Chapter 1

ENERGY-EFFICIENT DATA MULTICAST IN MULTI-HOP WIRELESS NETWORKS

Jennifer L. Wong, Giacamino Veltri and Miodrag Potkonjak
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA, US 90095

Abstract: ***Multi-hop wireless networks (MHWNs) are an emerging paradigm for bandwidth and energy-efficient wireless systems where each terminal communicates only with a few closely positioned neighbor nodes using low power communication schemes. High-rate (multimedia) data networks, sensor networks, and voice communications are seen as three main application domains of MHWNs. While MHWNs open many new research and economic opportunities, they simultaneously pose a number of new challenging technical problems. Among them, the fundamental role is reserved for energy- efficient data delivery.***

We address the problem of Data Multicast in Multi-Hop Wireless Networks that effectively captures requirements of all three of the scenarios. The problem focuses on how to minimize energy consumption while delivering data to all consumers that requested it. Since in the current and pending technologies communication dominates energy consumption, we aim to minimize the number of nodes that transmit data for the request data delivery problem instance. First, we formulate the problem and establish its computation complexity. Next, we develop an efficient heuristic which utilizes information about the geographical position of the nodes in the network to find the most energy-efficient communication path. Finally, we establish the effectiveness of the proposed approach by conducting comprehensive testing of our heuristic on typical ad-hoc networks and instances.

Key words: ***Wireless Ad-hoc Networks, Multi-Hop Networks, Sensor Networks, Data Dissemination.***

1. Introduction

In the last decade, the continuous high pace of technological advances has enabled the exponential growth of the Internet. We can trace the development of two implementation technologies as prime enablers of this growth. The first is the dramatic reduction in the cost of disks, or massive permanent storage. The second is the huge reduction in the cost of optical communication and simultaneous capacity increase. For example, the capacity of a hundred dollar disk increased by factor of 1,200 times in the last eleven years. At the same time, the bandwidth of optical cable has been doubling every nine months.

The Internet is a great educational, entertainment and economic resource which enables information to be available at the touch of a mouse. There is a wide consensus that the Internet will grow rapidly both in quantitative and qualitative terms. At the same time, it appears that we are on the brink of the next technological revolution that may have even higher impact. This revolution that will enable communication anytime, anywhere and a connection between physical and communication worlds is due to the advancement of wireless communication technology and sensors. For example, while in early 90's, wireless technology was mainly stagnant in the last six years before it started its exponential growth. With this advancement, there is currently a need for methodologies and technologies that will enable efficient and effective use of wireless network applications. The motivational factors pushing for these applications include the mobility of computational devices (e.g. cell phones and PDAs) and the ability to embed these devices into the physical world. These applications will make information easier to obtain and available at lower cost.

While the traditional wireless network architecture has been based on systems of static base stations, it appears that multi-hop network, where each node communicates with a few close nodes, is the most efficient in terms of energy saving and bandwidth reuse. In multi-hop networks each node communicates with other nodes that are geographically distant using intermediate nodes to build communication paths. The key constraint in this situation is energy.

Battery capacity and size limits the advancement and applications of these networks. In the last eleven years battery capacity has increased only by a factor of 2.7. Communication is the dominating energy consumption component in MHWN. As a result, the most effective way of energy saving is to power down all parts of then multi-hop network except those required for currently requested sensing and communication. Our goal in this work is to study how one can minimize energy consumption in multi-hop networks while satisfying the needs of all requested data transfers.

1.1 Motivation and Motivational Example

We illustrate the targeted problem using the following example in multi-hop networks. Recall that the network is multi-hop in the sense that every node cannot communicate with all other nodes in a single hop. As mentioned earlier, communication is the dominant cost, and therefore for each communication by a node in the network we assume a charge of ten units of energy. When a node is not communicating we assume a charge of one unit of energy. In the network, we have some number of data consumers and some number of producers. The consumers are programs, agents, humans or the Internet gateways which would like information from the network.

Producers are nodes in the network which detect events or have the stored information required to send to the consumers. An example of a producer could be a person in multimedia delivery who wants to broadcast the video to other users in the network. The goal is to transfer this information from the producers to the consumers in such a way that minimum energy is consumed in the network. Therefore, the goal is to find an energy-efficient way to use the minimum number of nodes when communicating.

There are at least two conceptually different and natural ways to abstract and specify the targeted problem: graph theoretic and geometrical. In graph theory problems, each of the nodes in the network corresponds to a node in the graph and is connected by an edge if communication is possible between the nodes. In many senses the graph theoretic approach is a proper and very efficient way to model MHWNs. However, there is at least one serious drawback. This approach allows for the construction of graphs which are not feasible in the geometrical or real life, sense.

