

Greedy is Good: On Service Tree Placement for In-Network Stream Processing

Zoe Abrams

Department of Computer Science
Stanford University, Stanford, CA
Email: zoea@stanford.edu

Jie Liu

Microsoft Research
Redmond, WA 98052
Email: liuj@microsoft.com

Abstract

This paper is concerned with reducing communication costs when executing distributed user tasks in a sensor network. We take a service-oriented abstraction of sensor networks, where a user task is composed of a set of data processing modules (called services) with dependencies. Communications in sensor networks consume significant energy and introduce uncertainty in data fidelity due to high bit error rate. These constraints are abstracted as costs on the communication graph. The goal is to place the services within the sensor network so that the communication cost in performing the task is minimized. In addition, since the lifetime of a node, the quality of network links, and the composition of the service graph may change over time, the quality of the placement must be maintained in the face of these dynamics. There exists a dynamic programming based algorithm for finding the optimum solution to the service placement problem. However, the algorithm is not flexible, often requiring the entire solution be recalculated in response to small, local changes. To address these challenges, we take a fresh look at what is generally considered a simple but poor performance approach for service placement, namely the greedy algorithm. We prove that a modified greedy algorithm is guaranteed to have cost at most 8 times the optimum placement. In fact, the guarantee is even stronger if there is a high degree of data reduction in the service graph. The advantage of the greedy placement strategy is that when there are local changes in the service graph or when a hosting node fails, the repair only affects the placement of services that depend on the changes. Simulations suggest that in practice the greedy algorithm finds a low cost placement. Furthermore, the cost of repairing a greedy placement decreases rapidly as a function of the proximity of the services to be aggregated.

I. INTRODUCTION

The possibly massive amounts of raw data and the large-scale, distributed, resource constrained nature of sensor networks motivate in-network processing that distills sensor data within the network before sending it to information consumers. A natural question to ask is where to place data processing modules (herein referred to as *services*, motivated by a service-oriented abstraction of the sensor network [1]), to achieve good overall performance and to conserve resources. The low power and usually mesh communication substrate brings a unique perspective to this problem. Communication in sensor networks is usually unreliable and costs a significant amount of energy, which motivates us to focus on reducing the cost of communication in achieving a user task. We use the notion of cost as a general abstraction here. It captures a combination of energy, bandwidth, and reliability concerns.

This problem is made more challenging in the presence of nodes and links that are unreliable. The network topology may change due to depleted batteries, node failure, or overcapacitated nodes. The cost of communication may also change if nodes are mobile or coupled closely with an evolving environment. In addition, the instantiation of services or the dependencies between services may change in situations such as conditional monitoring. For many applications, instead of sending changes to a central location that recomputes the optimum placement, it is desirable to *repair* the placements in the network using local and distributed algorithms. An ideal service placement strategy should be both *optimal*, in its placement quality, and *adaptable*, to changes in network and application topology.




Fig. 1. An illustration of the service placement problem.

As shown in Figure 1, we consider a connected network of sensor nodes that are distributed in the physical world. Each node can collect and process data as well as communicate with its neighbor nodes. A node also serves as a router to relay network traffic. A user can interact with the entire network through any node by, for example, sending queries and receiving answers. The communications have energy, bandwidth and reliability constraints, which introduce a cost for using communication along an edge per unit of data sent¹.

A user task is made up of services. These services form a *service composition graph* consisting of services and directed acyclic communication links that represent dependency constraints. Some services within a task must run on specific nodes. For instance, tasks that collect sensor data from a certain area must run on a node capable of performing the required sensing function. These services are said to be *anchored*. Other services are *floating* and can be placed on any node in the network. The goal is to find a *service placement* for floating nodes such that the total cost of running the task is minimized.

This problem formulation encompasses many sensor network applications, where the data collected by the sensors is viewed as streams, and user applications are operators applied on these streams. Since sensor data is distributed, yet usually spatially clustered, processing data near where it is collected can significantly reduce the amount of data transmitted. In a sensor database (e.g. Cougar [2] or TinyDB [3]), streams of data are collected at each sensor. Database operators such as MIN, MAX, and SUM can be applied on these data streams in the network to answer user queries. Since these operators can take an arbitrary number of operands and return only one number, they can significantly reduce network traffic when placed strategically along the paths that data are routed to the end user. In macro-programming paradigms such as regiment [4] and semantic streams [5], sensor streams are hierarchically constructed to inference high-level events and to trigger reactions. These user defined functions or inference units are placed in the network for resource efficiency and timely responses. In this paper we only consider service composition graphs that form a tree, where the sensors are leaves and the end information consumer is the root. Data only flows in one direction, from the sensors to the information consumer. Even with these assumptions, the formulation still covers a wide range of sensor network applications and can be applicable to other systems beyond sensor networks, such as overlay networks, workflow management for web service, and server farms.

A. Related Work

The *operator placement* problem, also called the *module assignment* or *task embedding* problem, for distributed tasks with precedence constraints, is one of the classical distributed computing problems [6], [7], [8], [9]. The problem shows up in many contexts such as overlay networks [10], grid computing [11], and streaming databases [12]. It has been show that the general operator placement problem is NP-hard, but polynomial time algorithms (e.g. based on dynamic programming) exist when the service graph is a tree [6].

¹We omit the cost of computation within a node.

The problem has recently gained increasing interests in sensor network communities due to the trend of interacting with an entire sensor network as one entity through, for example, database queries [13], [14], macroprogramming [4], [15], or service composition [5], [16]. While in the distributed computing context, the goal of operator placement is to minimize latency due to computation and communication, in sensor networks, the energy constraints and reliability of nodes must be taken into consideration.

The notion of filters with selectivity rates has also been studied previously. The work of [14], [17] considers optimum placement of filters with different selectivity rates so that the cost of execution and communication is minimized. However, their model is based on the notion that filters operate over a pre-existing aggregation tree and they exploit the freedom of re-ordering operators. Our work is different in that the data flow tree represents a work flow that must be executed in a specific order and must form an exact structure.

Finally, our approach involves finding a median in a distributed manner as a subroutine. Distributed algorithms for finding medians in general graphs [18] and trees [19] have been developed. These algorithms are designed to find a single median that has minimum average distance to every node in the network. In our scenario, we are interested in finding a minimum median for a particular subset of the nodes in the graph. One challenge tackled in this paper is the design of distributed algorithms that conserve energy by leveraging the proximity of nodes being aggregated.

This paper is organized as follows. In section II we formally define the service placement problem and outline a dynamic programming based optimum solution and an in-network relaxation heuristic. In section III-A, we introduce the greedy placement algorithm and show its performance bound for cases where data reduction rate at each service is high. In section III-B, we propose a modified greedy algorithm that can be used even when the data reduction rate does not satisfy the conditions in section III-A. Section III-C, describes how to implement the greedy algorithm in a distributed manner. We simulate and compare the performance of these algorithms in section IV.

