

Ad Hoc Networking:
Published in
'Systems', June 1999, pp 33-40.

Carlo Kopp, BE(Hons), MSc, PhD, PEng,
Defence Analyst and Consulting Engineer

Email: Carlo.Kopp@aus.net

© 1999, 2002, Carlo Kopp

February 19, 2002

Background Article

Systems, June, 1999

Disclaimer:

This document was compiled wholly from available public domain sources.

Distribution:

Permission is granted for distribution free of charge for educational purposes. Distribution for other purposes is not approved without the prior written consent of the author.

The Author:

Carlo Kopp, BE(Hons), MSc(Comp.Sci), PhD(Monash) is a former computer design engineer, embedded programmer, Unix systems programmer and a practising Unix systems consultant, with over 15 years of industry experience.

Revision Status:

First published in Systems, June 1999, pp 33-40.

`$Id: AdHocNetworks.tex,v 1.2 2002/02/19 21:54:50 carlo Exp carlo $`

Contents

1 Ad Hoc Networking	4
----------------------------	----------

1 Ad Hoc Networking

Ad Hoc mobile networking is the uncharted frontier of contemporary networking technology, and an research area which at this time is being heavily funded in the US. In essence, Ad Hoc networking is all about providing connectivity between mobile nodes, which have no supporting connections to the fixed networking infrastructure.

In an Ad Hoc mobile network, every node in the network carries its own router with it, and all nodes cooperate in carrying traffic.

The whole philosophy of the Ad Hoc networking model is a radical departure from the highly structured, and frequently hierarchical models employed for both local area and wide area networking, currently in use.

In the established, fixed infrastructure model, the routers and supporting functions such as name resolution are all embedded within the networking infrastructure. Communication between a pair of nodes requires that the nodes hand the traffic over to the routers, which then forward the traffic over multiple router hops until it arrives at the destination router, which then passes the traffic to the recipient node.

In this model, with the exception of the odd host tasked with acting as a router, mostly the routing function is performed by the network, and the nodes are essentially clients of the 'connectivity service' provided by the networking infrastructure. The model is hierarchical, insofar as traffic from small cells or subnets is concentrated as it flows up into the network, which aggregates the traffic associated with multiple virtual circuits or datagram connections, and carries it across specific point to point communications links to the geographical area within which the destination (or source) nodes are situated.

This paradigm has generally served us well, and indeed is the essence of the 'catenet' model used in ARPANET and now the Internet.

Because the topology of the network is unchanging, or changes at a very slow rate, mechanisms for name mapping, route discovery, and route maintenance can be very lethargic in their response times. Indeed, manual reconfiguration of the routing topology is commonly employed.

An Ad Hoc mobile network is essentially incompatible with this basic model, since it is highly time variant in topology.

At this point it is worth digressing into the practical, application oriented, aspects of the Ad Hoc network, since the technology promises many services which have hitherto been inconceivable.

The simplest Ad Hoc network can be envisaged as a wireless radio network between a collection of vehicles, ships, aircraft, or even people on foot, operating in a geographical area with no networking infrastructure. Many examples of such scenarios come to mind. A fleet of fishing vessels searching for schools of fish on the high seas, a seismic survey team in a remote area, a disaster relief operation, or aid operation, trying to function in an area which has been stripped by a natural disaster of its communications infrastructure, or if in the Third World, never had one in the first place. Scientists on field outings, or indeed even a class of school-children on an outing into a national park, all carrying laptops or wearables. Cars and trucks on country highways or freeways, with onboard Internet connectivity.

There are, of course, also a myriad of military applications involving the networking of aircraft, helicopters, tanks, ships, and even infantrymen with wearable computers.

The range of possible situations in which Ad Hoc networking can be exploited is huge, and this is not an understatement by any measure. A robust Ad Hoc networking scheme frees the individual from the geographical constraints of the fixed network. In this respect it is fundamentally different from established mobile networking, in which mobile nodes are tied down by the need to remain within the coverage of a wireless hub, connected to the fixed network infrastructure.

An Ad Hoc network, by its nature, provides mutual connectivity between cooperating peer nodes. Nodes which cannot directly communicate, are assisted by other nodes between which connectivity exists, and which can connect to the end nodes which intend to communicate. Therefore, every node in an Ad Hoc network must have the capability to perform as a router if its peers require it to do so.

