

Routing in HF Ad-Hoc WANs

Eric E. Johnson, Zibin Tang, Manikanden Balakrishnan, Huiyan Zhang, and Srugun Sreepuram
Klipsch School of Electrical and Computer Engineering¹
New Mexico State University

ABSTRACT

High frequency (HF) radio is commonly used to extend wireless communications beyond the line-of-sight range that limits the higher frequency bands. Despite such long-range coverage, however, indirect routing is sometimes required even in HF networks. In addition, HF radio and other wireless technologies are emerging as a means to interconnect wired subnetworks in various mobile and contingency applications. In this paper, we introduce WARRP, a Wireless Address Resolution and Routing Protocol that is specifically designed to integrate single-relay route discovery with address resolution in such ad-hoc WANs.

1. INTRODUCTION

High frequency (HF) radio offers beyond-line-of-sight wireless communications for applications ranging from extended line-of-sight within a naval battlegroup to global coverage supporting commercial and military aviation. The long-haul links available using transportable HF equipment also provide quick communications into disaster areas where the terrestrial infrastructure may have been severed or destroyed.

Despite this ability to communicate beyond line of sight, vagaries of propagation and other environmental effects can sometimes produce outages on some HF links while leaving others intact. Thus, reliability in HF networks is enhanced when indirect routing is supported [1]. Of course, most routes in an HF network usually require only a single link, and a single-relay routing mechanism should solve most of the remaining cases.

When an HF subnetwork is employed to interconnect IP-based networks, routing must be supported by an address resolution protocol that can determine the correspondence between HF link-layer addresses (e.g., STANAG 5066 [2]) and IP addresses. This is similar to the function performed by the Address Resolution Protocol (ARP) in Ethernet Local Area Networks (LANs); in fact ARP can be used in HF subnetworks whenever a

subnet mask can be used to determine whether the destination host is connected to the HF subnetwork. However, in some cases (e.g., naval battlegroups) the IP address of each HF node belongs to its respective LAN subnet; that is, the HF interfaces do not have IP addresses that are distinct from their LAN interfaces. In these cases, the IP addresses of HF subnet nodes must be treated as arbitrary, and subnet masking cannot be employed to determine whether any destination address is directly connected to the HF subnet.

In this paper, we will introduce an extension to ARP that integrates single-relay routing with address resolution in subnets having arbitrary IP addresses. The protocol also integrates naturally with LAN and WAN routers. While this protocol was conceived for use in HF networks, it is also under consideration for use in other wireless applications, and has therefore been termed the Wireless Address Resolution and Routing Protocol (WARRP).

We begin by describing a typical application for WARRP: a wireless WAN that links the shipboard LANs of a naval battle group. We then introduce the protocol and describe its packet structure and operation. The paper concludes with a discussion of other applications of WARRP.

1.1 Shipboard LAN

Figure 1 illustrates a notional shipboard LAN, router, and HF node. (In this and the following figures, truncated addresses are used for simplicity: the uppermost address bits are omitted, leaving only subnet and host parts. IP addresses are shown in regular font, while MAC addresses are circled italics.)




Figure 1: Notional Shipboard LAN

¹ This work supported by US Navy Space and Naval Warfare Systems Command, contract N660001-3287-2827LG.

- The LAN is the 4.0 IP subnet, in which the router port has been assigned the 4.1 address. Hosts in the 4.0 subnet use the router at 4.1 as their default gateway.
- The host running the HF protocols has been placed in a separate subnet (4.32), and has an IP address of “4.34” The STANAG 5066 (MAC) address assigned to this node is 13. Note that this host forwards IP datagrams between its two interfaces. Since HF provides the only connection off the ship in this case, the router uses 4.34 as its default gateway.
- COMSEC (not shown) is typically present between the HF host and the physical-layer HF block.

1.2 HF WAN

Figure 2 shows three ships interconnected by HF radio, forming an HF WAN. IP subnets have been set up on each ship. Subnet addresses may have been assigned independently, and there is no guarantee that a common prefix exists among those addresses. Note that the ship carrying the 6.x subnets also has a SATCOM connection to the Internet; the router port in the SATCOM subnet has IP address 12.23.

1.3 Routing Requirements in the HF WAN

A natural approach for routing within the HF WAN would be to operate the HF nodes as routers, with distinct IP addresses for each of their two interfaces (LAN and HF). The IP addresses in the HF subnetwork would be assigned in a single maskable address range. The HF nodes would operate as routers: when given an IP packet for delivery, the HF node would mask the destination IP address to determine whether it belongs to the

HF subnet. If so, ARP would be employed to determine the MAC address of that node; otherwise, the IP address would be sought in a routing table.