II. SERVICE PLACEMENT PROBLEM

We formally define the service placement problem as follows. An underlying network is given as an undirected communication graph $G = (V, E)$, where V is the set of nodes and $E \subset V \times V$ is the set of edges connecting the nodes². Let w_e on edges $e \in E$ be the weight on the edge, i.e., the cost of communicating one unit of data across the edge. Let u and v be two nodes, $w_{(u,v)}$ is the sum of the weights on the shortest path from u to v in G .

We are also given a service graph in the form of a rooted tree $T = (O, L)$ where O is a set of services, and $L \subset O \times O$ is the directed dependency links among the services. That is, $l = (q, p) \in L$ represents that the outputs of service q feed into the input of service p . We denote \mathcal{R} the root of the tree, and C_p the (direct) children services of p . For each $l \in L$, d_l represents the amount of data communicated on link l , and d_p^i the total amount of data fed into node p . That is $d_p^i = \sum_{q \in C_p} d_{(q,p)}$. The total data to be shipped out of p will be called d_p^o and $r_p = \frac{d_p^o}{d_p^i}$ denotes the data reduction rate at service p , defined as the data out of p divided by the data from all children into p .

We further make the following assumptions:

- The communication graph is connected.
- Only leaf services S and the root service \mathcal{R} are anchored in the communication graph. This assumption simplifies the problem but does not exclude situations where an interior node of the service graph is anchored, since we can consider that this interior service is the root of a subproblem and then merge subproblem solutions together. Merging is possible because placements below an anchored service are independent of the placements above it.

²To make presentation clear, we call elements of the communication graph *nodes* and *edges* and elements of the service graph are called *services* and *links*.

- The edge weights in the communication graph satisfy triangle inequality and are symmetric (that is, $w_{(u,v)} = w_{(v,u)}$).
- The computation cost is ignored. In particular, we assume the computational power on each node is sufficient to host the entire user task. That is, we do not constrain the capacity of the nodes, and can place multiple services on the same node.

Definition 1: Service Placement Problem: Find an onto function $f : O \rightarrow V$ satisfying anchor assumptions and such that

$$\sum_{p \in I} \sum_{q \in C_p} d_{(q,p)}^o \cdot w_{(f(q),f(p))} \quad (1)$$

is minimized, where I is the set of interior services in the service tree including the root, i.e. $I = O - S$.

We refer to the value of equation (1) as the cost of placement f , $f(p)$ the host of p , and f^* the function f that minimizes equation (1).

A. Optimum Dynamic Programming Algorithm

When the service graph is a tree, a polynomial-time algorithm for optimum placement exists. In fact, the problem can be solved using dynamic programming.

Define function $C(p, u)$ on services $p \in O$ and nodes $u \in V$ to be the minimum possible communication cost of routing all descendants of p to node u . For every leaf service $s \in S$ anchored at node v in the communication graph, define $C(s, v) = 0$; and $C(s, u) = \infty$ if $u \neq v$. Then, the function C can be computed recursively using the following equation:

$$C(p, u) = \sum_{q \in C_p} \min_{x \in V} (w_{(x,u)} d_{(q,p)} + C(q, x))$$

After computing $C(p, u)$, $\forall p \in O, u \in V$, define the map f^* with the set of pairs $p \in O, u \in V$ used in the recursive unfolding of $C(\mathcal{R}, v_{\mathcal{R}})$, where the root service is anchored at node $v_{\mathcal{R}}$. Then, f^* is the optimal placement. This algorithm has running time $O(|V|^2|O|)$. Although it finds the optimum quickly, it has two problems.

- 1) The algorithm is centralized. It requires the precise knowledge of all-pair shortest paths.
- 2) The algorithm is global. A change of weights in the communication graph or a change in the topology or data rates of the service graph could potentially trigger an entire recalculation for the function C .

B. An In-Network Relaxation Algorithm

When the total knowledge of the underlying network is unclear, researches have suggested relaxation-based placement heuristics. Relaxation can be centralized by abstracting the communication costs among the nodes into values in a metric space [10]; or it can be in-network [13]. In an in-network relaxation scheme, based on an initial placement, a service hosting node locally decides whether placing the service on a neighbor node can reduce the overall cost, assuming that no other services are moving at the same time. A local migration usually will have a chain effect that triggers up-stream or down-stream services to migrate. The placement tree iterates through these local adjustment and tries to settle at a minimum total cost.

Relaxation-based algorithms are extremely simple and adaptive. They can also handle cases where the service composition graph is not a tree. However, the quality of the placement highly depends on the initial placement of the services, since it can easily fall into local minima when the communication costs on edges are not uniform or when the network topology is irregular (e.g. with holes).




Fig. 2. An service placement scenario shows that the greedy placement can be arbitrarily bad without data reduction.

III. A DISTRIBUTED GREEDY ALGORITHM

One of the key challenges of service placement in sensor networks is the ability to adapt to changes in the communication and service graphs. In this section, we develop the a modified greedy placement algorithm and a decentralized adaptation strategy to repair a greedy placement. The algorithm is simple and efficient and has proven guarantees on performance, regardless of the number of nodes in the network and the complexity of the service composition tree. We first describe a straightforward greedy approach. We then build upon the greedy algorithm with a simple, yet potentially crucial, modification. This section ends with a description of a distributed repair strategy of the modified greedy algorithm.

A. The Greedy Algorithm

We now present an $O(|V| \cdot |O|)$ greedy heuristic that places services hierarchically. For each service p , if $f(q)$ is defined for all children q of p , then assign p to node

$$u = \operatorname{argmin}_v \left\{ \sum_{q \in C_p} d_{(q,p)} w_{(f(q),v)} \right\}$$

Since placement of each service only depends on the placement of its children, the greedy algorithm has the potential to be easily distributed and local changes in the service tree only affect the assignment of its ancestors.

However, the greedy algorithm can be arbitrarily more costly than the optimum solution. Figure 2 shows a binary tree service composition graph (a) that is to be placed on a communication graph (b), where the weights of communication are labeled on the edges, and ϵ is an arbitrarily small number. The leaf nodes D_i and the root R are anchored to the corresponding network node with the same label. The optimum placement routes the data inward, i.e. each service is assigned to the node in the communication graph with the same name. The greedy algorithm, on the other hand, assigns all the nodes outward, until reaching A'_i , and then routes the message all the way back inward to the root R .

