Stochastic Routing in Ad Hoc Wireless Networks

Christopher G. Lott and Demosthenis Teneketzis [1]
Department of EECS, The University of Michigan
1301 Beal Ave, Ann Arbor, MI 48109-2122
clott@eecs.umich.edu, teneketzis@eecs.umich.edu

Abstract

We investigate a network routing problem where a probabilistic local broadcast model for wireless transmission is used. We present results showing that an index policy is optimal for this problem. We extend the original model to allow for power control, and assert that the index nature of the optimal routing policy remains unchanged. We further allow time-varying system parameters in the original model, and discover conditions under which a time-varying index routing policy is optimal. Finally, we present a distributed implementation of the routing policy and provide results on its convergence properties.

1 Introduction

As communication networks become ever more prevalant, research to improve their design has become both more encompassing and more application specific. Protocol (policy) issues such as service priority, retransmission schemes, flow control, and routing are now studied for large distributed systems and for a variety of communication channel types. Further, it has long been recognized that integrating various source types, such as voice, video, email, web access, and business transaction data into one seamless transparent network results in both the highest efficiency and ease of use. In such networks the ability to provide appropriate Quality of Service (QoS) to each service type is of paramount importance. A further challenge is posed by wireless networks, where these service goals must be achieved with unreliable and time-varying channels, and where new concerns, such as energy consumption and channel interference, impose additional constraints.

The term *ad hoc* is often now applied to networks in which there is no central network controller, each node can itself act as a store-and-forward router, and in which the connectivity topology is time-varying (e.g. see [7]). Such a network is in contrast to, for example, a cellular network where each cell has a central base station through which all data is transmitted. We think of the network as having a number of messages, each of which is destined for some set of destinations. The general routing problem in an ad hoc network is to define a policy which, given the trajectory histories of all the messages, chooses which nodes should next transmit which message. It is also desired that this policy be implementable in a distributed fashion in the network, so the transmission decisions can be decided locally without knowledge of other parts of the network.

Many approaches can be taken to network routing optimization. A typical one is maximizing the overall throughput the network achieves. But in many cases for the wireless environment other considerations are at least as important. For example, often in a wireless network the energy source is locally stored in a battery at each node, and the major design goal is to achieve satisfactory communication while using up as little energy as possible. Also, there may be QoS issues, such as delivery timeliness and message priority, which affect the optimal service policy.

This paper briefly explores these design issues in network routing algorithms for ad hoc wireless networks, and provides a novel system model which allows for optimal design in a number of realistic situations. It is organized as follows. In the remainder of Section 1 we briefly discuss the literature available on routing in ad hoc wireless networks, and present a summary of the contributions of our paper. Section 2 contains a qualitative description of our transmission model. In Section 3 we present the basic time-invariant routing problem. In Section 4 we discuss the time-invariant routing problem with power control. In Section 5 we present the time-varying routing problem. In Section 6 we introduce a distributed algorithm for the time-invariant routing problem. Conclusions are discussed in Section 7. A full presentation of the topics discussed in this paper, including literature survey and references, model description, and the detailed analysis of stochastic routing for various models, can be found in [10].

1.1 Ad Hoc Wireless Network Routing Literature

Much research work on routing for ad hoc wireless networks has been published in recent years. Most of this

[1] This research was supported in part by ARO Grant DAAH04-96-1-0377, AFOSR Grants F49620-96-1-0028 and F49620-98-1-0370, and NSF Grant ECS-9979347.

work includes simulation results, comparing the proposed protocol with a generic link-state or distance vector protocol, or comparing specific algorithms against each other (a recent example is [3]). A minority of the published papers include analytic results, such as proofs of convergence or convergence rates. The work on routing for ad hoc wireless networks can perhaps most conveniently be categorized as either on-demand or route maintenance.

- *On-Demand Protocols*
  On-demand protocols are the algorithms which only update routing tables when a new message arrives. Examples include [4], [5], [8], [11], [13], [14]. and [15].

- *Route Maintenance Protocols*
  Route maintenance protocols periodically update topology information in the network independent of traffic flow. Examples include [2], [6], [12], [14]. [16], [17], and [18].