However, for reasons beyond the scope of this paper, it is not possible to assign a separate IP address for the HF interface on the HF nodes. Instead, each HF node uses a single IP address on both interfaces, operating more as a filtering repeater than as a router. As a result, the HF nodes cannot use a subnet mask to determine whether to look up the destination in a routing table versus an ARP cache.

In Figure 2, all outbound packets from hosts in the 4.0 subnet will be forwarded to the local HF node (4.34) by the local router. The HF node could simply broadcast each packet to all other HF nodes and listen for link-layer acknowledgements to determine whether a packet was accepted by another node; however, this is unreliable and places a heavy processing load on all HF nodes. Instead, we would prefer to maintain a table at each node that identifies the MAC address to which packets for various subnets should be sent. This table combines features of a routing table and an address resolution cache: it is organized by destination subnets and includes a default gateway (like a routing table), but returns a MAC address (like an ARP cache).

Now consider the router connected to the 6.0 subnet. Outbound packets arriving from LAN hosts need to be forwarded to the local HF node (6.98) if the destination is reachable via the HF WAN, and sent toward the Internet via the SATCOM link otherwise. This router needs a current list of subnets accessible via the HF WAN for this routing decision; the router may also advertise to external routers the subnets reachable via HF so that traffic from distant hosts can reach ship-board hosts.

The HF nodes connected to stub LANs (e.g., 4.34) could use an ARP-like reactive protocol to discover MAC addresses to which outbound traffic should be sent. However, routers that serve as gateways between the HF WAN and external WANs need to maintain current lists of accessible subnets so that traffic addressed to the stub LANs can be routed correctly. The HF nodes connected to such gateways therefore need to monitor accessible subnetworks proactively.

From this proactive mechanism for monitoring connectivity to reachable subnetworks, it is but a small step to support relaying around link outages. As we will see, the procedure for querying HF subnet

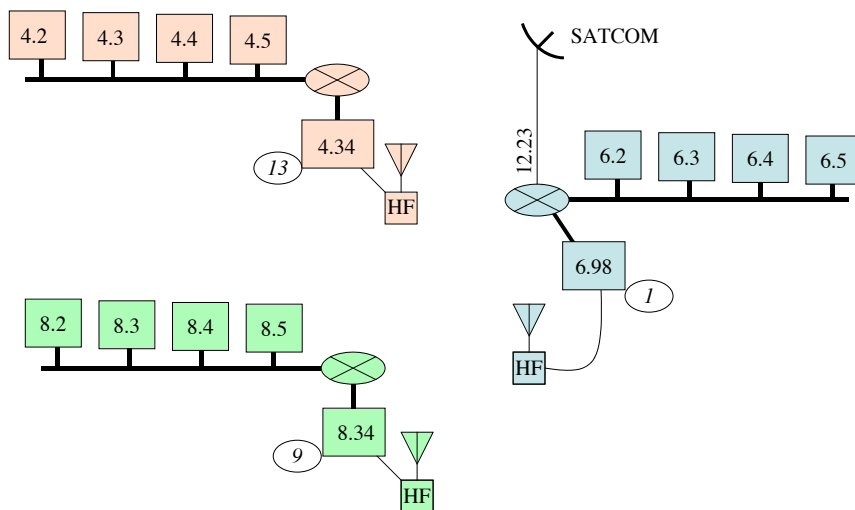


Figure 2: Notional HF WAN

members to find a single-relay path to another node integrates naturally with the ARP query procedure.

1.4 Routing in an Ad-Hoc WAN

A mobile ad-hoc network (MANET) may be defined as one in which neither the nodes composing the network, their topology, nor their connectivity is static. MANET research often considers single-host nodes whose addresses can be assigned dynamically. In the case of an HF WAN as described above, however, the mobile nodes are entire ships, which each carry one or more networks with static IP addresses. Routing in the HF WAN must support the arrival and departure of such ships. Furthermore, the networks linked by the HF WAN may have arbitrary addresses, so we cannot rely on subnet masks to determine whether or not a desired host can be reached via the HF WAN.

Thus, we need a routing approach that deals efficiently with the arrival and departure of entire autonomous systems, which may have arbitrary IP address ranges. To distinguish this case from the well-understood term MANET, we call such networks Ad-Hoc WANs.

The Wireless Address Resolution and Routing Protocol (WARRP) was developed specifically for the unusual requirements we find in Ad-Hoc WANs.

2. THE WARRP PROTOCOL

WARRP extends ARP to provide the additional capabilities required for ad-hoc WANs. We begin by describing the WARRP Protocol Data Unit (PDU), followed by a discussion of the WARRP protocol.