In the geometrical specification of the problem, each node is placed in 2 or 3-D space and communication between nodes is configured the same way as done in graph theory. The main complication with geometrical representation is that unit distances between nodes in the network have, in general, no correlation with the number of hops required for the nodes to communicate. Two nodes can be positioned in close vicinity of each other, but the only way to communicate between the nodes is to travel some very long multi-hop path. For example this could be due to a obstacle, such as a building of mountain, obstructing their communication path. Because of the limitations of both graph theoretic and geometrical representation, we believe it is advantageous and necessary to consider both of them when approaching the problem.

We now state the targeted problem of Data Multicast in Multi-Hop Wireless Networks. Informally, the goal is to transmit information from a node to a set of specific nodes in such a way that the minimal number of nodes have to communicate. Consider the following example shown in Figure 1-1.

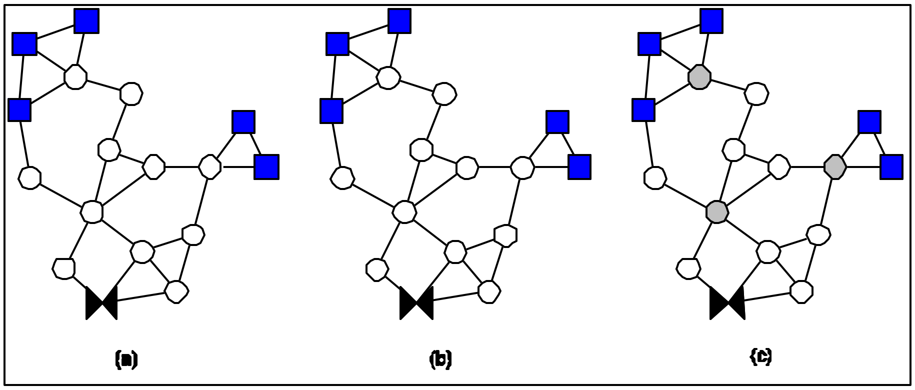

Figure 1-1. (a) Original Network. (b) Steiner Tree. (c) Path with Multicast considered.

In this example, the producers are denoted by black bowties. The consumers are denoted by squares and the circles indicate all other nodes in the network. In Figure 1-1(a), we see the given network. The goal is to connect the bowtie with the squares with the path with the minimal amount of communication.

One interesting observation is that the problem is similar, but not identical, to finding a Steiner tree. It is important to realize that a Steiner tree-based solution is not necessarily optimal. The Steiner tree problem consists of finding a minimum-weight tree connecting a designated set of vertices, called terminals, in a weighted graph or points in a space. The tree may include non-terminals, which are called Steiner vertices or Steiner points. In this case, we can assume that every edge in the network is of equal cost and that Steiner vertices and points are not allowed. If this is the case, then a possible Steiner tree is shown in Figure 1-1(b).

However in MHWNs the problem is slightly different. The key difference is the ability to utilize the benefits of multicasting in MHWNs. Broadcasting allows the communication of data to multiple nodes without additional cost. Therefore, the cost of a single node sending information to one of its neighbors is exactly the same as if it sent the information to multiple or all of its neighbors. Figure 1-1(c) shows a minimal path which makes use of multicasting. In the Steiner tree, 10 communications are necessary to send the information from the source to the destinations. However, when multicasting is efficiently utilized, only 8 nodes are needed to communicate. Each of the gray nodes in Figure 1-1(c) denote nodes which communicate only once, but to multiple nodes, therefore their outgoing edges are only counted as a single edge.

While one can easily solve the motivational problem by implicit enumeration, we will prove that the targeted problem is NP-complete. Note that the Steiner tree problem is also NP-complete. Therefore, the algorithmic goal of this work is to introduce a heuristic approach to the problem of Data Multicast in Multi-Hop Wireless Networks which makes use of multicasting and short path communications to minimize the communications in the network, and therefore save energy.

1.2 Objectives

This work was driven by a number of objectives.

- **To find and formulate an interesting, important and frequently occurring problem in MHWN**. We have formulated the problem of Data Multicast in Multi-Hop Wireless Networks. This problem has an interesting interpretation in a number of different types of Multi-Hop Wireless Networks. For example, the problem is a key component for applications such as event tracking and multimedia data delivery.
- **Establish the complexity of the problem.** We prove the complexity of the Data Multicast in Multi-Hop Wireless Networks problem to be NP-Complete by component design.
- **Develop an efficient heuristic for the problem** We developed a heuristic technique which leverages on the geographical position of the nodes in the network to determine the most energy-efficient multicast path.
- **Develop efficient testing techniques and instances for the problem.** Since the problem is NP-complete and there are no established benchmarks, it is difficult to directly evaluate the quality of obtained solutions.

1.3 Organization

We organize the remainder of the work in the following way. In the next section, we survey the related work. After introducing the related preliminary material in Section 3, we formally define the problem, establish its complexity and present the new heuristic algorithm. Finally, before concluding, we present comprehensive experimental results.

2. Related Work

The related work can be traced along three lines of research: ad-hoc wireless networks, data dissemination in multi-hop wireless networks, and the design and use of low power distributed systems.

