

Performance of Routing Protocols in HF Wireless Networks

Eric E. Johnson*, Manikanden Balakrishnan*, and Zibin Tang*
Klipsch School of Electrical and Computer Engineering
New Mexico State University

ABSTRACT

In this paper we compare the performance of the Wireless Address Resolution and Routing Protocol (WARRP) and the Internet standard Optimized Link State Routing (OLSR) protocol (RFC-3626) over various topologies of the HF Wide Area Network (HFWAN). WARRP was designed specifically for HF networks, in which network diameter seldom exceeds two hops, while OLSR is a general-purpose wireless routing protocol. We explore the effects of channel impairments and node mobility on routing overhead and network performance.

1. INTRODUCTION

High Frequency (HF) radio provides extended line-of-sight and beyond-line-of-sight wireless communications for defense, diplomatic, and commercial applications. HF networks are often fully connected (requiring no relaying), especially in surface-wave applications. However, the HF channel can be an unreliable propagation medium, with significant packet loss rates and other environmental effects, leading to intermittent link outages and even network partitions. Thus an indirect routing capability is an attractive feature for HF networks, although most of the routing issues can be addressed by a simple single-relay routing mechanism rather than the more complex routing protocols used in the Internet.

The Optimized Link State Routing (OLSR) protocol (RFC-3626) [1] is an optimization of the classical link state routing algorithm, designed to suit mobile ad-hoc networks. OLSR is generally appropriate for large and dense wireless networks, where it outperforms traditional link-state routing schemes due to its novel mechanisms for reducing routing overhead transmissions. The Wireless Address Resolution and Routing Protocol (WARRP) is an integrated address resolution and routing functionality originally developed for mobile HF WANs (interconnecting IP-based sub-networks), introduced in [2]. WARRP is an on-demand routing scheme that extends ARP to provide additional routing capabilities, tailored to the needs of an ad-hoc HF network.

In this paper we evaluate and compare the performance of OLSR and WARRP through a combination of

mathematical analysis and simulation. The analysis compares the volume of overhead transmissions generated by these protocols under various network topologies that approximate a typical HF WAN. The results quantify the suitability of these protocols to HF networks.

2. WIRELESS ROUTING PROTOCOLS

Routing is the function in a network that finds a path through one or more links (or subnetworks) for delivering user traffic from a source node to a destination node. Routing protocols for wireless networks face unique challenges, compared to routing in wired networks. In a wireless network, links are usually lower in bandwidth and more error-prone; node mobility can introduce frequent topology changes; and nodes may have energy and processing constraints if they are battery powered. As a result, novel routing protocols have been developed to address these unique challenges.

2.1 Optimized Link-State Routing (OLSR)

The Optimized Link-State Routing (OLSR) protocol [1] is a proactive routing protocol; that is, it continually monitors the connectivity among network nodes and floods connectivity summaries through the network to update routing tables at each node. In this section, we summarize the aspects of OLSR that bear upon the overhead burden that it places on the network.

2.1.1 Multipoint Relaying

A novel feature of OLSR is the use of a subset of network nodes to flood topology information. In previous link-state routing protocols, a node's connectivity is flooded to all other network nodes: the node sends its connectivity to all of its neighbors, each of which re-sends the packet to that node's neighbors, and so on until all nodes in the network have sent the packet once.

By contrast, OLSR nodes select a subset of their neighbors for relaying their connectivity throughout the network. Multipoint Relay (MPR) nodes are selected from a node's neighbors in such a way that the node can reach all of its one- and two-hop neighbors either directly or via relay by that node's MPR set.

* This work supported by the US Navy Space and Naval Warfare Systems Center under contract N66001-03-D-0042

2.1.2 Connectivity Monitoring

Each node in an OLSR network periodically broadcasts a *Hello* packet. In addition to identifying the sender of the packet, these packets summarize local connectivity at that node. Connectivity to each neighbor is reported as being unidirectional, bidirectional, or MPR (i.e., selected as a multipoint relay by the reporting node). Hello packets are not forwarded by the neighbors, but are used to update the connectivity table at each node that receives them; each node also maintains an MPR Selector Table that lists the nodes that have selected it as an MPR. A typical Hello interval is 2 seconds.