1.2 Contribution

We present what is, to the best of our knowledge, the first network routing protocol which uses a probabilistic local broadcast model for wireless transmission. This model allows for routing decisions to be made based on immediate feedback from actual transmissions. In contrast to most of the existing literature, where link-level communication is point-to-point, we fully utilize the result of each local broadcast and use the information provided by the immediate feedback to construct an optimal routing protocol. Such a protocol is well-suited for a distributed implementation, as it is based only on information local to a node. We present a method describing how the protocol is updated when topological changes occur in a distributed mobile network, and this demonstrates the protocol's practicality.

We extend the model to allow for transmission power control at each node, and show how the fundamental nature of this optimal protocol doesn't change. For this case, an optimal protocol effectively resolves the tradeoffs between fewer long hops vs. more short hops.

We also analyze a time-varying network, where reward for message delivery depends on time, and present conditions under which we can specify an optimal protocol. This model captures the important tradeoff between QoS requirements, such as delivery timeliness, and other criteria, such as energy efficiency.

2 Model Description

We consider the transmission of one message from a source to a destination. As is standard in many mobile radio environments, we assume each mobile has an omnidirectional antenna. Multiple neighboring nodes can in parallel receive a single transmission. We assume neighboring nodes can exchange control messages at minimal cost, implementing the transmission protocol.

The communication channel is degraded by numerous effects, such as blockage, shadowing, sky and receiver noise, and signal interference. Upon transmission failure, there is a rate of message retransmission, which is a function of message length, transmission rate, and protocol.

In our model we assume that the channel degradations can be lumped into just two categories: 1) "Fast" effects. Given what the controller knows, these effects are independent from one transmission to the next. Examples include shadowing, sky and receiver noise, and interference from other mobiles (assumed unpredictable). 2) "Slow" effects. These remain fixed over the entire time of transmission of a message to its final destination. This includes blockage and signal interference with long-term presence in the environment.

We incorporate these channel degradation effects into a probabilistic transmission model. The above assumptions then mean that transmission events are independent from one time to the next. However, reception at separate nodes can be correlated. These points are made explicit in the model described in the next section.

3 The Time-Invariant Stochastic Routing Problem

We begin by briefly defining notation and definitions for the system model under consideration, which we refer to as Model **(M)**.

N is the number of nodes in the network.

$\Omega = \{1, \ldots, N\}$, the set of all nodes. So $|\Omega| = N$.

$S \subseteq \Omega$ refers to a state of the system, defined as the set of nodes which have received the message.

$R : 2^{\Omega} \to \Re$ is the reward function, and $R_i := R(\{i\})$.

π is a Markov policy. We write $\pi(S) = i$ to indicate policy π transmits at node i when in state S. We write $\pi(S) = r$ to indicate policy π retires and receives reward $R(S)$ when in state S.

$V^{\pi}(S)$ is the expected reward when starting in state S under policy π. We often write V_i^{π} for $V^{\pi}(\{i\})$.

We write $P^i(S'|S)$ to indicate the probability of reaching state S' from state S when choosing i for transmission, $i \in S$.

We let $P_{ij} := \sum_{S:i,j\in S} P^i(S|i)$. j is called a *neighbor* of i if $P_{ij} > 0$. $\mathcal{N}(i)$ is the set of all neighbors of i, together with i itself.

Definition 3.1 *A function $f : 2^\Omega \to \Re$ is an* index function *on Ω if f satisfies*

$$f(S) = \max_{i\in S} f(\{i\}) \quad \forall S \subseteq \Omega \tag{1}$$

We begin with a centralized version of the stochastic routing problem, formulated as follows.

Problem $(\mathbf{P_1})$

We consider the transmission of a single message, from a given initial state S_o (i.e. a given set of nodes) to a set of destination states, in a wireless ad hoc network of N nodes described by Model **(M)**. Transmission instances occur at discrete time points. Each transmission from a given node i incurs a fixed energy cost c_i. At each transmission instance the transmitting node is chosen by a central controller that always knows the current state of the system (i.e. the set of nodes that have the message). Node transmissions are local broadcasts, that is, multiple neighbor nodes may all simultaneously receive the message. Given the node chosen to transmit, the probability that a given set of nodes receives the message is known and fixed. The central controller is informed, without any cost, as to which nodes receive the message. Nodes never forget the message after they receive it until the end of the transmission process. Each transmission event is assumed independent of those before and after. Control information flow between the nodes and the controller is considered free in energy and instantaneous in time. A reward function R is specified, where R must be an index function (see Def. 3.1). At any instance, the central controller can terminate the transmission process or choose to continue transmitting. The objective is to choose: (i) the node to transmit at each transmission instance, and (ii) the instance to terminate the transmission process, to maximize

$$E\left\{R(S_f) - \sum_{t=1}^{\tau} c_{i(t)}\right\} \tag{2}$$

where τ is the time when the transmission process is terminated, S_f is the state at τ, and $i(t)$ is the node chosen by the transmission policy at time t.