2.1 Multi-Hop Wireless Networks

Cellular local area networks have emerged to become a dominating architecture entity in wireless communication. Recent technology has lead to the development of wireless ad-hoc networks, a viable alternative with the potential to reduce deployment costs and increase energy and bandwidth use efficiency. Varieties of high impact applications have been envisioned for this type of network which have the potential to impact our daily lives [EST00,LIC60,WEI93].

Multi-hop sensor networks pose a need for new solutions to a number of design and efficient usage problems [KOU02, SLI01, WON02]. For example, there is a need to address problems such as new signal processing techniques [TEN96,POT00,YU99], operating systems with features such as low power design [DON97], robotics [SUK00], coverage [MEG01b], and quality of service [MEG01].

2.2 Data Dissemination in Multi-Hop Wireless Networks

Broadcasting information from a producer to a set of consumers is a problem that has received a great deal of attention in a number of computer science fields. For example, [ACH96] describes how Broadcast Disks methodology can be efficiently used to provide data to a set of consumers using satellite broadcast in such a way that each consumer waits least for his data. Therefore, the main goal is to organize broadcast in such a way that data in higher demand is more often broadcast. Dan et al. [DAN94] and many other research addressed questions related to efficient multimedia data delivery from disks. By far the most popular and widely used scheme for broadcasting data is Internet multicast [DEE90,BIR91]. There is very little similarity between that and our problem, due to very different set of constraints and objective function. Recently, data delivery and aggregation attracted the attention of theoretical computer science, and in particular the approximation algorithms community [GUH00,GUH00a,PRA91].

The work presented by Intanagonwiwat et. al in [INT01] shares some similarities with our work. For example, the goal of [Int01], like ours, is to reduce the amount of energy it takes for producers to communicate with consumers, and even more specifically, it attempts to minimize the number of communications required by using a Greedy Incremental Tree. However, the difference mainly lies in the formulation of the problem. In the work, the idea is to send data from multiple producers (referred to as sources) to a particular consumer (referred to as a sink). Along the way, messages from the producers can be combined (aggregated) to form a more compact message to send along to the consumer. Thus, the work deals with minimizing the size of the message received by the consumer. In our work, messages are generated by one producer and distributed to many consumers. Thus, the message sent by nodes will be the same, and therefore our goal is not to minimize the size of the message, but to minimize the amount of communications needed to reach all consumers.

2.3 Low Power and Distributed Systems

Energy consumption is one of the main concerns in wireless ad-hoc networks. Low power techniques are one of the key issues which need to be addressed at all levels of development. A number of different works have focused on minimizing power consumption at the routing level [SIN98,CHA99], and also for signal processing of mobile nodes [DON97,RAB00].

Distributed systems have developed over the last two decades, beginning with techniques such as interprocess communications and remote invocation, distributed file systems, and data replications [COU01,TAN95]. In the future, distributed systems will expand to support mobility of users over wireless and ad-hoc networks. Special types of innovative distributed systems are the World Wide Web and the emerging field of wireless ad-hoc networks. The design of distributed systems is also often addressed from the synthesis point of view in the design automation community [PRA91, YEN95].

3. Preliminaries

In order to make the paper self-contained, in this section we summarize all the main assumptions.

We assume the following MHWN architecture throughout the rest of the paper. For each node in the network, we assume a finite, standard communication transmitting and listening range. We say that two nodes are neighbors if they reside within their communication range. Communication in the network is assumed to be triggered by sensing or a request by one or more nodes in the network. These requests are much smaller than actual sent data and therefore are not considered during the optimization process.

Communication in current technologies is the dominant cost in terms of energy consumption in wireless networks. In today's technology, listening is often as expensive as communication. However, in the future this may not be the case. For the sake of this work, we assume listening for communication from other nodes is equally as costly as communicating information. According to these assumptions, we define two modes for each of the nodes, **on** and **off**. Other models which do not assume equality between listening and communicating may include a third mode which is **standby**, or a listening state.

We define the **on** mode to include when a node is communicating information and when it is listening or receiving information from another node. In this mode, the node is consuming a dominant amount of energy. The **off** mode is considered to be the minimum state in which the node consumes an absolute minimum amount of energy possible.

We recognize the importance of localized algorithms in MHWNs, and the development of a localized approach is a future goal. However in this case, we assume a centralized system. We see this work as an important and natural starting point for the development of a localized approach. Currently, we see two planes of communication in MHWNs: control and data. The control plane is the operating system of the network, where the data plane is the information and data being passed to various places in the network. The control plane is orders of magnitude smaller than the data plane, and therefore will require negligible communication costs. As a result of this, a centralized approach for communicating **on** and **off** states to each node in the network is a reasonable approach.

Our model of MHWNs is expanded to include two special types of nodes, producers and consumers. Producers are nodes in the network which have relevant information which needs to be passed on. Consumers are nodes which need the information which the producers have. Applications such as information dissemination, information aggregation, and event/object tracking, to name a few, all contain the notion of producers and consumers.