Let $0 \leq h \leq H$ be the depth of a service in the tree with $h_R = 0$, the amount of data sent out from each leaf be 1, and the data reduction rate at each intermediate node be r . Then the optimum assignment has cost

$$C^* = \sum_{h=1}^H 2^h (2r)^{H-h} = 2^H \sum_{h=1}^H (r)^{H-h}$$

while the cost of greedy assignment is

$$C = \sum_{h=1}^H 2^h (2r)^{H-h} (2^{H-h+1} - 1) = 2^H \sum_{h=1}^H (r)^{H-h} (2^{H-h+1} - 1)$$

Obviously, if $r \geq 1/2$, then $C/C^* \rightarrow \infty$ as $H \rightarrow \infty$, which indicates that the greedy algorithm can provide a solution with cost that is arbitrarily worse than the optimum. However, it is also interesting to observe that if $r < 1/2$, $C/C^* < \infty$. In particular, for $r = 1/4$, $C/C^* \leq 2$. That is, for this example, when the data reduction rate is $1/4$, the greedy algorithm is at most twice as bad as the optimum placement. We will now show that this is true for any problem instance.

We define the following notation:

- D_p are the descendants of p (excluding both p and the leaf nodes).
- $c(p)$ is the cost of routing data from the direct children nodes C_p to this node p in the greedy algorithm. In other words, for function f defined by the greedy algorithm, $c(p) = \sum_{q \in C_p} d_{(q,p)} w_{(f(q), f(p))}$. Similarly, $c^*(p)$ is the cost for the optimum placement.
- Recall that d_p^o is the total data to be shipped *out* of p and the total data fed into p is d_p^i .
- Recall $r_p = \frac{d_p^o}{d_p^i}$ is the data reduction rate at service p .
- π_{pq} (with q being a descendant of p) is the set of service nodes along the path from p to q , including q but not p . If $p = q$ then the set is empty.
- $|\pi_{pq}|$ is the cardinality of the set π_{pq} , i.e. the number of edges between q and p in the tree. If $p = q$ the length is zero.
- $r_{pq} = \prod_{h \in \pi_{pq}} r_h$ is the data reduction rate along the path from q to p , again including q but not including p . If $p = q$ the product over the empty set is 1.
- $\gamma^*(p)$ is the optimum cost of routing a total data amount d_p^i from leaves to p in the optimum solution. The total amount of data is distributed among the leaves proportional to the real traffic. Precisely, $\gamma^*(p) = \sum_{q \in (D_p \cup p)} r_{pq} c^*(q)$.

By definition, for each node p , the greedy algorithm picks a placement that minimizes the cost of sending data d_p^i from p 's child nodes to p . So, the cost $c(p)$ should be no greater than the cost of shipping the data from each $q \in C_p$ along the greedy assignment path all the way back to the leaf nodes, then along the optimum path according to $f^*(p)$. This implies the following inequality:

$$c(p) \leq \gamma^*(p) + \sum_{q \in D_p} r_{pq} c(q). \quad (2)$$

Using this inequality, we can derive the following lemma.

Lemma 1: For each node p , the greedy placement cost and the optimum placement cost satisfies:

$$c(p) \leq \sum_{q \in (D_p \cup p)} r_{pq} 2^{|\pi_{pq}|} c^*(q) \quad (3)$$

The proof is given in Appendix A.

Applying this lemma recursively, the cost of the entire tree is:

$$\begin{aligned} C(T) &= \sum_{p \in I} c(p) \\ &\leq \sum_{p \in I} \sum_{q \in (D_p \cup p)} r_{pq} 2^{|\pi_{pq}|} c^*(q) \\ &\leq \sum_{p \in I} \sum_{q \in (\pi_{\mathcal{R}_p} \cup \mathcal{R})} r_{qp} 2^{|\pi_{qp}|} c^*(p) \end{aligned}$$

If all data reduction rates are the same value R ,

$$\begin{aligned} C(T) &\leq \sum_{p \in I} \sum_{q \in (\pi_{\mathcal{R}_p} \cup \mathcal{R})} (2R)^{|\pi_{qp}|} c^*(p) \\ &\leq \sum_{p \in I} \frac{(2R)^{|\pi_{\mathcal{R}_p}|+2} - 1}{2R - 1} c^*(p) \end{aligned}$$

So, we have the following theorem.

Theorem 1: For a tree service placement problem, if the data reduction rate at every service satisfies $r \leq R < 1/2$, then

$$C(T) \leq \frac{1}{1 - 2R} C^*(T).$$

In particular, if $R = 1/4$, we have $C(T) \leq 2C^*(T)$. This result is independent of the size of the communication graph and the height of the service tree.

B. The Modified Greedy Placement Algorithm

When the data reduction rate is greater than or equal to $1/2$, the result from Theorem 1 does not apply. In this situation, we can cluster the original service tree and introduce “super-services” that have stronger data reduction rates. We then perform the greedy assignment algorithm on the modified service tree.

In particular, given a service tree $T = (O, L)$ and a desired data reduction rate R , we create a modified tree $T' = (O', L')$ by applying the following graph transformation:

1. Initialize $O' = \{s | s \in S\}$, $L' = \emptyset$.
2. For each service $p \in O$ with $C_p^T \subset O'$
 - If $r_p < R$: Add p to O' , add links from C_p^T to p into the set L' . Delete p from O . Here C_p^T are the children of p in T .
 - Otherwise, if $r_p \geq R$: For each $c \in C_p^T$, remove link (c, p) from L and add link (c, q) , where q is p 's parent in T . Delete p from O . Redefine $d_q^i = d_q^i - d_p^o + \sum_{c \in C_p^T} d_c^o$ and $r_q = \frac{d_q^o}{d_q^i}$.
3. If O contains only S and \mathcal{R} , end. Otherwise, go to Step 2.

The new tree T' thus created will have all the leaves of T , and all the data reduction rates less than R . We can then perform the greedy placement algorithm for T' .

Theorem 2: The optimum tree for routing T' costs at most $\frac{1}{R}$ times the optimum tree for routing T . Precisely, $C^*(T') \leq \frac{1}{R} C^*(T)$.

Proof: Take the original service tree T . For all nodes in $O - O'$, give the node a data reduction rate of 1 and call this tree \bar{T} . Clearly, $C^*(\bar{T}) \leq \frac{C^*(T)}{R}$ since the data rate at any node is increased by at most $\frac{1}{R}$. Now, any solution for \bar{T} can be used for T' to get the same total communication cost. ■

Theorem 3: For $R = 1/4$, the Modified Greedy Algorithm solution for T , $\hat{C}(T)$, has cost at most $8OPT$. Precisely, $\hat{C}(T) \leq 8C^*(T)$.

Applying the greedy algorithm to T' , $C(T') \leq \frac{1}{1-2R} C^*(T')$. Combining with theorem 2, $C(T') \leq \frac{1}{R-2R^2} C^*(T)$. Furthermore, $\hat{C}(T) \leq C(T')$, because we can route T using cost at most $C(T')$ by placing all excluded services at the host for the included service closest along the path to the root. We choose $R = 1/4$ to minimize $\frac{1}{R-2R^2}$, giving an approximation of 8. Again, this guarantee is independent of the height of the tree.




Fig. 3. A limited range flooding from three child nodes.