Mutual connectivity does not imply the ability to access the fixed infrastructure, and if connectivity to the fixed infrastructure is required, then at least one node in the Ad Hoc network must have the ability to connect to the fixed infrastructure and carry traffic into and out of the Ad Hoc network.

Connectivity between nodes in an Ad Hoc network can be provided by any channel which can carry traffic, and existing research encompasses everything from short wave radio, VHF, UHF, omnidirectional and directional microwave, and omnidirectional incoherent infrared or directional infrared and visible band laser technology. Any scheme which can carry a digital modulation without using a cable is a candidate for Ad Hoc networking applications, indeed the ideal Ad Hoc network doesn't really concern itself with the physical layer channel employed to carry ones and zeroes between participating nodes. Moreover, with proper protocol transparency, it need not be concerned with the type of traffic it carries, other than with the issue of the amount of bandwidth required for a given service.

We could envisage a situation in which you are travelling in your company car down a country highway, and you receive a 'phone' or 'video-phone' call from the office, which has hopped along a series of other cars and trucks along the highway, from the vehicle nearest to a fixed hub, in the nearest country town. The managing director has decided to change your assignment and you have to spend another three days in the middle of nowhere...

Bitten by a Taipan while bush-walking in North Queensland, you make a call on your wearable computer/communicator, which hops to a geologist's wearable, which hops to his 4WD, which in turn hops across a couple of cars on the highway, several miles away, and then hops to a fishing boat several miles off the coast, which happens to have a satcom link via which the message is hopped to a ground station in the vessel's home port, fifty kilometres away, from where the message is routed to an emergency service dispatcher. Since your wearable has a GPS receiver on it, the rescue helicopter is scrambled and within twenty minutes you are enroute to a hospital.

What mature and robust Ad Hoc networking offers is virtually universal connectivity, limited only by the link performance and routing delays of the participating nodes, and their connectivity to the established fixed network.

Too good to be true ? Perhaps, but certainly, as fanciful as these examples may be, they are all well within the bounds of today's technology, providing that suitable Ad Hoc routing protocols exist and are implemented.

The technological challenges of Ad Hoc routing are very much non-trivial !

The first 'package' of problems derive from the continuously varying topology, and potential throughput, of the Ad Hoc network.

Topology varies simply because some nodes will move in and out of wireless link range of other nodes in the network, be it through distance, or concealment behind terrain or other obstacles which prevent transmission, such as inclement weather or rain which soaks up microwaves and laser beams.

Network throughput will vary for two reasons. The first, and obvious reason, is that the larger the number of hops your traffic has to travel across to get to where it is going, the greater the routing delays you incur, which cumulatively add up to increase the latency of your link, and thus potential throughput, for a finite buffer size in participating nodes. Since the network topology is continuously changing, frequently in an unpredictable manner, the number of hops between you and your destination node will also vary. This has other implications we will discuss later.

The second reason why network throughput will vary is a consequence of Shannon's information theory, since for a constant power output and receiver sensitivity, as distance increases between

two wireless nodes, the signal/noise ratio declines and thus achievable link bit rate drops. Therefore, as the signal weakens, the range of potentially available services declines, or the bit error rate increases. This variation of throughput with time is usually referred to as 'fading', as much as this term is used and abused in various niches of communications theory.

At this time very few protocols for MAC layer connections exist which can adaptively adjust their throughput to accommodate variations in link performance. We are seeing the first steps with the IEEE 802.11 wireless Ethernet, where link quality degradation forces a reduction in link bit rate, albeit in large and discrete chunks. We have yet to see a genuinely robust protocol which dynamically 'rubberbands' the bit rate through the channel to achieve a desired balance of speed and bit error rate. In wireless networking, where power and bandwidth come at a big premium, every snippet of usable bit rate is valuable.

By far the biggest problem in current Ad Hoc networking research is that of routing in a situation where the topology of the network is changing continuously, somewhere within the network.

Static networks mostly use either Distance Vector (DV) or Link State routing algorithms, neither of which are spectacularly well suited to highly dynamic topologies.

Distance Vector, or Bellman-Ford schemes, such as those used in the DARPA packet radio protocol, RIP, XNS or IPX, are based on the idea of periodically broadcast tables of distances, typically in hops, between a node and all possible destinations. A necessary requirement is that the update rate is greater than the rate of topology change.

Link State schemes, such as those used in the OIS IS-IS or Internet OSPF routing protocols, rely on the broadcast of complete topology maps for the network.