2.1.3 Topology Distribution

At somewhat longer intervals (e.g., 5 seconds), each node selected as MPR by some node creates a Topology Control (TC) packet that contains its MPR Selector Set, and floods this TC information throughout the network using the MPR mechanism.

Every node in the network updates its topology table using the MPR Selector Sets received in TC packets. This topology table contains (*destination*, *last_hop*) pairs. The *destination* in each pair is reachable via the associated *last_hop* node. The *last_hop* node is typically an MPR for the *destination* node and has advertised this link in a TC packet. OLSR routes are built by stringing together these pairs to form a chain from the sender to the destination.

2.2 Wireless Address Resolution and Routing Protocol

The Wireless Address Resolution and Routing Protocol (WARRP) is a mixed proactive/reactive (“on demand”) routing protocol developed specifically for small-diameter networks. Using WARRP, nodes can query their neighbors for connectivity to other nodes. Responses to WARRP queries carry one of several codes:

1. Identity: the responder is the desired destination
2. Direct connectivity: the responder can relay to the destination
3. Gateway: the responder has no direct connectivity information, but is connected to other wide-area network(s) that may be able to reach the destination
4. Relay to gateway: the best the node can offer is to relay traffic to a gateway

WARRP is most effective when at least one response to a query is type 1 or 2, meaning that the distance in hops to the destination is 1 or 2.

One unusual feature of WARRP is its support for extensive subnetworks attached to, and reachable via, each wireless node. Such an environment occurs when mobile

platforms (ships, aircraft, etc.) carry multiple hosts on one or more LAN(s) that are connected to a wireless network via a router or a dual-homed node on the LAN. WARRP advertises all of the subnetworks that are reachable via the single wireless node.

Another unusual feature of WARRP is its integration of an ARP-like address resolution protocol with the routing functions: WARRP responses carry not only the IP addresses of the hosts and subnets reachable via the responding wireless node, but also the MAC address of that node.

WARRP can operate in strict on-demand mode, but it also includes a simple proactive feature that can improve its performance when topology changes are rare: as each node joins the network, it broadcasts a Topology Request, to which each neighbor responds with a complete list of subnetworks reachable via that wireless node. These subnet-to-MAC address mappings can be cached to provide proxy ARP functionality.

3. ANALYTICAL COMPARISON

Because the HF radio channel is characterized by low bandwidth and relatively long link turnaround times [Tt], we have chosen as a performance metric the rate of routing-related overhead transmissions by network members. While the number of bytes or packets is also of interest, the quantization of transmission time as integer multiples of the interleaver time (hundreds to thousands of milliseconds) means that minor differences in byte counts disappear in the necessary padding at the interleaver. It is straightforward to convert rate of transmissions to channel utilization when this interleaver quantization results in all packets having the same on-air duration.

Since we are comparing proactive and reactive protocols, the delay before message transmission can begin is also of interest, and is discussed at the end of this section.

We will compare the latency and overhead rates of our two routing protocols in a range of networking environments:

- Address resolution protocol: OLSR can use either ARP or the proxy address resolution mechanism based on WARRP topology requests; this subset of WARRP is here termed WARP (dropping the first R which stands for “routing). WARRP routing always uses its built-in address resolution mechanism.
- Node mobility: static topologies (no mobility) *versus* dynamic topologies

We begin by analyzing the two protocols, and then apply these models to our scenarios of interest. The following parameters are used in the models:

T_{Hello}	Hello packet interval (s)
T_{TC}	TC packet interval (s)
R_{Hello}	Total rate of Hello transmissions in the network
R_{TC}	Total rate of TC transmissions in the network
$N_{network}$	Number of nodes in network
N_{1-hop}	Average number of neighbors
N_{relay}	Average number of relay nodes in network (0 for a fully connected network)
N_{MPR}	Average size of MPR Set at relay nodes (0 for network diameter ≤ 2)
T_{stable}	Mean time between topology changes
T_{user}	Mean time between user messages in each direction for each pair of nodes
R_{user}	Total rate of user messages in the network
T_{ARP}	Address resolution protocol (ARP) cache timeout
R_{ARP}	Total rate of ARP transmissions in the network