C. Distributed Implementation of Greedy

In situations of node, link, or service changes, an existing placement may no longer be valid. The goal of this section is to describe an efficient way to find a new placement for a service, and all services that depend on it, in a distributed manner such that the communication cost of the new placement has the same performance guarantee as that of the modified greedy algorithm. We consider the situation where there is a change in a single service placement. The changed service is one component in a larger service graph that has already been placed. The change may come from the node that is hosting the service or from a change in the structure of the service tree itself.

A brute force implementation when service p changes is for every child of p to flood the entire network. Each flooding sets up the gradient of cost to the particular child node. To replace the service p , every node computes the total cost if it were to be the median for all p 's children. Then, a leader election process can be performed to choose the median with minimum total cost. This process would need to be repeated for every ancestor of p .

We now show how, when the child services are proximate, the median can be found without flooding the entire network. The more closely located the set of siblings, the less flooding is needed to find the minimum cost median. For many sensor network applications, sibling services are proximate since usually the users are interested in data that is relatively local.

The general idea of the algorithm is to flood some small area around the siblings until *some* median is found (not necessarily the minimum cost median). Then, the found median is used to limit the search space for the optimum median. Figure 3 shows a limited range flooding emanating from three child nodes A , B , and C , represented by the stars. The circles are flooding ranges from each child node with cost less than C_A , C_B , and C_C respectively. D is a node where the three flooding ranges intersect. It may not be the minimum cost median, but the minimum cost median will be within the *union* of the three flooding ranges, since each node, like E , that is outside the union will introduce total cost greater than $C_A + C_B + C_C$ (the cost of choosing D to be the median).

In the following algorithm, we assume that the nodes are reasonably time synchronized. We denote the children, siblings, and parent of a service p with F_p and call the hosts of services F_p the *immediate relatives* of p . A slight abuse of notation, we assume sending $f(F_p)$ in a message includes the ID numbers of the immediate relatives of p along with information about which host holds which service (similarly

for C_q). We assume that if p is hosted at node h , the broadcast distances to the immediate relatives of p are known to h . If this assumption cannot be made, then the ID numbers of the immediate relatives are known and can be used to determine distances using standard distributed algorithms such as doubling broadcast distance until the requested node responds. We also assume that when a message is sent from some location, it will first be received by a node along a shortest path. It is assumed that all nodes know α , a parameter of the algorithm that is used as the initial value in limited cost flooding.

The algorithm performs in two stages: a limited cost local flooding from each children to find an initial median and a flooding to the union of the first flooding zone to find the optimal median. The algorithm is defined by 4 states. Each state has an event that triggers entry into the state, some activity that the node performs while in the state, and an event that triggers exit from the state. Note that the states are not exclusive, that is, a node can be in more than one states at the same time. In addition, there are 5 types of messages that nodes might send during the course of the algorithm. Where applicable, the recipient and sender of the message is indicated using unique node ID numbers. The $newHost(p, h, f(F_p))$ message indicates that service p , with immediate relatives $f(F_p)$, has a new host with ID number h . Regardless of the state, a node hosting service p that receives a $newHost(p, h, empty)$ message sends a $newHost(p, h, f(F_p))$ message to h and deletes p from the set of services it hosts. Upon receipt of a $newHost$ message, all nodes update the new service location if the information is relevant for the services it hosts. To prevent the description from becoming cumbersome we have omitted details that can either be inferred or implemented using standard techniques. These details include the precise content of the messages and the manner of calculating distance to a node using a message received from that node.

INITIATE STATE

Entry Event: A change in the network occurs requiring a new placement of service q .

Activity: Choose arbitrary child $p \in C_q$ with host $f(p)$. Send $newHost(p, h, f(F_p))$ message to $f(p)$ and $f(q)$ with $h = f(p)$ and $f(F_p)$ empty. Essentially, this message tells the node hosting p to move the service p from itself to itself in order to initiate the greedy algorithm.

Exit Event: Send $newHost$ message.

LEAD STATE

Entry Event: A $newHost(p, h, f(F_p))$ message is received and the new host ID matches my own ID, $f(F_p)$ is not empty, and p is not the root of the service graph.

Activity:

- 1) Set q to be the parent of p in the service graph.
- 2) Broadcast $initiate(self, q, f(C_q))$ message to hosts of C_q .
- 3) Upon first receipt of a $medianValue(v, id)$ message, set w equal to the time elapsed between entrance into the leader state until the current moment. Wait for $5w$ additional time steps to pass. Choosing the $medianValue$ message with the smallest v , and set $f(q)$ (the new host for q) to be the id from this message. Send message $newHost(q, f(q), empty)$ to nodes in $f(F_p)$.

Exit Event: Send $newHost$ message.

FLOOD STATE

Entry Event: Receive $initiate(p, q, f(C_q))$ message, where my ID w is in $f(C_q)$.

Activity: Set $J = \alpha$. J is a cost bound on the reach of the message such that the message will be received by node u iff $w_{(w,u)}d_w^o \leq J$.

LOOP: Broadcast a $flood_{radius}(J, p, q, f(C_q))$ message. If a $terminate(q)$ message is not received within given timeout, let $J = g(J)$ where g is monotonically increasing in J , and return to LOOP (perform another bounded depth flooding).

Exit Event: Receive $terminate(q)$ message.

SEARCH STATE

Entry Event: Receive $flood_{radius}(J, p, q, f(C_q))$ message.


Activity:

- 1) Retransmit any $flood_{union}(union, p, q, f(C_q), v)$ message received.
- 2) If receive $flood_{radius}$ or $flood_{union}$ messages from all hosts of services C_q
 - If the all messages are $flood_{radius}$ messages, send $terminate(q)$ message to all nodes $f(C_q)$.
 - Compute $v = \{\text{the cost of hosting } q \text{ myself}\}$.
 - Compute $v^* = \{\text{the smallest from amongst } v \text{ and any median value received in any } flood_{union} \text{ message so far}\}$.
 - Send $medianValue(v, self)$ message to the leader p iff $v \leq v^*$
 - Send $flood_{union}(union, p, q, f(C_q), v^*)$ messages to all neighbors.


Exit Event: Send $flood_{union}$ messages to all neighbors.

Consider the distributed algorithm at work in the example from Figure 3. Node F will soon run out of batteries and the service p it is hosting must be placed at a different node. Node F chooses to assign leadership to node B , because B is hosting $q \in C_p$, an arbitrarily chosen child of p . Leader node B sends a message to the hosts of p 's children, A , B , and C , telling them to initiate their flooding. When node D hears from A , B , and C , it sends them all a $terminate$ message to stop their flooding processes. D computes the cost of hosting p itself, and sends this value to leader node B . Node D also forwards its median value and distances to A , B , and C on to its neighbors so that the gradients will continue to grow into the union space. Similarly, all of the nodes inside the union space will forward the gradients within the union while they are waiting for the gradient from all children of p . Once all gradients are received, they too will compute the cost of hosting p themselves and send the cost of hosting p themselves to the leader B in a $medianValue$ message if it is a better host than D . Once B receives the median value message from D , it expects to hear from other nodes in the union within 5 times the amount of time elapsed so far since weights are symmetrical. Leader B chooses a host for p that has minimum cost, say node K and sends this decision to the hosts of the immediate relatives of q so that they can update their information, including the old host site for p (node F). Node F fills in information about the hosts of F_p and retransmits this message to K so that K has all the information it needs to properly host p . Now K is the leader of a new median search for a node to host the parent of p , unless p is the root.