An additional assumption in our model is that the geographical position of each node is known according to a defined origin. A number of location discovery techniques have been developed [MEG01c,SAV01].

4. Design Approach

In this section we first informally and formally define the addressed problem. Next, we establish the computational complexity of the problem. Finally, we propose a new heuristic algorithm that leverages on both graph theoretic and geometric information to find an energy-efficient solution to the data multicast problem.

4.1 Data Multicast in Multi-Hop Networks

Informally, the problem of Data Multicast in Multi-Hop Networks can be defined as follows. The goal is to select the minimum number of nodes in the network to be **on**, such that there is a path of communication between the producer and each of the consumers. Intuitively, by selecting the minimum number of nodes we are maximizing the number of nodes which can be in the **off** mode, and therefore save the maximum amount of energy.

We define the Data Multicast in Multi-Hop Networks Problem formally using the Garey-Johnson format [GAR79].

Problem: *Data Multicast in Multi-Hop Networks*
Instance: Graph G =3D (V,E), subset {P: p_1, p_2, … , p_l] ? V, subset {C: c_1, c_2, … , c_m} ? V, positive integer K_2 ? |V|.
Question: Is there a subset V' ? V with |V'| ? K_2 where C ? V', at least one P_i ? V', and there exists an S, a sequence <s_1, s_2, … , s_n> where {s_i, s_{i+1}} ? E, s_1 ? P, and s_n ? C, for every C to at least one P?

4.2 Complexity

To justify our heuristical approach to the Data Multicast in Multi-Hop Networks problem, we prove using Karp's polynomial transform component design techniques that the problem is NP-complete. Specifically, we map the Minimum Cover problem to the Data Multicast in Multi-Hop Networks. For the sake of completeness, the formal definition of the Minimum Cover Problem is as follows:

Problem: *Minimum Cover*
Instance: Collection T of subsets of a finite set S, positive integer K_1 ? |C|.
Question: Does T contain a cover for S of size K_1 or less, i.e. a subset T' ? T with |T'| ? K_1, s.t. every element of S belongs to at least one member of T'?

Proof: We show that one can transform, in polynomial time, an arbitrary instance of Minimum Cover Problem to the Data Multicast in Multi-Hop Networks problem. Given an instance of Minimum Cover consisting of j subsets T_1, T_2, … , T_x, we construct a graph G =(V,E) such that C=S, |P|=1, and |V|= x+|S|+1. Consider, for example, the following Minimum Cover instance:

$$T = \{\{e_1, e_3, e_5\},\{e_2, e_3, e_4\}\{e_1, e_4\}\{e_3, e_4, e_5\}\}$$
$$S = \{e_1, e_2, e_3, e_4, e_5\}$$

We map the instance to the problem of Data Multicast in Multi-Hop Networks in the following way. We create a single vertex, P_1, which corresponds to the collection T. For every element in S, we create a vertex, c_i, and include each vertex in the subset C. Each subset in T represents a single vertex, v_i. The graph G, for the example given above is shown in Figure 1-2. We place P_1 at the top of the graph and create edges from P_1 to each of the *j* subsets of T, or vertices v_i. In the final row, each element of S is shown, represented by c_i. Each of the subset nodes, v_i, connects to the corresponding elements in which it contains (c_i vertices).

If a solution to the Data Multicast in Multi-Hop Networks problem can be found where K_2 is equal to some K, then a solution to the Minimum Cover problem can be found with K_1 = K-|C|-1.

Consequently, we have shown that the Minimum Cover problem polynomially transforms to the problem of Data Multicast in Multi-Hop Networks. It is easy to see that the Data Multicast problem is NP, because one can easily check using breadth-first search and enumerate to see if the proposed solution is correct. Therefore, the Data Multicast in Multi-Hop Networks is a NP-complete problem.

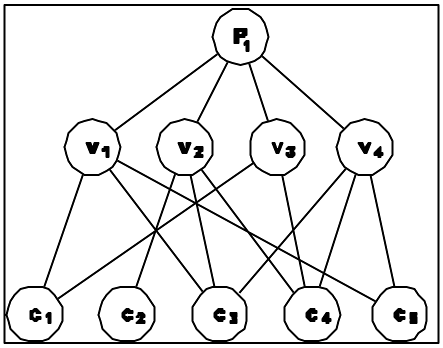

Figure 1-2. Map of Minimum Cover Problem to Data Multicast in Multi-Hop Networks

4.3 Line-directed Node Selection Heuristic

Our heuristic approach to the Data Multicast in Multi-Hop Networks problem considers a network with one producer and *n* consumers. It can be easily generalized by applying the same procedure to each producer in the general case.

Table 1-1. Line-Directed Node Selection Heuristic.