3.1 OLSR Overhead Analysis

OLSR generates two types of overhead packets: Hello and TC. In this section, we develop a model for the rate of overhead transmissions given network size and connectivity. Note that user traffic has no direct effect on the routing overhead due to OLSR. Address resolution overhead traffic, however, is driven by user traffic.

- Hello traffic is generated by every node at a fixed rate, with no relaying. Total network traffic for Hello packets is

$$R_{Hello} = \frac{N_{Network}}{T_{Hello}}$$

- TC traffic is also generated by each MPR node at a fixed minimum rate. (Although more frequent TC updates may triggered when a topology change is detected, there is a minimum interval between such TC updates, which we assume will not be significantly different from the TC interval in HF networks.) TC packets are relayed via the MPR scheme; a rough approximation of the number of retransmissions of each TC packet is $\approx N_{relay}N_{MPR}$.
- Total network traffic for TC packets comprises the initial transmissions and retransmissions

$$R_{TC} = \frac{N_{relay}}{T_{TC}} (1 + N_{relay}N_{MPR})$$

- User traffic is modeled here as being uniformly distributed among all pairs of network nodes. The time between successive messages sent in each direction be-

tween any pair of nodes is T_{user} . If user traffic to non-neighbor nodes requires two hops, the total rate of user messages in the network is

$$R_{user} = \frac{N_{Network}}{T_{user}} (N_{1-hop} + 2(N_{network} - 1 - N_{1-hop}))$$

- If ARP is used for address resolution, neighbor entries in the ARP cache must be refreshed *when used* if they are older than T_{ARP} . User messages are sent only to 1-hop neighbors. Each ARP request prompts a response so total network traffic for ARP packets is

$$R_{ARP} = \frac{2N_{Network}N_{1-hop}}{\max(T_{ARP}, T_{user})}$$

3.2 WARRP Overhead Analysis

In this section, we develop a model for the rate of overhead transmissions generated by WARRP. Apart from the initial Topology Requests and the responding Local Subnet Reports, WARRP generates overhead only when connectivity changes, and then only when a link outage is discovered during an attempt to send traffic over a link that has failed. In response to discovery of a link outage, a single WARRP Request is sent, and each 1-hop neighbor of the requesting node returns a WARRP Response.

The rate of discovering link outages depends on both the rate of topology changes and the rate of traffic on affected links (since link outages are discovered only by attempts to send user traffic). An upper bound on the rate of WARRP overhead transmissions assumes that every topology change results in a WARRP Request plus WARRP Responses from each 1-hop neighbor:

$$R_{WARRP} \leq \frac{N_{1-hop} + 1}{T_{stable}}$$

3.3 Overhead Results: Static Topology

We examine the routing overhead produced by our protocols under both static and dynamic topologies. The simpler case is a static topology, in which the network graph may be either fully or partially connected but does not change. For this analysis, we consider the two topologies shown in Figure 3:

- a completely connected 6-node network: every node is connected to every other node
- a 6-node network which lacks connectivity among nodes at the geographical extremes of the topology: the center nodes can reach every other node, but the nodes on either side cannot reach the nodes at the other side.




Figure 3a. Fully-connected topology




Figure 3b. Partially-connected topology

3.3.1 Routing Overhead Packet Rates

In the static topologies, WARRP generates no overhead packets after the network is formed.

OLSR generates a steady-state overhead rate due to Hello messages (as does ARP if used). For the partially connected topology an MPR node will be selected and it will send TC messages, but fewer ARP packets are sent than in the fully connected case. $R_{OLSR} = R_{Hello} + R_{ARP}$ is plotted as a function of T_{Hello} in Figure 4. OLSR with WARP address resolution is also shown in Figure 4.