For each service p placed in this process, $2|C_p|$ bounded depth floodings are needed. If this amount of flooding is still too expensive and the optimality of the solution can be further relaxed, we can use the first median found (node D in the example from Figure 3 as an approximation to the minimum median. In fact, we can improve the algorithm by using time varying flooding. That is, the forwarding of the bounded depth flooding message is delayed at each node proportional to the amount of data transmitted along the incoming link. That is, at node u delay $T = \tau d_u^o w_{(u,v)}$ where v is the node that first forwarded the message to u and τ is a parameter of the algorithm. In this way, the node where the flooding messages meet is closer to services sending more data. To see that the median found in this way can only be an approximation rather than the optimum, consider a node with two children. The minimum cost median should be placed at the child that is sending out more data (even if the child sends only slightly more data). However, the first median found is somewhere between the two children.




(a) Perturbed grid.



(b) Perturbed grid with a whole centered at (300, 600) with radius 200.



(c) Randomly distributed nodes.



(d) Randomly distributed nodes with a whole centered at (300, 600) with radius 200.

Fig. 4. Examples of network topologies used in the simulation.

IV. SIMULATIONS

We evaluate both the placement and adaptation costs of the greedy algorithm through simulation. The placement cost is the value of the placement given by equation (1) from section II, and the adaptation cost is the amount of energy expended in adapting the placement after a change or failure in the network. We use four kinds network topology: perturbed grid, perturbed grid with a hole, random topology, and random topology with a hole. Figure 4 shows examples of these topologies. In a square region of 1000x1000, various numbers of sensors are deployed. The communication range of each node is chosen uniformly from $[0.5\phi, 1.5\phi]$, where ϕ is a parameter to control different network densities. The cost on the edges are uniformly distributed in $[10, 15]$. The root of the service tree is always set at the bottom left corner of the network, i.e., at coordinate (0, 0).

A. Placement Cost

We use two set of examples to study placement performance in the simulations: a sparse tree that is inspired by a parking garage application, and a regional data collection tree inspired by TinyDB. In the first example, we compare the greedy placement results with the optimum placement and the in-network relaxation based placement from section II-B.

1) *Sparse service tree*: We place 64 nodes in the field and pick $\phi = 100$. We use a service graph as shown in Figure 5(a) as a representative for a class of distributed event detection applications. This particular graph is inspired by the parking garage scenarios in [1] and [20]. Assume that a set of wireless webcams are installed in a parking structure. A set of sensors (e.g. break beams or RFID readers, labeled as S_1, \dots, S_5 in the figure) are deployed near the entrance that detect (service S_6) and identify (service S_7) incoming vehicles. Based on the driver's parking ' preference, a set of cameras (S_8 and S_9) near

the desired parking place are cued. Images are pieced together (service S_{10}) and open slots are counted (service S_{11}). The locations of the open spots are returned to a display (service S_{12}).




Fig. 5. A service tree example inspired by a parking garage scenario.

The output data rate from each sensor and service is labeled on the links. The data reduction rates $r_{S6} = 1/5$, $r_{S7} = 10$, $r_{S10} = 5/21$, and $r_{S11} = 1/50$. Mapping to the sensor network layout, we assume $S1, \dots, S5$ are near $(700, 700)$; cameras $S8$ and $S9$ are near $(900, 100)$, and the root service S_{12} is at the origin $(0,0)$. Note that not all data reduction rates are less than $1/4$. We convert the service graph into the one shown in Figure 5(b) using the modified greedy algorithm.

Figure 6 shows an example of the final result from different placement algorithms. In this assignment $C_{greedy}/C_{opt} = 1.11$, $C_{modified}/C_{opt} = 1.02$, and $C_{relax}/C_{opt} = 2.19$.

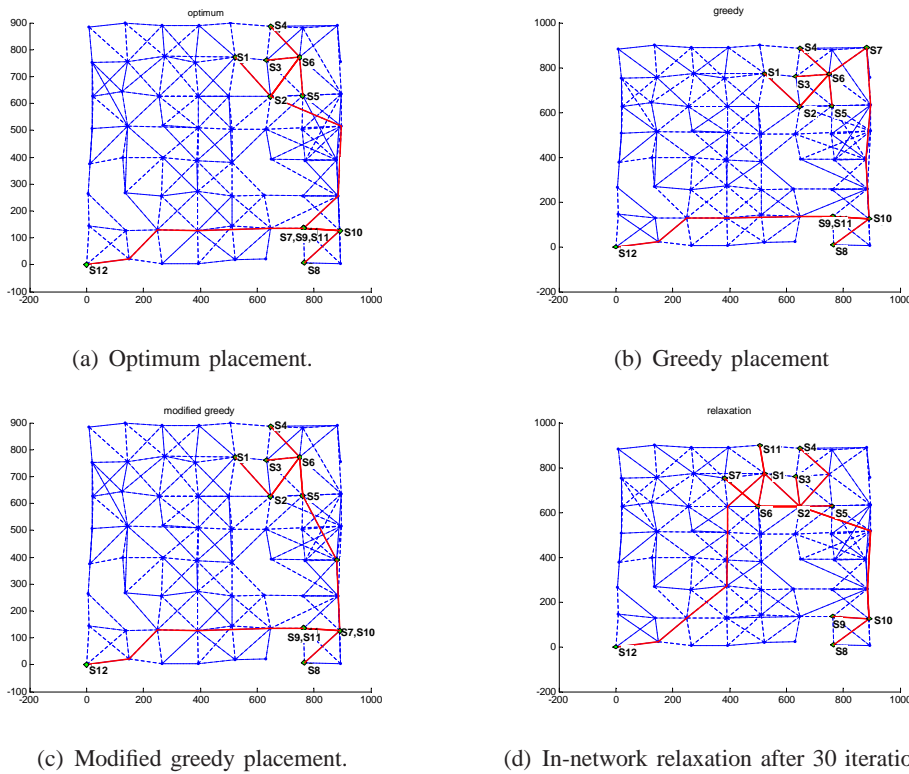


Fig. 6. Examples of placement results under different algorithms.

To collect performance statistics, we run each algorithm on 32 randomly generated networks in each topology category and compute the ratio of the total cost over the optimal cost. Figure 7 shows the statistics of the results. Each network topology is shown in a separate subfigure. The min, median, and

max ratio are plotted. In all these examples, the modified greedy placement performs better than greedy placement, which is better than in-network relaxation.