2.1 WARRP PDU

The standard ARP packet is shown in Figure 3. The function of ARP is to communicate the correspondence of network-layer protocol addresses (usually IP addresses) with link-layer “hardware” addresses (for example, Ethernet addresses). Thus we find in the ARP packet fields for Protocol and Hardware Addresses of the sending node and the target. The Opcode field distinguishes ARP requests from responses.

The remaining fields serve to make ARP sufficiently general to handle a wide range of link-layer and network-layer protocols: the Address Space fields contain codes that identify the address space in use at each layer (e.g., Ethernet and IP), while the Length fields specify the number of bytes in the respective addresses.

HW Addr Space		Proto Addr Space
HW Len	Proto Len	Opcode
Sender HW Address...		
Sender HW Address		Sender Proto Addr...
Sender Proto Addr		Target HW Address
Target HW Address		
Target Proto Address		

Figure 3: ARP Packet (showing Ethernet and IP addresses)

The key feature missing from ARP that is needed for address resolution in Ad-Hoc WANs is a means to link a single MAC (hardware) address with one or more subnets that are reachable via that MAC address. However, it is straightforward to simply add a list of subnet addresses to the standard ARP PDU. The usual encoding for specifying a subnet pairs a network address that belongs to that subnet with a subnet mask, a bit vector that has ones in the network address bit positions that are significant in distinguishing subnet members from non-members.

The WARRP packet format (Figure 4) accommodates a variable number of subnet specifications by appending them to a fixed-format packet similar to the ARP packet.

- A Count field in the header indicates the number of subnets appended to the packet.
- Authentication is included to prevent ARP spoofing.
- The routing functions of WARRP are implemented using the Opcode and GW (gateway) fields. The GW flag is set to 1 when a node can forward packets toward the Internet (i.e., beyond its local LANs and the ad-hoc WAN).
- As in ARP, the WARRP Opcode field distinguishes requests and responses. However, WARRP has a broader range of Opcodes than ARP, as described in the next section.

HW Addr Space		Proto Addr Space		
HW Len	Proto Len	Count	GW	Opcode
0000	Sender HW Address			
Sender Proto Addr				
0000	Target HW Address			
Target Proto Address				
Authentication...				
Subnet Address				
Subnet Mask				

-
-
-

Figure 4: WARRP Packet (showing STANAG 5066 and IP addresses)

2.2 WARRP Protocol Operation

The principal function of both ARP and WARRP is finding the MAC address to which traffic for a known network address should be sent. In ARP, this is accomplished by broadcasting a request that contains the network address of the target, with the target hardware address left blank. The sender protocol and hardware addresses are also sent in the ARP request, so that the target (and all other nodes that receive the broadcast) can record this address correspondence for future use.

ARP responses, which fill in the correct target hardware address, are normally generated by the target node, although another node can be programmed to send a proxy response on behalf of the target node if necessary.

To reduce the overhead of ARP requests and responses, nodes cache address translations from recent requests and responses. ARP cache entries are invalidated after some time, however, so that stale address translations (resulting from replacing a network interface card, for example) do not remain in the cache. Like ARP, nodes employing WARRP will also cache resolved addresses, although a different approach is used for replacing cached address and routing data.

WARRP is intended for use in ad-hoc WANs, in which a sender cannot employ subnet masking to determine whether or not the target is directly reachable via the WAN. Instead, senders broadcast WARRP requests for *any* target not found in the senders WARRP cache.

The ad-hoc nature of the network demands a flexible response mechanism that supports the normal ARP concept, along with single-relay routing and the discovery of gateways for forwarding packets beyond the local LANs and WAN. For this reason, any member of the ad-hoc WAN can generate a WARRP response to provide information that may be helpful. Table 1 lists the range of WARRP responses that may be sent, in decreasing order of value to the sender of a WARRP request.

WARRP Responses in Token-Passing Networks. The naval battlegroup WANs in which WARRP will first be implemented use a token-passing MAC protocol [3] for channel access control. In this structured, contention-free system, a WARRP Request can elicit a response from each node in the network during the current token rotation. Each node will send the most useful response (highest-numbered Opcode) from Table 1.