3.3.2 Routing Overhead Channel Utilization

Each HF transmission requires overhead including link turnaround time and the modem synchronization preamble at the beginning of a transmission. However, if routing packets are appended to transmissions already required (e.g., tokens in a token-passing network), this overhead is already borne by those transmissions, and the routing protocol is responsible only for the additional channel time needed for the routing packets.

Transmission times with HF modems are quantized by the interleaver time in use. With current HF technology (MIL-STD-188-110B or STANAG 4539 modems), the shortest interleaver available is 120 ms. A reasonable data rate for surface-wave channels is 6400 bps, which results in a quantum for packet size of 96 bytes (including headers from the link layer and others as required).

An OLSR packet carries both IP and link headers, totaling about 44 bytes (for STANAG 5066 data link protocol). A WARRP packet carries just the link-layer header, or about 24 bytes. For our 6-node networks, both protocols will send packets somewhat smaller than a single interleaver, so on average a routing packet will increase a transmission by $T_{pkt} = 120$ ms. The channel utilization due to routing protocol overhead is therefore approximately

$$U_{routing} = T_{pkt} R_{routing}$$

where the “routing” subscript is OLSR or WARRP.

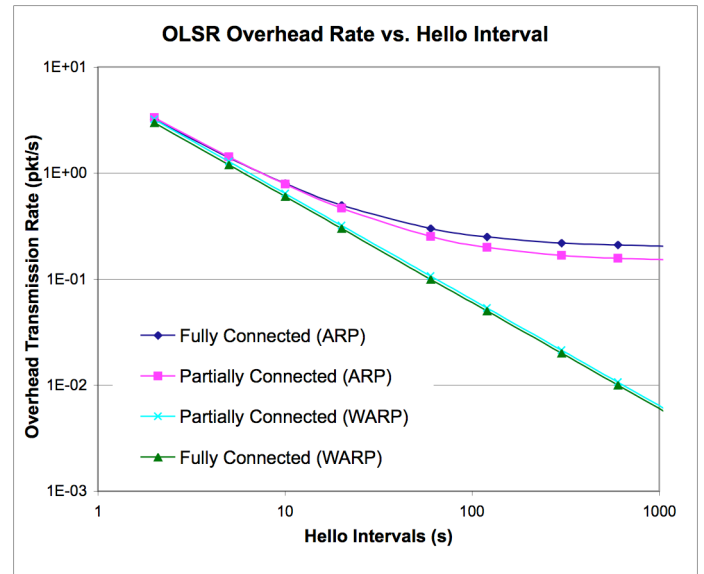


Figure 4. Static Analysis of OLSR

From Figure 4, we see that any Hello interval of up to 10 s results in routing overhead of more than 10%, and that to reduce routing overhead to 1% requires a hello interval of at least 100 s and use of WARP.

3.4 Overhead Results: Dynamic Topologies

To simplify the analysis of dynamic connectivity, we consider two simple topologies:

- a network in which nodes leave and join at a known average rate, and
- a network which flips back and forth between the fully connected and partially connected topologies analyzed in the previous section.

3.4.1 Leave/Join Topology

In our first dynamic topology, at intervals averaging T_{stable} , a node joins or leaves a fully-connected network. Joining and leaving occur with equal probability, and the average number of nodes in the network is $N_{network}$.

The average rates of Hello and ARP transmissions for OLSR are proportional to $N_{network}$, so their average rates do not depend on T_{stable} . Since the network is fully-connected, TC packets will not be generated.

In contrast, WARRP is strongly affected by the changes in topology. On average, half of the topology changes result in a new node joining the network, resulting in a Topology Request plus WARRP Responses from each 1-hop neighbor. Thus, we have WARRP overhead at the following rate:

$$R_{WARRP} = \frac{N_{1-hop} + 1}{2T_{stable}}$$

In Figure 5, we compare the overhead of WARRP to that of OLSR (with 1-minute and 5-minute Hello intervals, and using ARP or WARP for address resolution). In this scenario, WARRP generates less overhead than OLSR whenever the Hello interval is adequate to track the changing topology ($T_{\text{Hello}} \leq T_{\text{stable}}$).