Input: Wireless Network, S, with a single producer and *n* consumers
Output: On/Off state for each node in the network.
Algorithm:

```
for each node i in S
    i.Mode = Off;
Turn-On-Node(Producer);
while (Consumers are Off)
    for each node i in S
        N = 0, Sum = 0, NextNode = NULL, Min = INFINITY;
        If (i.Mode = Influenced)
            for each node C that is a consumer in S
                D = NodeLineDistance(i, C,Producer);
                If (D ?0)
                    Sum = D + Distance(i,C);
                    N = N + 1;
            Sum = Sum / N;
            If (Sum ? Min) and (Sum ? 0)
                NextNode = i;
                Min = Sum;
    Turn-On-Node(NextNode);
MST-MaxLeaf(S);
PruneMST(S);
```

The algorithm consists of three main ideas. The first is to use line-directed information to select nodes which will communicate between the producer and consumers. The second is to construct the minimal spanning tree, from all the selected nodes, with the largest number of leaves. The leaves represent selected nodes which are not necessary in order for the consumers and the producer to communicate. The last idea is to prune from the tree all the leaves which are not consumers or the producer.

The line-directed algorithm calculates what we call a sphere of influence. For a given node, *n*, all of its neighbors are included in the nodes sphere of influence. If *n* communicates to any node, all of the other neighbors can detect this communication, assuming the nodes are **on** or in the **sensing** mode. We then consider the neighbors of a node which we know to be on, to be influenced by that node, because communication to them is obtained at no additional cost.

For each consumer in the graph, we introduce a line between it and the producer, which we will call guidelines. These guidelines never change, and are used to determine the path from the producer to the consumers.

We begin the line-directed algorithm by placing all the nodes in the network in the **off** mode except for the producer. We continually turn nodes **on** which are selected to be on the path to the consumers, until all the consumers are **on** and each have a path to the producer.

In order for a node to be turned **on**, and form a path to the consumer, we determine the sphere of influence of the current nodes which are in the **on** mode. For each of the nodes in the sphere of influence, we sum In order for a node to be turned **on**, and form a path to the consumer, we determine the sphere of influence of the current nodes which are in the **on** mode. For each of the nodes in the sphere of influence, we sum up the perpendicular distance from the node to each of the guidelines and the distance from the node to the consumer. Note that once a consumer is **on** and connected to the producer, we remove its guideline from future consideration. Also, if the perpendicular intersection point of the node and the guideline is not on the line segment between the producer and the consumer, we disregard it. We normalize the calculated sum of each of the nodes in the sphere of influence and select the node with the minimum value. The selected node is turned **on** and each of the node's neighbors is included in the sphere of influence. If any of the consumers are neighbors of the node, they are also placed in the **on** mode.

It is often the case, that the best node at a single iteration of the algorithm is not the best node to use when constructing communication paths with the minimum number of nodes. As a result, some nodes maybe unnecessarily **on** in the network. Therefore, we find a minimal spanning tree on the final state to help eliminate these nodes.

Once the algorithm has found some path(s) from the producer to all the consumers, we build the minimal spanning tree with the maximal number of leaf nodes (MST-MaxLeaf) from all the nodes which are **on**. Each node which is a leaf node can later be pruned under the assumption that it is not needed on the minimal path between the producer and consumers. We use a greedy algorithm which begins with the producer and builds outwards. At each step, the MST-MaxLeaf algorithm examines all the neighbors of all the leaf nodes in the tree and selects the neighbor which has the greatest number of neighbors not in the tree.

Once we have constructed the minimal spanning tree with the maximum number of leaf nodes, we then iteratively prune or turn **off** all nodes which are leaves in the tree at the current iteration.

For example, consider the MHWN shown in Figure 1-3. We represent the network as a grid, where each intersection represents a node, and each edge corresponds to a communication edge. A grid design of the network is only used for simplification and example purposes. The line-directed algorithm can be applied to an arbitrary MHWN.

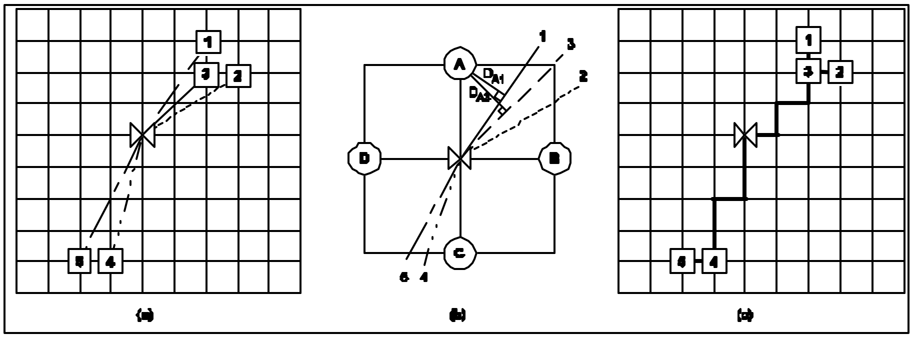

Figure 1-3. (a) Example of the line-distance algorithm. (b) Example of first iteration. (c) Final solution found by the line-distance algorithm.