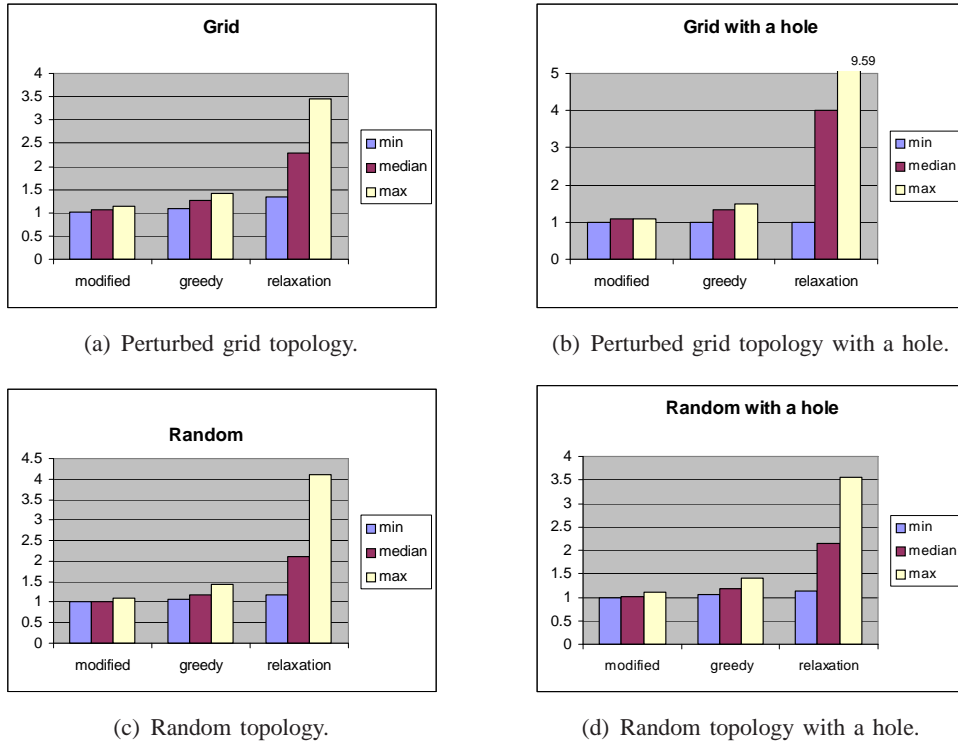


Fig. 7. Statistics for the cost ratio between different placement algorithms and the optimum placement.

2) *Data aggregation*: The next set of experiments test the greedy placement algorithm in data aggregation applications. The network is set up similar to Figure 4(a) and a base station is placed at the bottom left corner. The tasks are aggregation queries, such as MAX, computed over a subset of sensor inputs. We compare the performance of greedy placement with TinyDB [3] like aggregation trees, which we call TAG trees. A TAG tree is built without prior knowledge of what sensor data will be collected. An aggregator, like MAX, is placed at every node on a predetermined routing tree built using the surge routing protocol³. Since these aggregators can take an arbitrary number of inputs and produce only one output with the same data type, their data reduction rate equals the inverse of the number of children the node has in the TAG tree.

In these data aggregation scenarios, the optimum placement of aggregators is a minimum-cost Steiner tree, which is a well studied NP-hard problem [21]. In this section, we apply the greedy placement algorithm to an n -ary balanced tree. We compare our placement cost with the cost of a TAG tree. By using an n -ary tree, we effectively create aggregation services with data reduction rate $1/n$.

In particular, given a network of N nodes in a square field, we select $m = \lfloor \sqrt{N} \rfloor$ nodes as source sensors. We place these sources “evenly” along the anti-diagonal line (from upper left to lower right), so that we give the greedy algorithm enough space (i.e. the upper-right half of the field) to make sub-optimal choices. To achieve this effect, we divide the field into an $m \times m$ grid, and select one sensor from every anti-diagonal grid. We run 32 experiments for each $N = 64, 128, 256, \text{ and } 512$ under a perturbed grid topology. In each configuration we vary the data reduction rate R (i.e. the branching factor in the aggregation tree) to be 2, 4, 6, and 8.

³e.g. from TinyOS (<http://www.tinyos.net>)

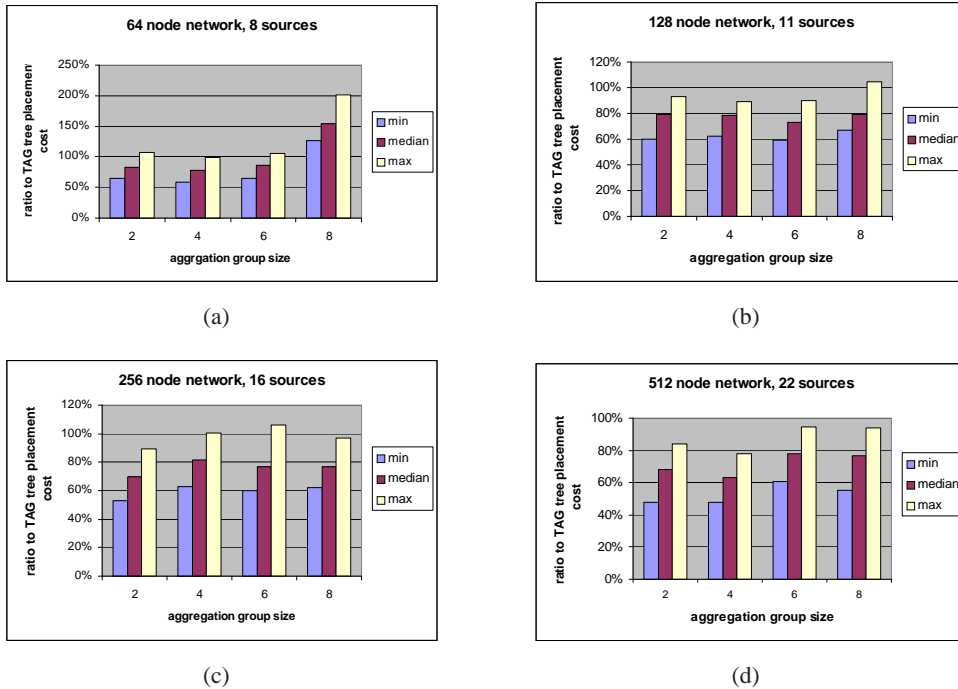


Fig. 8. Statistics on the ratio between the greedy algorithm cost on an x -ary tree and the cost when using a predefined routing tree.

Figure 8 plots the ratios between the cost of the greedy algorithm placement on trees with varying branching factors and the cost of placement using a predefined TAG tree. In almost all cases, the greedy placement performs better than the TAG tree. The result is not surprising, since by breaking up the aggregator, we have created a specific tree for the specific set of data sources. A notable exception is in Figure 8(a) with group size 8. Since there are only 8 data sources in the configuration, by building a 8-ary tree, we are not doing any in-network aggregation.