Figure 5: OLSR vs WARRP Overhead: Leave/Join Scenario

3.4.2 Sides Connected/Disconnected Topology

In our second dynamic scenario, the connections between nodes on the sides of the network are periodically lost and recovered. The interval between changes averages T_{stable} .

In the case of OLSR, we need $T_{\text{Hello}} \ll T_{\text{stable}}$ if the routing protocol is to be able to detect the changes and proactively compute indirect routes. Each time a side node discovers that it has lost connectivity to the other side nodes (usually the result of consecutive missing Hello packets from those nodes), it will select a MPR to relay to them (one of the end nodes in Figure 3b), and will identify that MPR in its next Hello packet. The MPR will then begin sending TC packets, which will inform all other side nodes that this MPR can reach every other node. Thus, only one MPR will usually be selected.

When the side nodes recover connectivity to the other side (by receiving a Hello), the side nodes will clear their MPR sets, but transmission of TC packets by the former MPR will continue for $3 T_{\text{TC}}$ to void previously stored topology data. Thus, the additional routing overhead generated by OLSR for each connectivity change in this scenario amounts to a series of TC packets sent from the se-

lected MPR node. The total routing overhead produced by OLSR in this scenario is approximately

$$R_{\text{OLSR}} = R_{\text{Hello}} + R_{\text{TC}} \left(\frac{T_{\text{stable}} + 3T_{\text{TC}}}{2T_{\text{stable}}} \right)$$

WARRP will not generate Topology Request and Response exchanges in this case, since none of the nodes are entering the network. However, WARRP Requests (and WARRP Responses) will be generated whenever traffic cannot be sent over links that previously worked. In the worst case, each *loss* of connectivity will be discovered, and will prompt each disconnected side node to send a WARRP Request, which will result in WARRP Responses from each of its 1-hop neighbors. Thus, the upper bound on WARRP overhead in this scenario is

$$R_{\text{WARRP}} \leq \frac{N_{\text{disconnected}} N_{\text{side-nodes}} (N_{1\text{-hop}} + 1)}{\max(T_{\text{user}}, 2T_{\text{stable}})}$$

where $N_{1\text{-hop}}$ is the reduced 1-hop neighborhood when the sides are disconnected. These results are plotted in Figure 6, showing overhead for WARRP and for OLSR with 1-minute and 5-minute Hello intervals.




Figure 6: OLSR vs WARRP Overhead: Sides Connected/Disconnected Scenario

3.5 Latency Analysis

The higher overhead of proactive protocols such as OLSR is the price paid for immediate readiness to route packets. This contrasts with reactive protocols, which must some-

times first find a route before user packets can be sent onward toward their destination. In this section, we briefly examine the latencies introduced by OLSR, ARP, and WARRP in HF networks.

For OLSR, a route can be computed immediately from the routing table that is constantly maintained by OLSR. Delays in traffic delivery will occur only if the routing table is incorrect (a link has failed within the past $3 T_{\text{hello}}$), or an ARP cache entry has expired. (The latter source of delays can be eliminated if WARP or another proxy ARP mechanism is employed.)

Normally, a node using OLSR will be able to begin packet transmission after just the channel access delay. In a token-passing network, the average MAC delay is just half of the token rotation time. The additional delay for an ARP response would be a single token rotation time. Routing failures after a connectivity change would interrupt traffic until the routing table converges to the new topology; this could require as long as $3 T_{\text{hello}} + T_{\text{TC}}$.

WARRP, due to its proactive Topology Request mechanism, also maintains a routing/address resolution table, although it does not actively maintain it until a routing failure occurs. Thus the best-case latency for WARRP is also just the channel access time. When the WARRP table is incorrect, the defect is repaired by a WARRP request and WARRP responses. In a token-passing network, this process requires a single token rotation.