3.1 Analysis of Problem $(\mathbf{P_1})$

For the analysis of Problem $(\mathbf{P_1})$, we can restrict attention to Markov policies without any loss of optimality. We need the following definitions.

Definition 3.2 *A Markov policy π is a* priority *policy if there is a strict priority ordering of the nodes s.t. $\forall i$ we have $\pi(S \cup \{i\}) = \pi(\{i\}) = i$* or *$r, \forall S \subseteq \Omega_i$, where Ω_i is the set of nodes of priority lower than i.*

Definition 3.3 *A priority policy π is called an* index *policy if V^π is an index function on Ω.*

The following theorem is the main structural result of this section. It states that we need only consider index policies for Problem $(\mathbf{P_1})$.

Theorem 3.1 (Index Policy) *[10] There is an optimal Markov policy π for Problem $(\mathbf{P_1})$ which is an index policy.*

We present an algorithm which computes the optimal index policy.

Algorithm 1 *(A Dijkstra-Type Algorithm for an Index Policy)*

Define the sets $\mathcal{A}$ and $\mathcal{X}$ as follows.

Initially: *$\mathcal{A}$ contains the nodes of highest reward in arbitrary order (there must be at least one such node); the action taken by the optimal index policy π on these nodes is r. $\mathcal{X}$ is the unordered complement (w.r.t. Ω) of the set of nodes of highest reward.*

During the construction of optimal policy π: *$\mathcal{A}$ contains a priority list of a set S of nodes, $S \subset \Omega$, together with the action specified by π on each node in S. $\mathcal{X}$ is the unordered complement (w.r.t. Ω) of $\mathcal{A}$.*

The algorithm proceeds as follows.

1. *For each $i \in \mathcal{X}$, compute V_i^π from*
$$V_i^\pi = \max\left\{\frac{-c_i + \sum_{S\supset\{i\}:\pi(S)\neq i} P^i(S|i)V^\pi_{\pi(S)}}{\sum_{S\supset\{i\}:\pi(S)\neq i} P^i(S|i)}, R_i\right\} \tag{3}$$
assuming that i is the node of next highest index in π.
2. *Choose $i \in \mathcal{X}$ with the highest value of V_i^π, with ties broken arbitrarily. Append this node to the list $\mathcal{A}$ as the next priority node, together with the action specified by (3). Remove i from $\mathcal{X}$.*
3. *If $\mathcal{X}$ is empty, stop. If not, goto step 1.*

In Step 1 the right-hand-side of (3) computes the best expected reward for node i, assuming i is the node of next highest index in π. This computation is feasible because π is a priority policy.

The above algorithm also resembles Klimov's algorithm [9] and has the following feature.

Theorem 3.2 *[10] For Problem (P), Algorithm 1 produces an optimal index policy.*

Note that the computational complexity of Algorithm 1 is $O(N^2)$.

4 The Time-Invariant Power Control Problem

We define the following variant to Problem $(\mathbf{P_1})$.

Problem $(\mathbf{P_2})$

We consider Problem $(\mathbf{P_1})$, with the following addition. At each time step the central controller chooses a node for transmission, from among the nodes with the message, and a power level from among the allowable levels for that node. To each node and power level is associated a transmission cost and a probability that a given set of nodes receives the message. At any time, the controller can decide to stop transmitting and receive the reward for its current state, or continue the transmission process. We seek a policy which maximizes (2) under the conditions of Problem $(\mathbf{P_1})$ and the above addition.