B. Adaptation Cost

In this section, we evaluate the performance of our distributed greedy algorithm and, in particular, the effectiveness of using bounded depth flooding to find the minimum cost weighted median. There are three key metrics to evaluate: the flooding zone union, the number of *medianValue* messages sent, and the quality of the initial median. The flooding zone union is defined as the number of nodes reached in the flooding processes, or, the number of nodes that enter the SEARCH STATE. The smaller this zone, the less energy and network traffic is spent forwarding flooding messages. The number of *medianValue* messages sent is the number of nodes in the flooding zone union that have cost smaller than the cost of the first median found at the intersection of the individual flooding zones. A node only sends its median value to the leader node if it is less than the median value found initially. Therefore, the fewer the number of nodes that would cost less than the initial median, the fewer *medianValue* messages will be sent. Reducing the number of *medianValue* messages conserves valuable energy and communication resources. Finally, the quality of the initial median is defined as the ratio between the minimum cost median and cost of the initial median found. The better the quality of the initial median, the more promise there is in using it as a substitute for the optimum median when communication resources are severely constrained.

We run the simulation in networks with 512 nodes randomly scattered in a physical space of size 1000×1000 . In a square region of size 200×200 at the center of the space, we randomly place a number of child services who need to find a host for their parent node. These children are in the center of the space so that the flooding process can reach a large number of nodes without hitting the edge of the network.

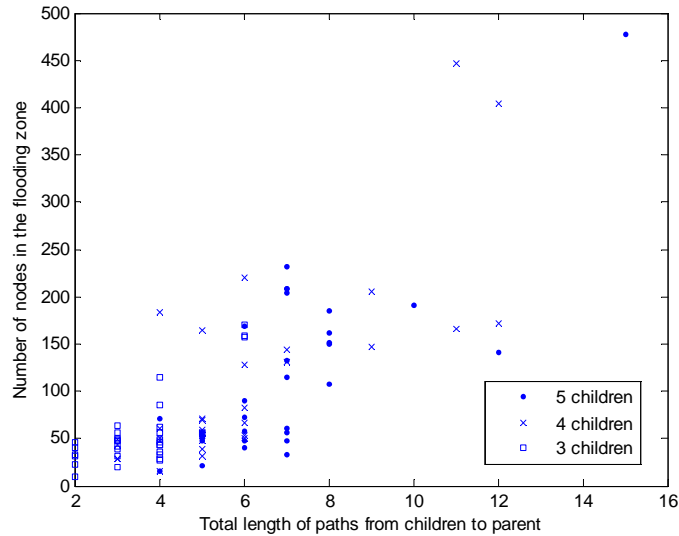


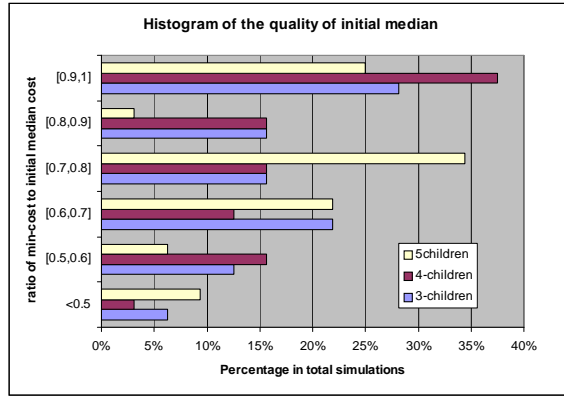
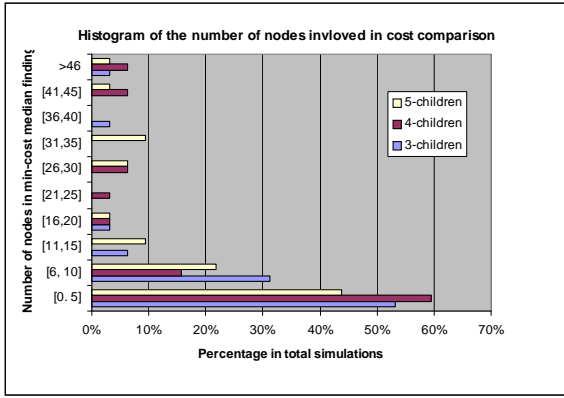
Fig. 9. The number of nodes in the flooding zone vs. how close the children are.

All children services have the same amount of outgoing data. The number of children is a simulation parameter, chosen from $\{3, 4, 5\}$. We run 32 simulations for each parameter.

Figure 9 shows the size of the flooding zone union under different runs. The X-axis is the total length of the communication paths from child hosts to the chosen parent host. This is used to measure the closeness of the child hosts. The smaller this number, the closer the child hosts are, and thus the smaller the flooding zone union should be. The unweighted median cost also serves as a lower bound on the amount of communication that must be expended in finding this median (assuming the only way to determine distances is through direct communication between nodes). Since the flooding zones stop growing when all children gradients intersect, the more children, the larger the flooding zone union tends to be. The significant reduction in the size of the flooding zone union as the proximity of the children nodes increases, emphasizes that the performance of the distributed greedy algorithm is dependent on the proximity of the children nodes.

Figure 10(a) shows a histogram of 96 simulations, bucketed according to the number of *medianValue* messages in the simulation. In the majority of the simulations, less than 10 nodes had a cost that was smaller than the cost of the initial median. Less than 9% of the simulations had more than 46 nodes send *medianValue* messages. Out of a total of 512 nodes, this shows a significant reduction in communication and processing costs as a result of the low cost of the initial median. Figure 10(b) shows the distribution of the cost ratio between the optimum median and the initial median. In the majority of the simulations, the ratio is more than $\frac{2}{3}$, with about $\frac{1}{3}$ of the simulations resulting in an optimum median that is more than .9 the initial value. Although there is still more to understand about the full implications of using the initial median in place of the optimum, these results indicate that this option has real potential.

It should be pointed out that sometimes the objectives of minimizing the flooding zone union and minimizing the cost of the initial median are in conflict. For instance, it is often the case that setting the individual flooding radii to be the same for all nodes, independent of the amount of outgoing data at that node, will result in a smaller union flooding zone but will increase the value of the initial median found. Further exploration into the tradeoffs between these two metrics and their relative importance in various application scenarios is an interesting direction for future work.



(a) The distribution of the number of nodes whose cost is smaller than the cost of the initial median. (b) The distribution of the cost ratio between the minimum median and the initial median.

Fig. 10. The number of *medianValue* messages and the quality of the initial median.

V. CONCLUSION

In this paper we studied the performance bound for a modified greedy algorithm for service tree placement and proposed a distributed scheme to adapt to network and service tree changes. We recommend using the the optimum dynamic programming solution when all information is centralized, the network is relatively stable, or initially when a task is deployed. When the network is dynamic, the greedy algorithm is a fast, simple, distributed, and efficient (thus, “good”) alternative with both provable guarantees and easy adaptivity.

