L_{v,u}$ be the subgraph induced by the set of nodes on the left of and excluding the path $\pi(v, u)$ in the subtree T_v , and $R_{v,u}$ be the subgraph induced by the set of nodes on the right of and including the path $\pi(v, u)$, as shown in Figure 1(a). It is easy to see T_v can be divided into three distinct subgraphs, viz., $L_{v,u}$, T_u , and $R_{v,u}$.

DP Algorithm. Consider a subtree T_v and a node x on the leftmost branch of T_v . Let us assume that all the nodes on the path $\pi(v, x)$ (including v and x) have already been selected as caches. Let $\tau(T_v, x, \delta)$ denote the optimal access cost for all the nodes in the subtree T_v under the *additional* update cost constraint δ , where we do *not* include the cost of updating the already selected caches on the path $\pi(v, x)$. Below, we derive a recursive equation to compute $\tau(T_v, x, \delta)$, which would essentially yield a dynamic programming algorithm to compute $\tau(T_r, r, \Delta)$ – the minimum value of the total access cost for the entire network tree T_r under the update cost constraint Δ .

Let O_v be an optimal set (not including and in addition to $\pi(v, x)$) of cache nodes in T_v that minimizes the total access time under the additional update cost constraint δ . Let u be the leftmost deepest node of O_v in T_v , i.e., the node u is such that $L_{v,u} \cap O_v = \emptyset$ and $T_u \cap O_v = \{u\}$. It is easy to see that adding the nodes along the path $\pi(v, u)$ to the optimal solution O_v does not increase the additional update cost incurred by O_v , but may reduce the total access cost. Thus, without loss of generality, we assume that the optimal solution O_v includes all the nodes along the path $\pi(v, u)$ as cache nodes, if u is the leftmost deepest node of O_v in T_v .

Recursive Equation. As described above, consider an optimal solution O_v that minimizes $\tau(T_v, x, \delta)$, and let u be the leftmost deepest node of O_v in T_v . Note that O_v does not include the nodes on $\pi(v, x)$. Based on the definition of u and possible cache placements, a node in $L_{v,u}$ will access the data item from either the nearest node on $\pi(v, u)$ or the nearest node on $\pi(v, x)$. In addition, any node in T_u will access the data item from the cache node u , while all other nodes (i.e., the nodes in $R_{v,u}$) will choose one of the cache nodes in $R_{v,u}$ to access the data item. See Figure 1(b). Thus, the optimal access cost $\tau(T_v, x, \delta)$ can be recursively defined in terms of $\tau(R_{v,u}, p(u), \delta - d_{u, \pi(v, x)})$ as shown below. Below, the quantity $d_{u, \pi(v, x)}$ denotes the shortest distance in T_v from u to a node on the path $\pi(v, x)$ and hence, is the additional update cost incurred in updating the caches on the path $\pi(v, u)$. We first define $S(T_v, x, \delta)$ as the set of nodes u such that the cost of updating u is less than δ , the additional update cost constraint. That is, $S(T_v, x, \delta) = \{u | u \in T_v \wedge (\delta > d_{u, \pi(v, x)})\}$.

Now, the recursive equation can be defined as follows.

$$\tau(T_v, x, \delta) = \begin{cases} \sum_{i \in T_v} a_i \times d_{i, \pi(v, x)} & \text{if } S(T_v, x, \delta) = \emptyset \\ \min_{u \in S(T_v, x, \delta)} \\ \left(\begin{array}{l} \sum_{i \in L_{v, u}} a_i \times \min(d_{i, \pi(v, u)}, d_{i, \pi(v, x)}) \\ + \sum_{i \in T_u} a_i d_{iu} \\ + \tau(R_{v, u}, p(u), \delta - d_{u, \pi(v, x)}) \end{array} \right) \end{cases}$$