Thus, the combination of proactive and reactive features in WARRP brings its expected latency for user messages (in token-passing networks) very close to that of OLSR: both usually require only one half the token rotation time, with an occasional extra token cycle to update the WARRP table or ARP cache, respectively.

4. SIMULATIONS

Simulations using the NetSim HF network simulation environment were used to validate the analytical results. We used networks of six nodes in the topologies described in the previous section.

- User traffic was generated independently at each node using a memoryless interarrival distribution with mean interarrival time of T_{user} .
- TCP was not used in these simulations.
- OLSR or WARRP was used to maintain routing tables.
- Address resolution from IP to MAC addresses employed ARP, WARP, or WARRP.
- A simple ARQ protocol was employed, which declared a link down when packet(s) to a destination were unacknowledged after one token rotation.

- The MAC layer used HFTP (token passing).
- A simple go/no-go channel model was used, in which packets were either received error-free or lost entirely. The channel model included a feature to switch connectivity tables at fixed or stochastic intervals.

4.1 Static Scenarios

The static scenarios used in our simulation experiments are the same as in the analysis section: a fully connected network of 6 nodes, and a partially connected network in which the two sides cannot reach each other directly.

As expected WARRP introduced no overhead in the static topologies after the initial round of Topology Requests. OLSR also behaved more or less as expected from the analysis. Analytical and simulation results for OLSR are compared in Figure 7 (with 90% confidence intervals) with Hello intervals ranging from 1 to 20 minutes.

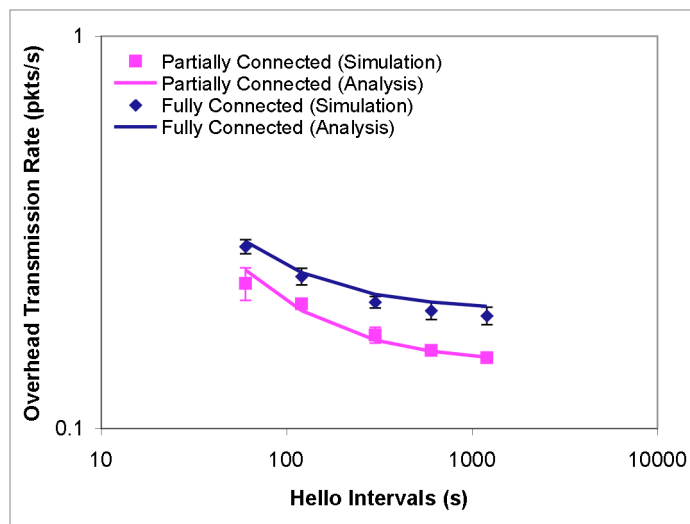


Figure 7. OLSR/ARP Static Scenario Results (Simulation vs Analysis)

4.2 Dynamic Scenarios

The dynamic scenarios used for the analytical model were also replicated in the simulator.

Figure 8 compares the results of simulation and analysis for the leave/join scenario. In every case, simulation and analysis agree to within the 90% confidence intervals. Much of the variation in the simulation results arises from interactions with the MAC layer, which allows each node to transmit only once each token rotation (about 60 s).

The wider relative variation seen in the WARRP simulations results is partly exaggeration on the log-scale graph due to the low level of overhead in WARRP, and partly a reflection of explicit interaction between WARRP and

HFTP: the time required to add nodes to the ring as they join determines the rate of WARRP Topology Requests and Responses. This ring joining time varies from run to run because nodes arrive asynchronously to the process of soliciting successors by active ring members (see [3] for further explanation of this process).