4.1 Analysis of Problem $(\mathbf{P_2})$

The fundamental nature of Problem $(\mathbf{P_1})$ is not altered by adding in a choice of power levels. By appropriate definition of the state-space (each power level becomes a node), we can convert Problem $(\mathbf{P_2})$ into one that is equivalent to Problem $(\mathbf{P_1})$. Thus, the nature of the optimal policy of Problem $(\mathbf{P_1})$ is maintained. The details of the specification of the appropriate state-space for Problem $(\mathbf{P_2})$ and its equivalence to Problem $(\mathbf{P_1})$ can be found in [10].

5 The Time-Varying Stochastic Routing Problem

We define the following variant to Problem $(\mathbf{P_1})$.

Problem $(\mathbf{P_3})$

We consider Problem $(\mathbf{P_1})$, with the following modifications. We allow the parameters in Model **(M)** to be time-varying. That is, at node i the transmission cost is $c_{i,t}$ $(c_{i,t} > 0, \forall t, i)$, the transition probability is $P_t^i(S'|S)$, and the reward function is $R_{i,t}$. The overall reward function is written R_t, and is an index function on Ω at each t. We further assume the existence of a time τ such that $R_{i,t} = 0$, $\forall t \geq \tau, \forall i$. At any time, the controller can either continue the transmission process, or idle for one unit of time, or stop the process and receive the reward for its current state. We seek a policy which maximizes (2) under the above conditions.

By allowing time-varying system parameters, the model of Problem $(\mathbf{P_3})$ captures the timeliness and delivery quality aspects of transmission in ad hoc networks.

5.1 Discussion of Problem $(\mathbf{P_3})$

Without any loss of optimality we restrict attention to Markov policies. We notate a time-varying Markov policy, defined over a backward time interval $1, 2, \ldots, w$, as $\pi^{\mathbf{w}} := (\pi_1, \pi_2, \ldots, \pi_w)$, where π_w refers to the Markov policy at time-to-go w. We call $\pi^{\mathbf{w}}$ a *time-varying index policy* when π_w is an index policy at each w(cf. Def. 3.3).

We present an algorithm that computes an optimal time-varying index policy for Problem $(\mathbf{P_3})$ whenever such a policy exists.

Algorithm 2 *The algorithm is defined inductively on the time-to-go w.*

When $w = 1$, the optimal policy π_1 is to retire and obtain the reward of the node which has received the message and has the largest $R_{i,\tau-1}$. This π_1 is an index policy.

Assume that for some arbitrary $w - 1 < \tau$ we have defined $\pi^{\mathbf{w-1}}$, an optimal time-varying index policy. We show how an optimal index policy $\pi^{\mathbf{w}}$ is computed, assuming a certain system condition (specified below) is satisfied at each node computation. If the condition is not satisfied, the algorithm halts, and an optimal time-varying index policy does not exist.

To compute the policy at w, we start with two sets $\mathcal{A}$ and $\mathcal{X}$ that are defined in precisely the same way as in Algorithm 1.

The algorithm proceeds in three steps.

1. *Let $R^{\mathcal{X}} := \max_{i \in \mathcal{X}} R_{i,\tau-w}$. Define $\pi_k^{\mathbf{w}}$ to be a policy which chooses $k \in \Omega$ at w, and then is identical to $\pi^{\mathbf{w-1}}$ thereafter. Compute $V^{\pi_i^{\mathbf{w}}}(\mathcal{X})$ for each node $i \in \mathcal{X}$ using*
$$V^{\pi_i^{\mathbf{w}}}(\mathcal{X}) = -c_{i,\tau-w} + \sum_{S \supseteq \mathcal{X}} P_{\tau-w}^i(S|\mathcal{X}) V^{\pi^{\mathbf{w-1}}}(S) \quad (4)$$
Define $U = \max_{i \in \mathcal{X}} V^{\pi_i^{\mathbf{w}}}(\mathcal{X})$.
2. *Consider two cases.*

Case 1: $R^{\mathcal{X}} < U$

Define $D = \{i \in \mathcal{X} : V^{\pi_i^{\mathbf{w}}}(\mathcal{X}) = U\}$. Suppose there exists $i \in D$ that satisfies either Condition 1 *or Condition 2 stated below. Set $\pi^{\mathbf{w}}(S) = i$ for all $S \subseteq \mathcal{X}$ such that $i \in S$; append i to $\mathcal{A}$ and remove it from $\mathcal{X}$.*

If there exists no $i \in D$ that satisfies either Condition 1 or Condition 2, halt. A time-varying optimal policy does not exist.