In the above recursive equation, the first case corresponds to the situation when the additional update constraint δ is not sufficient to cache the data item at any more nodes (other than already selected cache nodes on $\pi(v, x)$). For the second case, we compute the total (and minimum possible) access cost for each possible value of u , the leftmost deepest additional cache node, and pick the value of u that yields the minimum total access cost. In particular, for a fixed u , the first term corresponds to the total access cost of the nodes in $L_{v, u}$. Note that for a node in $L_{v, u}$ the closest cache node is either on the path $\pi_{v, x}$ or $\pi_{v, u}$. The second and third terms correspond to the total access time of nodes in T_u and $R_{v, u}$ respectively. Since the tree T_u is devoid of any cache nodes, the cache node closest to any node in T_u is u . The minimum total access cost of all the nodes in $R_{v, u}$ can be represented as $\tau(R_{v, u}, p(u), \delta - d_{u, \pi(v, x)})$, since the remaining available update cost is $\delta - d_{u, \pi(v, x)}$ where $d_{u, \pi(v, x)}$ is the update cost used up by the cache node u . The overall time complexity of the above dynamic programming algorithm can be shown to be $O(n^4 + n^3\Delta)$ by careful precomputation.

4 General Graph Topology

The tree topology assumption makes it possible to design a polynomial-time optimal algorithm for the cache placement problem under update cost constraint. In this subsection, we address the cache placement problem in a general graph topology. In the general graph topology, the cache placement problem becomes NP-hard. Thus, our focus here is on designing polynomial-time algorithms with some performance guarantee on the quality of the solution.

As defined before, the total update cost incurred by a set of caches nodes is the minimum cost of an optimal Steiner tree over the set of cache nodes and the server; we refer to this update cost model as the Steiner tree update cost model. Since the minimum-cost Steiner tree problem is NP-hard in general graphs, we solve the cache placement problem in two steps. First, we consider a simplified multiple-unicast update cost model and present a greedy algorithm with a provable performance guarantee for the simplified model. Then, we improve our greedy algorithm based upon the more efficient Steiner tree update cost model.

4.1 Multiple-Unicast Update Cost Model

In this section, we consider the cache placement problem for general network graph under a simplified update cost model. In particular, we consider the

multiple-unicast update cost model, wherein we model the total update cost incurred in updating a set of caches as the sum of the individual shortest path lengths from the server to each cache node. More formally, the total update cost of a set M of cache nodes is $\mu(M) = \sum_{i \in M} d_{si}$, where s is the server. Using this simplified update cost model, the cache placement problem in general graphs for update cost constraint can be formulated as follows.

Problem Under Multiple-Unicast Model. Given an ad hoc network graph $G = (V, E)$, a server s with the data item, and an update cost Δ , select a set of cache nodes $M \subseteq V$ ($s \in M$) to store the data item such that the total access cost $\tau(G, M) = \sum_{i \in V} a_i \times \min_{j \in M} d_{ij}$ is minimum, under the constraint that the total update cost $\mu(M) = \sum_{i \in M} d_{si} < \Delta$.

The cache placement problem with the above simplified update cost model is still NP-hard, as can be easily shown by a reduction from the k -median problem. A number of constant-factor approximation algorithms have been proposed [7, 15] for the k -median problem which can also be used to solve the above cache placement problem. However, all the constant-factor approximation algorithms are based on the assumption that the edge costs in the network graph satisfy the triangular inequality. Moreover, the proposed approximation algorithms for k -median problem cannot be easily extended to the more efficient Steiner tree update cost model. Below, we present a greedy algorithm that returns a solution whose “access benefit” is at least 63% of the optimal benefit, where access benefit is defined as the reduction in total access cost due to cache placements.

Greedy Algorithm. In this section, we present a greedy approximation algorithm for the cache placement problem under the multiple-unicast update cost constraint in general graphs, and show that it returns a solution with near-optimal reduction in access cost. We start with defining the concept of a benefit of a set of nodes which is important for the description of the algorithm.