In this example, we have 5 consumers (squares) and 1 producer (bow tie). In Figure 1-3(a) the network is shown along with the guidelines for each of the consumers in the network. According to the line-distance algorithm, the producer is **on** and the four nodes, A, B, C, and D, which can communicate with the producer are in the sphere of influence (shown in Figure 1-3(b)). For A, B, C, and D, we determine perpendicular distance for each of the nodes to the guidelines and their distance to each of the consumers. Recall, that if the perpendicular distance to the guideline does not fall on the line segment between the consumer and the producer, then the measurements are not taken into account. For example, for nodes A and B, we calculate the perpendicular distance and physical distance to consumers 1, 2, and 3. When calculating the sum for nodes C and D, we only consider consumers 4 and 5. The node with the minimal normalized sum will be turned **on** and its neighbors will be added to the sphere of influence, in this case, node B is turned **on**. In the next iteration, calculations will be made for each of B's neighbors. Once paths to each of the consumers have been established, the MST-MaxLeaf is established and then iteratively pruned. In this case, no additional nodes were **on** therefore there was no need to find the MST-MaxLeaf or to prune. The minimal communication path to all consumers is shown in Figure 1-3(c).

5. Experimental Results

The proper evaluation of the effectiveness of the developed heuristic for the Data Multicast problem poses some technical and logistic problems.

Technical, because the problem is NP-complete and therefore, in general we do not know how to find the optimal solution. Logistic, because the problem is not previously studied and therefore, we cannot compare the new algorithm against the previous on standard benchmarks.

In this situation, there are two sound, but conceptually different ways to evaluate an algorithm. The first involves creating instances of the problem which the optimal solution is known, yet the solution is not obvious and does not favor any one particular algorithm (especially the algorithm in question). Applying this technique to ours problem means that before we can determine the effectiveness of our algorithm, we first have to create instances of our problem in which the optimal solution is known. Creating these instances leads to several difficulties in particular when one tries to generate a structurally diverse set of examples.

Therefore, we have taken a different approach to evaluate the new algorithm - the minimum sharp bound. The idea is to calculate bounds in such a way that the optimal solution cannot be better and yet the bounds are as close as possible to the optimal solution. To accomplish this goal, we evaluate two separate statistics. Our problem is to determine the minimum number of communications such that all consumers can receive the information from the producer. So, the first statistic used to evaluate our algorithm is calculated by determining a lower bound on the number of communications it would take for the producer to talk to each consumer separately. Then, we determine the number of communications it would take our solution to talk to each consumer separately. We determine the excess number of communications used by our algorithm and express it as a percentage greater than the minimum calculated previously, and finally, we take the average over all the paths. To calculate the minimum number of communications required for a producer to communicate with a consumer, we simply calculate the breadth-first-search shortest path from the producer to the consumer. To determine how many communications are required for the producer to communicate with the consumer in our solution, we calculate a breadth-first-search shortest path using only the nodes that we calculated to be **on**. Thus, the closer our solution is to the breadth-first-search shortest path, the closer our solution is to optimal. However, it is important to note that in many cases, due to the branching of communications, the breadth-first-search shortest path from the producer to a consumer is not the optimal solution, and so, having more communications than the breadth-first-search shortest path requires does not indicate that our solution is poor.

The second statistic presented is the percentage of sensors that need to be on and are not consumers; obviously, consumers must be on if they are to hear the message communicated by the producer. To calculate this second metric, we simply determine the number of sensors turned on by our algorithm and divide by the total number of sensors. Once again, it is important to note that turning on the minimum number of sensors does not necessarily guarantee the minimum number of communications, and again, this is due to the branching of communications.

Table 1-2. Selected Experimental Results for the Line-Directed Algorithm.

# of Sensors	**# of Consumers**	**Communication Radius**	**Ave. # Sensors on per Path**	**Ave. Excess Communications**
200	10	0.15	11	25.2%
400	100	0.25	25.25	3.8%
600	210	0.075	43.5	19.8%
800	200	0.125	25.125	16%
1000	250	0.075	29	30.4%

We ran several test cases using our algorithm with the number of sensors and consumers, and sensor communication radius as the parameters. The number of sensors varied from 200 to 1000 in increments of 200, the percentage of consumers varied from 10 to 250, and the sensor communication radius varied from 0.05 to 0.3 (assuming that all sensors were located in a 1x1 square). Table 1-2 displays a subset of the results. It is apparent that while the quality of the solution varies relative to the lower bounds, it is consistently competitive. In Figure 1-4 and 1-5 we show two examples of networks with ten consumers and 300 nodes randomly placed, with uniform communication radius of 0.10.

6. Conclusion

We addressed the problem of efficient Data Multicast in Multi-Hop Wireless Networks. The goal is to minimize the energy consumption while delivering data to all consumers in a multi-hop wireless network that requested it. We formulate the problem and proved that the problem is NP-complete. The technical highlight of the paper is an efficient heuristic that utilizes information about the geographical position of the nodes in the network to find the energy-efficient communication path. The experimental results indicate that the new algorithm produces solutions that are close to sharp lower bounds.