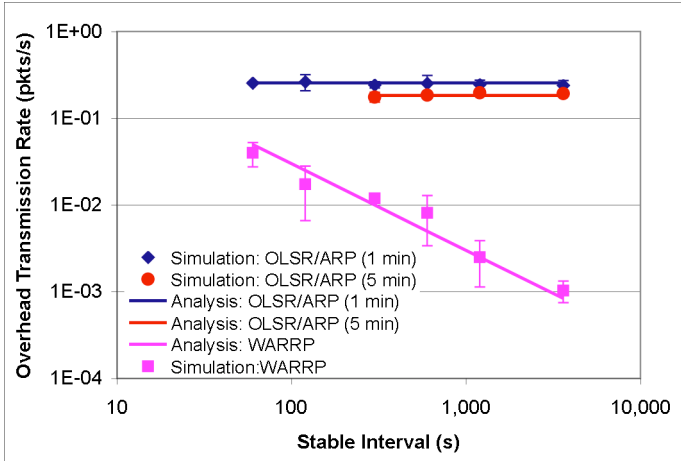


Figure 8. Leave/Join Scenario Results (Simulation vs Analysis)

In Figure 9, the results for simulation and analysis of the Sides Connected/Disconnected Scenario are compared. As before, we see good agreement between simulation and analysis, with larger relative variation in the WARRP results. In this case, WARRP appears sensitive to variations in the time required by the ARQ/token MAC combination to determine that a link has failed.

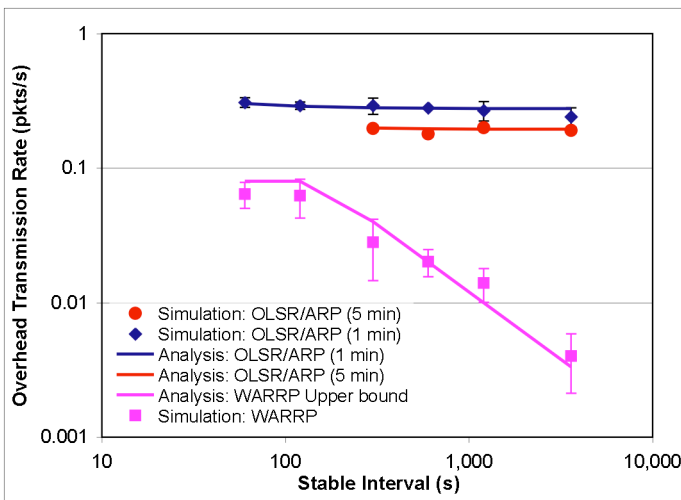


Figure 9. Sides Connected/Disconnected Results (Simulation vs Analysis)

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have examined the performance characteristics of two different philosophies in wireless network routing protocols in the HF radio environment: proactive routing, exemplified by OLSR, and application-specific reactive routing, exemplified by WARRP. In general, the long range of HF radio should provide rich connectivity and will often render routing a minor concern. This would suggest that a proactive routing protocol would be overkill unless its update rate is reduced to approximate the expected rate of topology change in the network.

Our results confirm this intuition, and show that even with low update rates, OLSR generated more routing overhead than the reactive WARRP protocol. Nevertheless, OLSR is well-known and software implementations are readily available, and its overhead may be sufficiently low to permit its use in networks that are not saturated with user traffic.

Future refinements of this investigation will replace the simple ARQ and stateless traffic generators reported here with a full implementation of the STANAG 5066 protocol, TCP, and scripted user traffic generators that provide traffic profiles typical of military HF applications. Detailed investigation of routing latencies would be valuable, as would investigation of networks of diameter 3, beyond which WARRP is probably not a good solution.

Trials of the OLSR within the Allied HF WAN software stack are anticipated in the near future, perhaps followed by similar trials using an implementation of WARRP.

6. REFERENCES

1. T. Clausen, *et al.*, "Optimized Link State Routing Protocol, IEEE INMIC, Pakistan, 2001.
2. E.E. Johnson, *et al.*, "Routing in HF Ad-Hoc WANs," *Proceedings of MILCOM 2004*, Monterey, CA, October 2004.
3. E.E. Johnson, *et al.*, "Robust Token Management for Unreliable Networks," *Proceedings of MILCOM 2003*, Boston, MA, 2003.
4. E.E. Johnson, M. Balakrishnan, and Z. Tang, "Impact Of Turnaround Time On Wireless MAC Protocols," *Proceedings of MILCOM 2003*, Boston, MA, 2003.