Definition 1. (Benefit of Nodes) Let A be an arbitrary set of nodes in the sensor network. The *benefit* of A with respect to an already selected set of cache nodes M , denoted as $\beta(A, M)$, is the decrease in total access cost resulting due to the selection of A as cache nodes. More formally, $\beta(A, M) = \tau(G, M) - \tau(G, M \cup A)$, where $\tau(G, M)$, as defined before, is the total access cost of the network graph G when the set of nodes M have been selected as caches. The *absolute benefit* of A denoted by $\beta(A)$ is the benefit of A with respect to an empty set, i.e., $\beta(A) = \beta(A, \emptyset)$.

The *benefit per unit update cost* of A with respect to M is $\beta(A, M)/\mu(A)$, where $\mu(A)$ is the total update cost of the set A under the multiple-unicast update cost model. \square

Our proposed Greedy Algorithm works as follows. Let M be the set of caches selected at any stage. Initially, M is empty. At each stage of the Greedy Algorithm, we add to M the node A that has the highest benefit per unit update cost with respect to M at that stage. This process continues until the update

cost of M reaches the allowed update cost constraint. The algorithm is formally presented below.

Algorithm 1. Greedy Algorithm

Input: A sensor network graph $V = (G, E)$.

Update cost constraint Δ .

Output: A set of cache nodes M .

BEGIN

$M = \emptyset$;

while $(\mu(M) < \Delta)$

Let A be the node with maximum $\beta(A, M)/\mu(A)$.

$M = M \cup \{A\}$;

end while;

RETURN M ;

END.

◇

The running time of the above greedy algorithm is $O(kn^2)$, where k is the number of iterations and n is the number of nodes in the network. Note that the number of iterations k is bounded by n .

Performance Guarantee of the Greedy Algorithm. The Greedy Algorithm returns a solution that has a benefit at least 63% of that of the optimal solution. The proof techniques used here are similar to the techniques used in [11] for the closed related view-selection problem in a data warehouse. Due to the space limitation, we omit the proof here.

Theorem 1 *Greedy Algorithm (Algorithm 1) returns a solution M whose absolute benefit is of at least $(1 - 1/e)$ times the absolute benefit of an optimal solution having the update cost (under the multiple-unicast model) of at most that of M .*

■

4.2 Steiner Tree Update Cost Model

Recall that the constant factor performance guarantee of the Greedy Algorithm described in previous section is based on the multiple-unicast update cost model, wherein whenever the data item in a cache nodes needs to be updated, the updated information is transmitted along the individual shortest path between the server and the cache node. However, the more efficient method of updating a set of caches from the server is by using the optimal (minimum-cost) Steiner tree over the selected cache nodes and the server. In this section, we improve the performance of our Greedy Algorithm by using the more efficient Steiner tree update cost model, wherein the total update cost incurred for a set of cache nodes is the cost of the optimal Steiner tree over the set of nodes M and the server of the data item.

Since the minimum-cost Steiner tree problem is NP-hard, we adopt the simple 2-approximation algorithm [10] for the Steiner tree construction, which constructs a Steiner tree over a set of nodes L by first computing a minimum

spanning tree in the “distance graph” of the set of nodes L . We use the term 2-approximate Steiner tree to refer to the solution returned by the 2-approximation Steiner tree approximation algorithm. Based on the notion of 2-approximate Steiner tree, we define the following update cost terms.

Definition 2. (Steiner Update Cost) The *Steiner update cost* for a set M of cache nodes, denoted by $\mu'(M)$, is defined as the cost of a 2-approximate Steiner tree over the set of nodes M and the server s .

The *incremental Steiner update cost* for a set A of nodes with respect to a set of nodes M is denoted by $\mu'(A, M)$ and is defined as the increase in the cost of the 2-approximate Steiner tree due to addition of A to M , i.e., $\mu'(A, M) = \mu'(A \cup M) - \mu'(M)$. \square

Based on the above definitions, we describe the Greedy-Steiner Algorithm which uses the more efficient Steiner tree update cost model as follows.

Algorithm 2. Greedy-Steiner Algorithm

Same as Algorithm 1 except μ is changed to μ' . \diamond

Unfortunately, there is no performance guarantee of the solution delivered by the Greedy-Steiner Algorithm. However, as we show in Section 5, the Greedy-Steiner Algorithm performs the best among all our designed algorithms for the cache placement problem under an update cost constraint.