Table 1: WARRP Messages

Opcode	Meaning	Target Protocol Address Field	Target MAC Address Field	Count Field	Subnetworks
5	I am the target	Target Addr	Target MAC	# subnets @ responder	Responder subnets
4	I can relay to target	Target Addr	Relay MAC	# subnets @ responder	Responder subnets
3	I am a gateway (GW)	Target Addr	GW MAC	# subnets @ responder	Responder subnets
2	I can relay to a GW	Target Addr	Relay MAC	# subnets @ responder	Responder subnets
1	Local subnet report	Responder Addr	Responder MAC	# subnets @ responder	Responder subnets
0	Topology Request	Broadcast Addr	(blank)	# subnets @ requester	Requester subnets
≤ 0	WARRP Request	Target Addr	(blank)	0	

When the token completes its rotation and returns to the requester, that node can simply send its traffic to the MAC address that was in the response with the highest Opcode. Ties in Opcode should be resolved in favor of the first such response; this will tend to result in the most efficient data transfer in a token-based network, as described later.

WARRP Responses in Contention-Based Networks.

When a contention-based protocol such as IEEE 802.11 DCF [4] or DCHF [5] is used for channel access control, a burst of responses to a WARRP request would tend to congest the channel. To address this problem, the sender of a WARRP Request can restrict the range of responses that are sent by setting the Opcode field to a value less than 0. Nodes that receive a WARRP Request with an Opcode less than 0 respond only if the sum of their opcode and the received Opcode would be at least 0. Thus, a request with an Opcode of -4 indicates that only the target itself or a relay to that target should respond.

Topology Requests. An optional mechanism is included in WARRP for proactively identifying subnets that are present at network member nodes. A Topology Request lists the subnets present at the sending node, and requests similar lists from other network members. Each responding network member sends its list in a Local Subnet Report. All network members update their WARRP caches as they receive these subnet lists in both Topology Requests and Local Subnet Reports. (These self-reported subnet/MAC address pairs are exempt from aging in the WARRP cache, discussed below.)

A Topology Request might be sent, for example, when a new member joins a network. This serves to update existing members' caches with the newly available subnets, and rapidly fills the new member's cache with the subnets that are present elsewhere in the network.

Nodes sending Topology Requests and responses also advertise their ability to serve as a gateway to external networks by setting the GW flag.

WARRP Cache Aging. The cost in network bandwidth for WARRP requests will often be high, so entries in the WARRP cache

should be retained until they are invalidated by repeated communications failure. Upon such repeated failure, a new WARRP request is sent. If the target reports a new MAC address, the old entry is overwritten. If the most useful response is from a Relay node, the old cache entry is flagged as failed, and the MAC address of the relay node is appended, but the former MAC address is not

overwritten. Furthermore, a timeout is set for the Relay entry that prompts a new WARRP request after some time; if direct connectivity to the target is discovered either by receiving traffic directly from the target, or by receiving a WARRP response after the timeout, the cache entry is updated to drop the Relay MAC address.

Integration with LAN Routing Protocols. HF nodes exchange routing information (e.g., OSPF) packets with their local routers via their LAN interfaces. (Note that OSPF packets do not traverse the HF WAN.) Using this interior gateway protocol, the HF nodes learn of all local subnets from their local routers, and inform those routers of networks reachable via the HF WAN. Those distant networks (typically shipboard LANs) will be advertised as high-cost stub networks.

Note that the GW (gateway) flag in WARRP packets provides only an opaque indication that a node can serve as a gateway to external networks; it provides no information about distant networks accessible via the gateway. HF nodes store and use this information, but do not inform their local routers about gateway services at distant nodes. This ensures that a router with a SATCOM connection, for instance, will have no knowledge of HF connectivity to other SATCOM terminals, so it can never advertise connectivity to distant Internet networks via HF as a transit network.

3. DISCUSSION

In this section, we present a few examples of WARRP in use in an Ad-Hoc WAN. For these examples, we use a group of ships, labeled with the subnets they carry.

3.1 Topology Request

In Figure 5, ship 173 (i.e., the HF node with MAC address 173) has just joined the HF WAN, and broadcasts a Topology Request. Assuming that all nodes are within radio range of each other, every node in the network will receive and cache the following WARRP messages:

Opcode	MAC	GW	Subnets
0	173	0	12
1	21	1	145
1	33	0	17
1	79	0	42
1	23	1	19, 67
1	63	0	123




Figure 5: Topology Request

3.2 WARRP Request

Now, let's assume that node 173 hears the responses from all nodes except 79, and that node 79 can reach only nodes 33 and 23. If a host on ship 173 tries to send a packet to host 42.12, it will reach node 173, which will broadcast a WARRP Request naming host 42.12 as the target. In the course of the current token rotation, the following WARRP messages will be received at node 173:

Opcode	MAC	GW	Target
0: Req	173	0	42.12
3: GW	21	1	42.12
4: Relay	33	0	42.12
(no response in this slot)			
4: Relay	23	1	42.12
2: Relay to GW	63	0	42.12