Conditions 1 and 2 are the following:

Let $j = \pi_{w-1}(\mathcal{X})$.

Condition 1 $i = j$

Condition 2 *$i \neq j$ and $\exists B \subseteq \Omega$ such that both of the following relations hold.*

$$V^{\pi^{\mathbf{w-1}}}(\{k\}) \geq V^{\pi^{\mathbf{w-1}}}(\{j\}) \;\; \forall k \in B \quad (5)$$

$$\sum_{S \supseteq \{i\}: B \cap S \neq \emptyset} P^i_{\tau-w}(S|i) = 1 \quad (6)$$

Case 2: $R^{\mathcal{X}} \geq U$

Set $\pi_w(S) = r$, $\forall S \subseteq \mathcal{X}$ such that $i \in S$, where i is any node such that $R^{\mathcal{X}} = R_{i,\tau-w}$; append i to $\mathcal{A}$ and remove it from $\mathcal{X}$.

3. *If $\mathcal{X}$ is empty, an optimal index policy $\pi^{\mathbf{w}}$ for time w has been completely specified. Otherwise, go to Step 1.*

Relations (5) and (6) mean the following: when node i transmits at w (i.e. with w time units to go), a node that is at least as good as j at $w-1$ is reached with probability 1.

The following theorem summarizes the main result about Problem ($\mathbf{P_3}$).

Theorem 5.1 *[10] A time-varying index policy is optimal for Problem ($\mathbf{P_3}$) if and only if Algorithm 2 terminates with $w = \tau$.*

A full discussion of our work on time-varying systems appears in [10].

6 A Distributed Algorithm for Problem ($\mathbf{P_1}$)

We present an algorithm which computes the optimal solution for Problem ($\mathbf{P_1}$), but which has the further characteristic that computations at each node use only information directly from neighbor nodes. This property is critical to the distributed implementation of the optimal policy in an ad hoc wireless network. We claim convergence of the algorithm to the optimal node ordering and value function under the following constraints.

1. Each node i keeps a current estimate, denoted by V_i, of its own optimal expected reward value. Each node i also stores the best estimate (i.e. most recently transmitted) of each neighbor value, denoted $V_{i,j}$, where $j \in \mathcal{N}(i)$.
2. Information transfer among neighboring nodes consists only of the current V_i value of the transmitting node.
3. Each node's V_i information is transmitted asynchronously.
4. A node's V_i update, defined below, is also asynchronous.
5. It is assumed that each node has knowledge of its own $P(S|i)$ update structure. For example, i estimates $\hat{P}(S|i)$ based on all its communications, both control signals and messages.
6. The energy required to run the algorithm is not included in finding the optimal solution for Problem ($\mathbf{P_1}$).

Specifically, we define the following

Algorithm 3 *At each event time, any number of the following two events can occur.*

Event 1 *A node i receives V_j from a neighbor j, $j \in \mathcal{N}(i)$.*

Event 2 *A node i recomputes V_i using the current $V_{i,j}$ values, as follows.*

$$V_i^n = \max\left\{\max_{\tilde{\pi}}\left\{\frac{-c_i + \sum_{\mathcal{N}(i)\supseteq S\supset\{i\}:\tilde{\pi}(S)\neq i} P^i(S|i)V^n_{i,\tilde{\pi}(S)}}{\sum_{\mathcal{N}(i)\supseteq S\supset\{i\}:\tilde{\pi}(S)\neq i} P^i(S|i)}\right\}, R_i\right\} \quad (7)$$

The maximization in (7) is over all local priority orderings $\tilde{\pi}$ of i and its neighbors.

Events 1 and 2 occur infinitely often at each node i.

In [10], we prove the following. As long as the value functions are distinct for each node, then Algorithm 3 converges to these value functions in finite time. If two nodes have identical value functions, then value function convergence is guaranteed only asymptotically, though an optimal policy is determined in finite time.

The above results apply for any set of valid initial V_i estimates, and hence convergence is achieved in diverse situations, such as after abrupt changes in network topology. More specifically, when there are abrupt network topological changes, convergence of Algorithm 3 occurs under the following condition: the topological changes are a transient phenomenon, and message transmissions, receptions, as well as value function updates according to (7) take place during these changes.