4.3 Distributed Implementation

In this subsection, we design a distributed version of the centralized Greedy-Steiner Algorithm (Algorithm 2). Using similar ideas as presented in this section, we can also design a distributed version of the centralized Greedy Algorithm (Algorithm 1). However, since the centralized Greedy-Steiner Algorithm outperformed the centralized Greedy Algorithm for all ranges of parameter values in our simulations, we present only the distributed version of Greedy-Steiner Algorithm. As in the case of centralized Greedy-Steiner Algorithm, we cannot prove any performance guarantee for the presented distributed version. However, we observe in our simulations that solution delivered by the distributed version is very close to that delivered by the centralized Greedy-Steiner Algorithm. Here, we assume the presence of an underlying routing protocol in the sensor network. Due to limited memory resources at each sensor node, a proactive routing protocol [20] that builds routing tables at each node is unlikely to be feasible. In such a case, a location-aided routing protocol such as GPSR [17] is sufficient for our purposes, if each node is aware of its location (either through GPS [12] or other localization techniques [2, 5]).

Distributed Greedy-Steiner Algorithm. The distributed version of the centralized Greedy-Steiner Algorithm consists of rounds. During a round, each non-cache node A estimates its benefit per unit update cost, i.e., $\beta(A, M)/\mu'(A, M)$,

as described in the next paragraph. If the estimate at a node A is maximum among all its communication neighbors, then A decides to cache itself. Thus, during each round, a number of sensor nodes may decide to cache the data item according to the above criteria. At the end of each round, the server node gathers information from all the newly added cache nodes, and computes the Steiner tree involving all the selected cache nodes till the round. Then, the remaining update cost (i.e., the given update cost constraint minus the current update cost of the Steiner tree involving the selected cache nodes) is broadcast by the server to the entire network and a new round is initiated. If there is no remaining update cost, then the server decides to discard some of the recently added caches (to keep the total update cost under the given update cost constraint), and the algorithm terminates. The algorithm is formally presented below.

Algorithm 3. Distributed Greedy-Steiner Algorithm

Input: A network graph $V = (G, E)$.

Update cost constraint Δ .

Output: The set of cache nodes M .

BEGIN

$M = \emptyset$;

while $(\mu'(M) < \Delta)$

Let \mathcal{A} be the set of nodes each of which (denoted as A)
has the maximum $\beta(A, M)/\mu'(A, M)$ among its
non-cache neighbors.

$M = M \cup \mathcal{A}$;

end while;

RETURN M ;

END.

◇

Estimation of $\mu'(A, M)$. Let A be a non-cache node, and T_A^S be the shortest path tree from the server to the set of communication neighbors of A . Let $C \in M$ be the cache node in T_A^S that is closest to A , and let d be the distance from A to C . In the above Distributed Greedy-Steiner Algorithm, we estimate the incremental Steiner update cost $\mu'(A, M)$ to be d . The value d can be computed in a distributed manner at the start of each round as follows. As mentioned before, the server initiates a new round by broadcasting a packet containing the remaining update cost to the entire network. If we append to this packet all the cache nodes encountered on the way, then each node should get the set of cache nodes on the shortest path from the server to itself. Now, to compute d , each node only needs to exchange the above information with all its immediate neighbors.

Estimation of $\beta(A, M)$. A non-cache node A considers only its “local” traffic to estimate $\beta(A, M)$, the benefit with respect to an already selected set of cache nodes M . The local traffic of A is defined as the data access requests that use A as an intermediate/origin node. Thus, the local traffic of a node includes its own data requests. We estimate the benefit of caching the data item at A as