In a highly dynamic wireless network, such protocols run into a number of difficulties:

- topologies may be highly redundant, with some nodes being in the situation of being able to connect to a very large number of neighbours, while others see very few neighbours.
- bandwidth is scarce and cannot be wasted
- battery power on portable equipment is a finite resource which cannot be wasted
- high rates of topology change require high update rates

Link State, and Distance Vector routing schemes fall foul of these issues since they distribute a lot of routing information, and with high rates of topology change this will eat into bandwidth and thus battery power, moreso in highly redundant topologies, where much of the

information is effectively wasted. Maintaining a current routing table on a node which does not communicate much with its neighbours is a drain on critical resources for no return.

From a perspective view, the routing problem really decomposes into two problems. One is that of 'route discovery', the other is that of 'route maintenance' whereby the validity of discovered routing information is maintained.

Topologies in Ad Hoc networking are an issue within themselves. In essence, nodes may move around with no clear geometrical interrelationship, or may form clusters associated with groups of individuals or vehicles moving around in relatively close mutual proximity, or nodes may also form 'linear topologies', when vehicles travel down roads, railways, shipping lanes or air routes.

In a general sense, routing models are either based upon a 'Flat Topology' model, or a 'Hierarchical Topology' model. In the former, all nodes are peers, in the latter one node within a cluster gathers traffic on behalf of 'lesser' nodes in the cluster, and is responsible for carrying this traffic in and out of the cluster. The Hierarchical Topology is in many respects an offshoot of the static networking model, and has generally not been a popular research area, since it can result in less than optimal routing behaviour. Flat Topologies are in most situations the best approach, since they can provide for redundant paths in and out of cells, and should protocol support exist, load balancing across multiple links.

Routing models can also be divided in yet another way:

'Proactive Routing' is any scheme which continuously monitors the topology and maintains current routing tables regardless of instantaneous demand. DV and LS schemes fall into this category. While routing information is always available for a sender, the network is being continuously flooded with routing management traffic, much of which is unused.

'Reactive Routing' is any scheme where routing information is gathered only on demand. In such schemes, a route is discovered only when needed, and thus routing management traffic is kept to its bare minimum. Reactive schemes have been most popular to date, since they minimise the route management traffic overheads.

At this time research effort in Ad Hoc networking is mostly confined to two large research projects, both centred in the US. The IETF is sponsoring the MANET (Mobile Adhoc NETWORKING) project, while the US DoD's DARPA is funding the GLOMO (GLOBAL MOBILITY) program.

MANET is aimed primarily at extending the existing Internet protocol suite to accommodate Ad Hoc routing functions, within the context of IP connectivity. The physical channels are only of peripheral interest, in how they constrain the behaviour of the network.

GLOMO has a much broader scope, and is aimed primarily at developing the technology base

for universal, transparent, connectivity between military platforms, and troops, in the absence of a fixed infrastructure. Therefore a large component of GLOMO research effort delves into areas other than routing alone, such as wireless link and antenna design, and the topology is not limited to flat topology models alone. In this sense, Ad Hoc routing problems are only a component of the GLOMO effort.

The complexity of the Ad Hoc routing problem is reflected in the volume of research currently being conducted, no less than six different schemes are being researched.

The Zone Routing Protocol (ZRP), proposed by Haas, employs a proactive route discovery scheme within a local 'zone' immediate to a node, but uses a reactive scheme for connections outside the 'zone'.

The Destination Sequenced Distance Vector (DSDV) protocol, essentially a variant of the RFC 1058 RIP protocol, as proposed by Perkins and Bhagwat, uses elements of the well established RIP protocol, but adds sequence numbers to routing tables to eliminate routing loops, and uses triggered updates to propagate topology changes when these are discovered, in addition to RIP like periodic updates. DSDV is loop free, and is designed to respond quickly when changes in topology occur in between periodic update cycles. It is a proactive routing protocol with some reactive features.

The Dynamic Source Routing (DSR) protocol, proposed by Johnson, extends the source route model used in the current IP suite, and is a purely reactive protocol. Every packet contains an ordered list of intermediate routing nodes, every node maintains a route cache, and if a route does not exist in the cache, a 'route request' packet is broadcast and propagated along until it hits the destination, or a node which knows of the destination, upon which a reply packet is sent to the requesting node. Intermediate nodes add their address along the way, and update their caches with eavesdropped routes. Routes are maintained by watching for lost packets, upon which another route discovery must be performed. The DSR model is an extension of the route discovery scheme in the RFC 2002 mobile IP protocol, based in turn upon the existing RFC 791 Loose Source Record & Routing protocol.