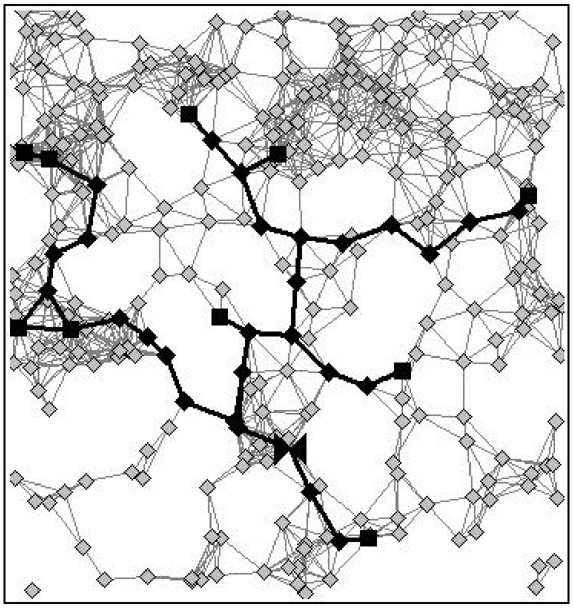

Figure 1-6. A Multi-Hop Wireless Network of 300 Sensors with Communication Radius of 0.1 and 10 Consumers.

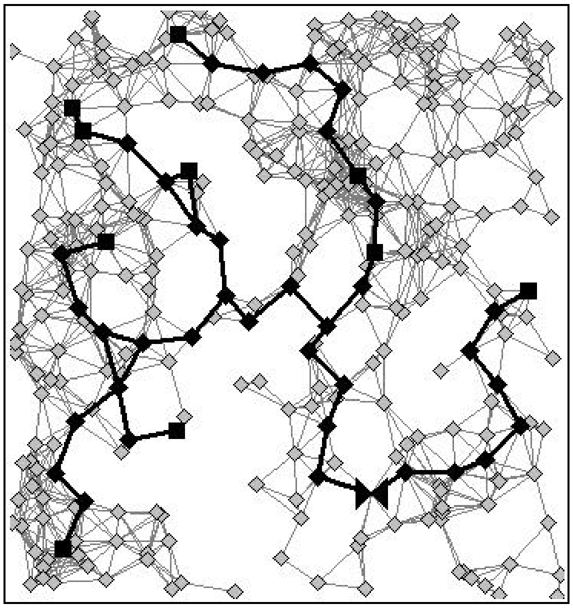

Figure 1-7. A Multi-Hop Wireless Network of 300 Sensors with Communication Radius of 0.1 and 10 Consumers.

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. ANI-0085773. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

References

[ACH96] S. Acharya, M. J. Franklin, S. B. Zdonik, “Prefetching from Broadcast Disks”, ICDE, Feb. 1996, pp. 276-285.

[BAR98] A. Bar-Noy, S. Guha, J. Naor, B. Schieber, “Message Multicasting in Heterogeneous Networks”, ACM Symposium on Theory of Computing (STOC), 1998, pp. 448-453.

[BIR91] K. Birman, A. Schiper, P. Stephenson, “Lightweight Causal and Atomic Group Multicast”, ACM Transactions on Computer Systems, Aug. 1991, Vol. 9, n. 3, pp. 272-314.

[CHA99] J.H. Chang, L. Tassiulas, “Routing for Maximum System Lifetime in Wireless Ad-hoc Networks”, Allerton Conference on Communication, Control, and Computing, Sept. 1999.

[COU01] G. Coulouris, J. Dollimore, T. Kindberg, “Distributed Systems: Concepts and Design”, Harlow, England, New York, Addison-Wesley, 2001.

[DAN94] A. Dan, D. Sitaram, P. Shahabuddin, “Scheduling Policies for an On-Demand Video Server with Batching”, ACM Multimedia, Oct. 1994, pp. 391-398.

[DEE90] S. Deering, D. Cheriton, “Multicast Routing in Datagram Inter-networks and Extended Lans”, ACM Transactions on Computer Systems, May 1990, Vol. 8, pp. 85-110.

[DON97] M.J. Dong, K.G. Yung, W.J. Kaiser, “Low Power Signal Processing Architectures for Network Microsensors”, International Symposium on Low Power Electronics and Design ACM, Aug. 1997, pp. 173-177.

[EST00] D. Estrin, R. Govindan, J. Heidemann, “Embedding the Internet: Introduction”, Communications of the ACM, May 2000, Vol. 43, n. 5, pp. 38-42.

[GAR79] M.R. Garey, D. S. Johnson. “Computers and Intractability: A Guide to the Theory of NP-completeness, W.H. Freeman and Company, New York, 1979.