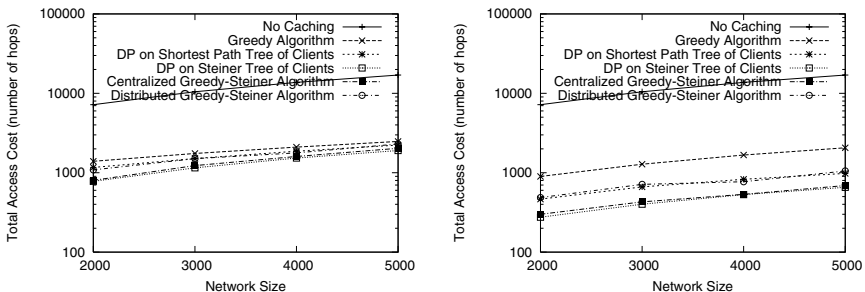
$\beta(A, M) = d \times t$, where t is the frequency of the local traffic observed at A and d is the distance to the nearest cache from A (which is computed as shown in the previous paragraph). The local traffic t can be computed if we let the normal network traffic (using only the already selected caches in previous rounds) run for some time between successive rounds. The data access requests of a node A during normal network traffic between rounds can be directed to the nearest cache in the tree T_A^S as defined in the previous paragraph.

5 Performance Evaluation

We empirically evaluate the relative performances of the cache placement algorithms for randomly generated sensor networks of various densities. As the focus of our work is to optimize access cost, this metric is evaluated for a wide range of parameters such as number of nodes and network and network density, etc.

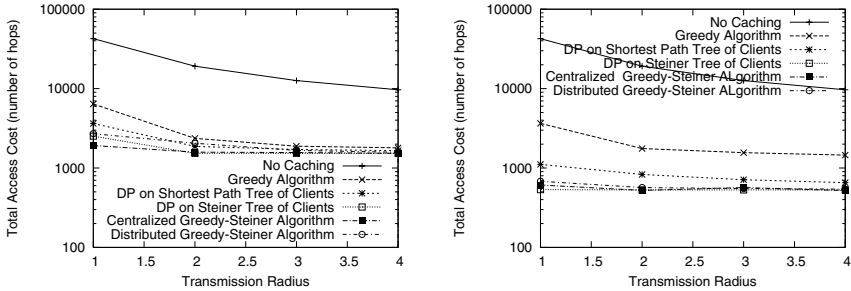
We study various caching schemes (listed below) on a randomly generated sensor network of 2,000 to 5,000 nodes in a square region of 30×30 . The distances are in terms of arbitrary units. We assume all the nodes have the same transmission radius (T_r), and all edges in the network graph have unit weight. We have varied the number of clients over a wide range. For clarity, we first present the data for the case where number of clients is 50% of the number of nodes, and then present a specific case with varying number of clients. All the data presented here are representative of a very large number of experiments we have run. Each point in a plot represents an average of five runs, in each of which the server is randomly chosen. The access costs are plotted against number of nodes and transmission radius and several caching schemes are evaluated:

- *No Caching* – serves as a baseline case.
- *Greedy Algorithm* – greedy algorithm using the multiple-unicast update cost model (Algorithm 1).



(a) Update cost = 25% of the near-optimal Steiner tree cost. (b) Update cost = 75% of the near-optimal Steiner tree cost.

Fig. 2. Access cost with varying number of nodes in the network for different update cost constraints. Transmission radius (T_r) = 2. Number of clients = 50% of the number of nodes, and hence increases with the network size.



(a) Update cost = 25% of the near-optimal Steiner tree cost. (b) Update cost = 75% of the near-optimal Steiner tree cost.

Fig. 3. Access cost with varying transmission radius (T_r) for different update cost constraints. Number of nodes = 4000, and number of clients = 2000 (50% of number of nodes).

- *Centralized Greedy-Steiner Algorithm* – greedy algorithm using the Steiner tree-based update cost model (Algorithm 2).
- *Distributed Greedy-Steiner Algorithm* – distributed implementation of the Greedy-Steiner Algorithm (Algorithm 3).
- *DP on Shortest Path Tree of Clients* – Dynamic Programming algorithm (Section 3.1) on the tree formed by the shortest paths between the clients and the server.
- *DP on Steiner Tree of Clients* – Dynamic Programming algorithm (Section 3.1) on the 2-approximate Steiner tree over the clients and the server.