The Temporally Ordered Routing Algorithm (TORA), proposed by Park and Corson, is another reactive route discovery protocol, which uses a 'link reversal' model in route discovery. A route discovery query is broadcast and propagates out through the network until it hits a destination or a node which knows a route to the destination. A parameter, termed 'height', which is a measure of the responding node's distance to the sought destination node, is then returned to the querying node. As the query response propagates back, each intermediate node updates its TORA table with the route and 'height' to the destination node. The querying node then uses the 'height' to select the best route. TORA has an interesting property that it frequently chooses the most convenient route, rather than the shortest route, and in doing so attempts to minimise the routing management traffic overhead.

The Ad hoc On Demand Distance Vector (AODV) protocol, proposed by Perkins, blends elements of the DSR and DSDV protocols, using the DSR reactive route discovery and maintenance models, in combination with the sequence number and periodic update features of the DSDV protocol.

An extensive simulation comparing DSDV, DSR, TORA and AODV was conducted at CMU and results published late last year. The conclusions from these simulations were by all means very interesting.

The simulation compared two important metrics, the ratio of delivered packets to sent packets, and the routing management overheads.

Highly dynamic networks, where most nodes are continuously moving, are the limiting situation for an Ad Hoc network and by far the most demanding scenario to be handled.

Under these conditions, the DSR and AODV protocols delivered 98% of packets, while the DSDV protocol foundered with a 70-75% delivery ratio. A 2% packet loss is respectable performance for the environment in question, and if we compare the packet loss ratio, then DSDV lost about 12-15 times as many packets in comparison with DSR and AODV.

Routing overhead performance was also interesting, and DSR won yet again. Interestingly, the DSDV protocol was the least sensitive to the mobility of nodes, requiring an almost constant overhead regardless of node mobility. In comparison, DSR incurred 1/2 the overhead of DSDV, while AODV incurred triple the overhead, and TORA four times the overhead of DSDV. Between the best and worst, this is an eightfold difference in overhead traffic.

Another interesting metric was an analysis of how sensitive protocol performance was to increasing network activity, measured by the number of active nodes.

Packet delivery performance was again best for AODV and DSR, with a degradation of about 1% with a tripling of load. The same tripling of load dragged TORA down to about 40% delivery ratio, making it virtually dysfunctional.

Routing overhead traffic variations with load showed that DSDV was the most robust, with a virtually constant overhead. DSR and AODV displayed similar behaviour, with overhead growing roughly in proportion to load. TORA performed poorly, with a tripling of load resulting in 25 times the routing traffic overhead.

In terms of which protocols did best in choosing optimal routes, DSR and DSDV did best.

It is perhaps a tribute to Occam, that ostensibly the simplest of the protocols, DSR, appeared to deliver the best all round performance. It also demonstrates a strong case for a reactive vs a proactive route discovery scheme.

Other variations upon these schemes are also being explored. The Location Aided Routing (LAR) protocol extends a DSR or AODV protocol, by using GPS satellite navigation data to limit the route discovery search to nodes which are known to be relevant.

In this context it is worth mentioning the work of Imielinski, who is proposing the inclusion of GPS geographical location data into IP addressing, to provide an alternate addressing scheme for generic IP routing. As yet this idea does not appear to have been incorporated into the MANET or GLOMO protocols.

Perhaps the best question to ask at this time, is 'Quo Vadis, Ad Hoc Networking?', pun intended. Clearly the research effort is beginning to converge and we can expect to see a final protocol choice fairly soon. A good guess is that it will be a variant of DSR/AODV, crossed with the source route based Mobile IP protocol.

In the longer term, the protocols devised for MANET may be the solution to the growing unreliability of the fixed Internet infrastructure. As its complexity continues to increase, the sheer number of routing nodes against the failure rate or crash rate of a typical routing node means that beyond some point, the rate of topology changes through router crashes will get ahead of the topology update rates accommodated by existing link state and distance vector algorithms such as OSPF and RIP. Adoption of a protocol derived from the MANET effort would defeat this problem.

Ad Hoc networking therefore promises much more than the obvious.