The simulations in section IV-A.2 suggest that performance improves when the routing tree is defined based on knowledge about the modules to be aggregated. As would be expected, tailoring the routing to the specific user task leads to significant reductions in energy consumption. This motivates further exploration into the possibility of a new paradigm in which routing structures are defined in the network, in real time, and can evolve, adapt, and change as the task at hand changes.

APPENDIX

A. We now prove Lemma 1 from section III-A using induction.

Proof: Base Case: By definition of greedy placement, $c(p) \leq c^*(p)$ for nodes directly above the leaves.

Inductive Hypothesis: Equation (3) is true $\forall q \in D_p$, then from (2)

$$\begin{aligned}
c(p) &\leq \gamma^*(p) + \sum_{q \in D_p} r_{pq} c(q) \\
&\leq c^*(p) + \sum_{q \in D_p} r_{pq} (c(q) + c^*(q)) \\
&\leq c^*(p) + \sum_{q \in D_p} r_{pq} \left(\sum_{o \in (D_q \cup q)} r_{qo} 2^{|\pi_{qo}|} c^*(o) + c^*(q) \right) \\
&\leq c^*(p) + \sum_{q \in D_p} r_{pq} c^*(q) + \sum_{q \in D_p} \sum_{o \in (D_q \cup q)} r_{po} 2^{|\pi_{qo}|} c^*(o) \\
&\leq c^*(p) + \sum_{q \in D_p} r_{pq} c^*(q) + \sum_{q \in D_p} r_{pq} c^*(q) \sum_{o \in \pi_{pq}} 2^{|\pi_{qo}|} \\
&\leq c^*(p) + \sum_{q \in D_p} r_{pq} c^*(q) + \sum_{q \in D_p} r_{pq} c^*(q) (2^{|\pi_{qp}|} - 1) \\
&\leq c^*(p) + \sum_{q \in D_p} P_{pq} 2^{|\pi_{pq}|} c^*(q) \\
&\leq \sum_{q \in (D_p \cap p)} P_{pq} 2^{|\pi_{pq}|} c^*(q)
\end{aligned}$$

■

REFERENCES

- [1] J. Liu and F. Zhao, "Towards semantic services for sensor-rich information systems," in *Proc. of the 2nd Intl. Workshop on Broadband Advanced Sensor Networks (Basenets'05)*, Boston, MA, October 2005.
- [2] P. Bonnet, J. Gehrke, and P. Seshadri, "Towards sensor database systems," *Lecture Notes in Computer Science*, vol. 1987, pp. 3–14, 2001. [Online]. Available: citeseer.ist.psu.edu/article/bonnet01towards.html
- [3] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong, "TAG: a Tiny AGgregation Service for Ad-Hoc Sensor Networks," in *OSDI*, December 2002.
- [4] R. Newton and M. Welsh, "Region streams: Functional macroprogramming for sensor networks," in *Proceedings of the First International Workshop on Data Management for Sensor Networks (DMSN)*, Toronto, Canada, August 2004.
- [5] K. Whitehouse, F. Zhao, and J. Liu, "Semantic streams: a framework for the composable semantic interpretation of sensor data," in *Proc. European Workshop on Wireless Sensor Networks (EWSN'06)*, Zurich, Switzerland, Feb. 2006.
- [6] S. H. Bokhari, "A shortest tree algorithm for optimal assignments across space and time in a distributed processor system," *IEEE Trans. Software Eng.*, vol. SE-7, no. 6, pp. 583–589, 1981.
- [7] M. G. Norman and P. Thanisch, "Models of machines and computation for mapping in multicomputers," *ACM Computing Surveys*, vol. 25, no. 3, pp. 263–302, 1993.
- [8] B. Veltman, B. J. Lageweg, and J. K. Lenstra, "Multiprocessor scheduling with communication delays," *Parallel Computing*, vol. 16, no. 2-3, pp. 173–182, 1990.
- [9] A. Billionnet, "Allocating tree structured programs in a distributed system with uniform communication costs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 5, no. 4, pp. 445–448, 1994.
- [10] P. Pietzuch, J. Ledlie, J. Shneidman, M. Roussopoulos, M. Welsh, and M. Seltzer, "Network-aware operator placement for stream-processing systems," in *Proc. 22nd Intl. Conference on Data Engineering (ICDE'06)*, Atlanta, GA, April 2006.
- [11] J. Cao, S. A. Jarvis, S. Saini, and G. R. Nudd, "GridFlow: Workflow management for grid computing," in *Proc. Intl. Symposium on Cluster Computing and the Grid (CCGrid'03)*, Tokyo, Japan, May 2003, pp. 198–205.
- [12] K. Munagala, S. Babu, R. Motwani, and J. Widom, "The pipelined set cover problem," in *Proc. Intl. Conf. Database Engineering (ICDE'05)*, Tokyo, Japan, April 2005.
- [13] B. J. Bonfils and P. Bonnet, "Adaptive and decentralized operator placement for in-network query processing," in *Proc. Information Processing in Sensor Networks (IPSN'03)*, Palo Alto, CA, April 2003, pp. 47–62.
- [14] U. Srivastava, K. Munagala, and J. Widom, "Operator placement for in-network stream query processing," in *Proc. 24th ACM Symposium on Principles of Database Systems (PODS'05)*, Baltimore, MD, June 2005.
- [15] A. Bestavros, A. D. Bradley, A. J. Kfoury, and M. Ocean, "snBench: A development and run-time platform for the rapid deployment of video sensornet applications," in *Proc. of 2nd Intl. Workshop on Broadband Advanced Sensor Networks (Basenets'05)*, Boston, MA, October 2005.

- [16] P. B. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan, "Irisnet: An architecture for a world-wide sensor web," *IEEE Pervasive Computing*, vol. 2, no. 4, pp. 22–33, 2003. [Online]. Available: <http://www.intel-iris.net/papers/pervasive-03.pdf>
- [17] P. T. Uriel Feige, Laszlo Lovasz, "Approximating min-sum set cover," in *Proc. 5th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX 2002), Rome, Italy*, September 2002, pp. 94–107.
- [18] E. Korach, D. Rotem, and N. Santoro, "Distributed algorithms for finding centers and medians in networks," *ACM Trans. Program. Lang. Syst.*, vol. 6, no. 3, pp. 380–401, 1984.
- [19] S. C. Bruell, S. Ghosh, M. H. Karaata, and S. V.Pemmaraju, "Self-stabilizing algorithms for finding centers and medians of trees," *SIAM Journal on Computing*, vol. 29, no. 2, pp. 600–614, 1999.
- [20] J. Campbell, P. B. Gibbons, and S. Nath, "IrisNet: An internet-scale architecture for multimediasensors," in *Proc. of ACM Multimedia (MM'05), Singapore*, November 2005.
- [21] V. Vazirani, *Approximation Algorithms*. Springer, 2001.