Figure 2 shows the effect of the number of nodes; the transmission radius (T_r) is fixed at 2. Figure 3 shows the effect of T_r ; a network of 4,000 nodes is chosen for these experiments and T_r is varied from 1 to 4. The general trend in these two sets of plots is similar. Aside from the fact that our algorithms offer much less total access cost than the no-caching case, the plots show that (i) the two Steiner tree-based algorithms (DP on Steiner Tree of Clients and Centralized Greedy-Steiner Algorithm) perform equally well and the best among all algorithms except for very sparse graphs; (ii) the Greedy-Steiner Algorithm provides the best overall behavior; (iii) the Distributed Greedy-Steiner Algorithm performs very closely to its centralized version.

6 Conclusions

We have developed a suite of data caching techniques to support effective data dissemination in sensor networks. In particular, we have considered update cost constraint and developed efficient algorithms to determine optimal or near-optimal cache placements to minimize overall access cost. Minimization of access cost leads to communication cost savings and hence, energy efficiency. The

choice of update constraint also indirectly contributes to resource efficiency. Two models have been considered – one for a tree topology, where an optimal algorithm based on dynamic programming has been developed, and the other for the general graph topology, which presents a NP-hard problem where a polynomial-time approximation algorithm has been developed. We also designed efficient distributed implementations of our centralized algorithms, empirically showed that they performs well for random sensor networks.

References

1. B. Badrinath, M. Srivastava, K. Mills, J. Scholtz, and K. Sollins, editors. *Special Issue on Smart Spaces and Environments*, IEEE Personal Communications, 2000.
2. P. Bahl and V. N. Padmanabhan. Radar: An in-building RF-based user-location and tracking system. In *INFOCOM'00*.
3. P. Berman and V. Ramaiyer. Improved approximation algorithms for the steiner tree problem. *J. Algorithms*, 1994.
4. S. Bhattacharya, H. Kim, S. Prabh, and T. Abdelzaher. Energy-conserving data placement and asynchronous multicast in wireless sensor networks. In *MobiSys'03*.
5. N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5), 2000.
6. D. Estrin C. Intanagonwiwat, R. Govindan and J. Heidemann. Directed diffusion for wireless sensor networks. *IEEE/ACM TON*, 11(1):2–16, February 2003.
7. M. Charikar and S. Guha. Improved combinatorial algorithms for the facility location and k-median problems. In *IEEE FOCS'99*, pages 378–388.
8. M. Chu, H. Haussecker, and F. Zhao. Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. *IEEE Journal of High Performance Computing Applications*, 2002.
9. D. Estrin, R. Govindan, and J. Heidemann, editors. *Special Issue on Embedding the Internet*, Communications of the ACM, volume 43, 2000.
10. E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM J. Appl. Math.*, 16:1–29, 1968.
11. H. Gupta. *Selection and Maintenance of Views in a Data Warehouse*. PhD thesis, Computer Science Department, Stanford University, 1999.
12. B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins. *Global Positioning System: Theory and Practice*. Springer-Verlag Telos, 1997.
13. Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *MOBICOM'00*.
14. K. Jain and V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *J. ACM*, 48(2), 2001.
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. *Journal of the ACM*, 48(2):274–296, 2001.
16. K. Kalpakis, K. Dasgupta, and O. Wolfson. Steiner-optimal data replication in tree networks with storage costs. In *Proceedings of IDEAS*, pages 285–293, 2001.
17. Brad Karp and H.T Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *MOBICOM'00*.

18. P. Krishnan, D. Raz, and Y. Shavitt. The cache location problem. *IEEE/ACM TON*, 8:568–582, 2000.
19. B. Li, M. J. Golin, G. F. Italiano, and X. Deng. On the optimal placement of web proxies in the internet. In *INFOCOM'99*.
20. C. E. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing (dsv) for mobile computers. In *SIGCOMM'94*.
21. R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang. Distributed topology control for wireless multihop ad-hoc networks. In *INFOCOM'01*.
22. J. Xu, B. Li, and D. L. Lee. Placement problems for transparent data replication proxy services. *IEEE JSAC*, 20(7):1383–1397, 2002.