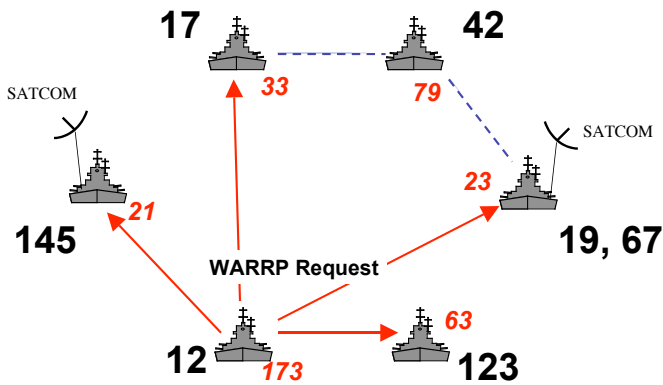


Figure 6: WARRP Request

3.3 Relayed ARQ

Since no direct link (WARRP Opcode 5) was found to the target, node 173 will send the packet via one of the nodes that offered to relay to the target: either node 33 or node 23. Choosing the earlier of the two (node 33) results in choosing a relay that precedes the destination in the token rotation order, and allows the packet to be sent indirectly in a single token rotation. Link-layer acknowledgements are returned link by link as each node receives the data, so complete data and ack exchanges are possible in every token rotation.

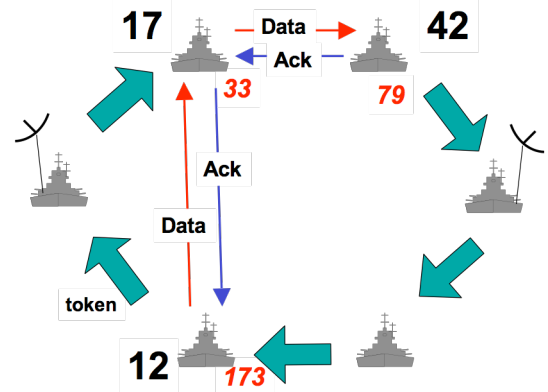


Figure 7: Relayed ARQ

4. CONCLUSIONS

The ad-hoc WAN idea is interesting, but not especially attractive. We would usually prefer to create a wireless LAN in which nodes are dynamically assigned subnet addresses, and ARP is used to resolve addresses. Existing wireless routing protocols could then be employed to manage indirect routes through the network as required. However, in special circumstances this is not possible, and WARRP was developed to integrate efficiently the address resolution and routing needs of such networks.

The proposed routing mechanism for Ad-Hoc WANs relies upon a variation of the Internet Address Resolution Protocol (ARP) for identifying reachable networks as well as resolving IP subnet-to-MAC (e.g., STANAG 5066) address correspondence. This Wireless Address Resolution and Routing Protocol (WARRP) runs within the HF hosts.

The unique features of WARRP-based routing are as follows:

- The IP addresses of hosts reachable via an Ad-Hoc WAN cannot be determined by applying a subnet mask to the IP address of the sending node. Therefore, nodes may send WARRP requests for any destination IP address.
- WARRP packets can include a list of subnets in addition to the usual ARP fields. All IP addresses in the subnetwork range obtained by applying a mask to a Protocol Address are reachable at the corresponding MAC address.
- Non-contiguous IP address ranges at a node may be announced in a single WARRP packet.
- WARRP requests and responses are always sent to the broadcast MAC address so that the sending node subnet/MAC address correspondence is announced to all Ad-Hoc WAN nodes.
- Authentication is included in each WARRP packet.
- A field in the WARRP packet indicates broadband connectivity to the Internet at the sending node. Such nodes are candidates for service as default gateway for an Ad-Hoc WAN
- WARRP supports single-relay routing to work around link outages in Ad-Hoc WANs.

REFERENCES

1. E.E. Johnson, R.I Desourdis, and M Rager, "Simulation of MIL-STD-187-721C Automated HF Networking," *Proceedings of 1996 Ionospheric Effects Symposium*, Alexandria, VA, 1996..
2. NATO Standardization Agreement 5066: *Profile for High Frequency (HF) Radio Data Communications*, version 1.2, NATO Standardization Activity reference 0114-C3/5066, 27 January 2004.
3. E.E. Johnson *et al*, "Robust Token Management For Unreliable Networks," *Proceedings of MILCOM 2003*, Boston, MA, 2003.
4. ANSI/IEEE Std 802.11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 1999.
5. E.E. Johnson, M. Balakrishnan, and Z. Tang, "Impact Of Turnaround Time On Wireless MAC Protocols," *Proceedings of MILCOM 2003*, Boston, MA, 2003.