[GUH00] S. Guha, N. Mishra, R. Motwani, L. O'Callaghan, “Clustering Data Streams”, Symposium on Foundations of Computer Science, Nov. 2000, pp. 359-66.

[GUH00a] S. Guha, A. Meyerson, K. Munagala, "Hierarchical Placement and Network Design Problems", IEEE Symposium on Foundations of Computer Science, 2000, pp. 603-612.

[INT01] C. Intanagonwiwat, D. Estrin, R. Govindan, J. Heidemann, "Impact of Network Density on Data Aggregation in Wireless Sensor Networks", International Conference on Distributed Computing Systems, Nov. 2001, pp. 1-16.

[KOU02] F. Koushanfar, M. Potkonjak, A. Sangiovanni-Vincentelli, "Fault Tolerance in Wireless Ad Hoc Sensor Networks", IEEE Sensors, June 2002.

[LIC60] J. C. R. Licklider, "Man-Computer Symbiosis", IRE Transactions on Human Factors in Electronics, Vol. 1, March 1960, pp. 4-11.

[MEG01] S. Meguerdichian, F. Koushanfar, G. Qu, M. Potkonjak, "Exposure In Wireless Ad Hoc Sensor Networks", International Conference on Mobile Computing and Networking, July 2001, pp. 139-150.

[MEG01b] S. Meguerdichian, F. Koushanfar, M. Potkonjak, M. Srivastava, "Coverage Problems in Wireless Ad-Hoc Sensor Networks," IEEE Infocom, April 2001, Vol. 3, pp. 1380-1387.

[MEG01c] S. Meguerdichian, S. Slijepcevic, V. Karayan, M. Potkonjak, "Localized Algorithms In Wireless Ad-Hoc Networks: Location Discovery and Sensor Exposure", ACM Symposium on Mobile Ad Hoc Networking and Computing, Oct. 2001, pp. 106-116.

[POT00] G. J. Pottie, W. J. Kaiser, "Wireless Integrated Network Sensors", Communications of the ACM, 2000, Vol. 43, n. 5, pp. 51-58.

[PRA91] S. Prakash, A. Parker, "Synthesis of Application-specific Multiprocessor Architectures", IEEE/ACM Design Automation Conference, June 1991, pp. 8-13.

[RAB00] J. M. Rabaey, M. Potkonjak, F. Koushanfar, S. Li, T. Tuan, "Challenges and Opportunities in Broadband and Wireless Communication Designs", IEEE/ACM International Conference on Computer Aided Design, Nov. 2000, pp. 76-82.

[SAV01] A. Savvides, C. C. Han, M. B. Srivastava, "Dynamic Fine-Grained Localization in Ad-Hoc Wireless Sensor Networks", International Conference on Mobile Computing and Networking, July 2001, pp. 166-179.

[SIN98] S. Singh, M. Woo, C.S. Raghavendra, "Power-aware Routing in Mobile Ad-hoc Networks", ACM/IEEE International Conference on Mobile Computing and Networking, June 1998, pp. 181–190.

[SLI01] S. Slijepcevic, M. Potkonjak, "Power Efficient Organization of Wireless Sensor Networks", IEEE Intl. Conf. on Communications, June 2001, Vol. 2, pp. 472-476.

[SUK00] G. S. Sukhatme, M. J. Mataric, "Embedding Robots into the Internet", Communications of the ACM, May 2000, Vol. 43, n. 5, pp. 67-73.

[TAN95] A. Tanenbaum, “Distributed Operating Systems”, Prentice Hall, Englewood Cliffs, N.J., 1995.

[TEN96] D. L. Tennenhouse, V. G. Bose, “The SpectrumWare Approach to Wireless Signal Processing”, Wireless Networks, 1996, Vol. 2, n. 1, pp. 1-12.

[ROD98] V. Rodoplu, T. H. Meng, “Minimum Energy Mobile Wireless Networks”, IEEE International Conference on Communications, June 1998, Vol. 3, pp. 1633–1639.

[WEI92] M. Weiser, “The Computer for the 21st Century”, Scientific American (International Edition), Sept. 1991, Vol. 265, n. 3, pp. 66-75.

[WEI93] M. Weiser, “Some Computer Science Issues in Ubiquitous Computing”, Communications of the ACM, July 1993, Vol. 36, n. 7, pp. 74-84.

[WON02] J. L. Wong, and M. Potkonjak, “Search in Sensor Networks: Challenges, Techniques, and Applications”, International Conference on Acoustics Speech and Signal Processing, May 2002.

[YEN95] T.-Y. Yen, W. Wolf, “Communication Synthesis for Distributed Embedded Systems”, International Conference on Computer-Aided Design, Oct. 1995, pp. 288-294.

[YU99] T. Yu, D. Chen, G. Pottie, K. Yao, “Blind Decorrelation and Deconvolution Algorithm for Multiple-input Multiple-output System”, The International Society for Optical Engineering, July 1999, Vol. 3807, pp. 200-209.

Index