As a result of this condition, at the end of the changes each node has a set of initial estimates of its value function and of its neighbors' value functions. Initiated with these estimates, Algorithm 3 converges to the correct value functions and an optimal routing policy. This result implies that Algorithm 3 will continue to compute optimal routing polices most of the time even when network topological changes occur continuously, provided that the rate (time-scale) at which these changes occur is much slower than the algorithm's rate of convergence.

Also in [10] we use the specific problem structure to provide an efficient algorithm for the computation in (7), and we relate Algorithm 3 and our method of proof to the distributed Bellman-Ford algorithm [2], and to the more encompassing theoretical work in [1].

7 Conclusion

We formulated a network routing problem which uses a probabilistic local broadcast model for wireless transmission. We showed that an index policy is optimal for this problem. We extended the model to allow for power control, and showed that the index nature of the optimal policy remains unchanged. Subsequently we allowed time-varying system parameters in the model, and discussed conditions under which a time-varying index policy is optimal. Finally, we presented a distributed implementation of the optimal routing policy for the original problem, and provided results on its convergence properties. A more detailed version of this work is available in [10], where further results are presented and the claims made in this paper are proved.

Future work in this area can take many forms. A fuller accounting of optimal policies for the time-varying problem needs to be made. For the distributed algorithm, important questions remain, such as convergence rate, counting to infinity, and packet looping. Distance vector algorithms have been in a long process of evolution to deal with these problems, and some of these ideas can perhaps also be used here. Also, methods to estimate the local transmission probabilities at each node should be studied. Finally, a study of the multicast problem for the model of this paper is of great interest, and could lead to more interesting optimization formulations.

References

[1] D. Bertsekas, "Distributed Dynamic Programming", *IEEE Transactions on Automatic Control* **27**:3, p.610-616, June 1982

[2] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, 1992

[3] J. Broch, D. Maltz, D. B. Johnson, Y. C. Hu, J. Jetcheva, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols", *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, October 25-30, 1998

[4] M. Corson and A. Ephremides, "A Distributed Routing Algorithm for Mobile Wireless Networks", *Wireless Networks* **1**, p. 61-81, 1995

[5] E. Gafni and D. Bertsekas "Distributed Algorithms for Generating Loop-Free Routes in Networks with Frequently Changing Topology", *IEEE Transactions on Communications* **29**:1, p.11-18, 1981

[6] J. Garcia-Luna-Aceves and S. Murthy "A Path-Finding Algorithm for Loop-Free Routing", *IEEE/ACM Transactions on Networking*, v 5 n 1, p.148-160, Feb 1997

[7] Z. Haas "Guest Editorial, Wireless Ad Hoc Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1329-1330, August 1999

[8] D. B. Johnson, "Routing in Ad Hoc Networks of Mobile Hosts", *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, p.158-163, December 1994

[9] G. Klimov, "Time Sharing Service Systems I", *Theory of Probability and its Applications*, **19**, p.532-551

[10] C. Lott and D. Teneketzis, "Stochastic Routing in Ad Hoc Wireless Networks", University of Michigan Control Group Report CGR 00-02 (in preparation)

[11] D. Maltz, J. Broch, J. Jetcheva, D. Johnson "The Effects of On-Demand Behavior in Routing Protocols for Multihop Wireless Ad Hoc Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1439-1453, August 1999

[12] J. Moy, "OSPF Version 2", RFC 1247 July, 1991

[13] V. Park and M. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks", *Proceedings of INFOCOM '97*, p.1405-1413, April 1997

[14] M. Pearlman and Z. Haas "Determining the Optimal Configuration for the Zone Routing Protocol", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1395-1414, August 1999

[15] C. Perkins "Ad Hoc On Demand Distance Vector Routing", Internet draft, draft-ietf-manet-aodv-00.txt, November 1997. Work in progress.

[16] C. Perkins and P. Bhagwat "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers", *Proceedings of the SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, p. 234-244, August 1994

[17] M. Pursley, H. Russell, P. Staples "Routing for Multimedia Traffic in Wireless Frequency-Hop Communication Networks", *IEEE Journal on Selected Areas in Communications* **17**:5, p.784-792, May 1999

[18] V. Rodoplu and T. Meng "Minimum Energy Mobile Wireless Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1333-1344, August 1999