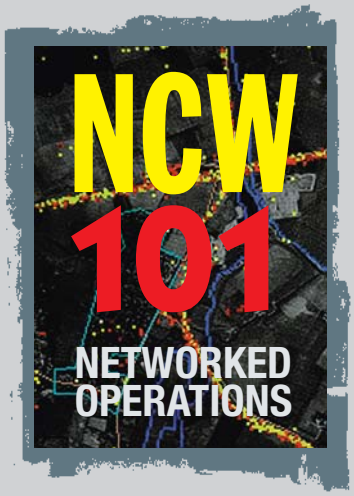

Ad Hoc Networking

NCW 101 part 4

Dr Carlo Kopp

Networks are becoming a shared and systemic single point of failure for modern combat forces. As more systems are networked, more systems become dependent upon the network to perform their tasks. Take away the network, and chaos is likely to rapidly ensue in the absence of alternative channels for gathering and distributing information. Opponents of networking have seized upon this idea, yet an emerging technology will very soon change much of the vulnerability inherent in current networking technologies. This emerging technology is the ad hoc network.

Ad hoc networks are designed to be self forming, self healing, distributed and lack any centralised control facilities to target. They are the technology underpinning the Joint Tactical Radio System (JTRS) Wideband Networking Waveform (WNW) and Tactical Targeting Network Technology (TTNT), intended as replacements for the legacy JTIDS/MIDS system. Much more resilient than legacy networks, ad hoc networks will over time displace all established technologies in military networking.

The idea of an ad hoc network emerged during the 1970s, with numerous DARPA-sponsored studies and trials of 'packet radio' systems leading to the Packet Radio Network (PRnet) in 1972, and later the Survivable Radio Network (SURAN) and Low-cost Packet Radio (LPR) efforts during the 1980s. During the 1990s ad hoc networks emerged in the commercial sector, as a byproduct of cheap radio-frequency wireless interfaces entering the commodity computing market.

Established network architectures, such as the JTIDS/MIDS system and the global Internet using TCP/IP suite protocols, are highly structured. JTIDS/MIDS (refer NCW101 Part 3) relies on a precisely controlled time slot system, requiring synchronisation. The Internet protocols rely upon a hierarchical and highly structured tree like topology of connections, in which many critical services such as name-to-address mapping are essentially centralised.

To best appreciate how ad hoc networks function, and how they acquire their resilience, it is necessary to explore the established Internet model in more detail.

The Internet is the evolutionary offspring of research conducted during the early 1960s. This research established that it was feasible to create computer networks spanning global geographical footprints, by using a technique called 'packet routing', devised by then PhD student Leonard Kleinrock, at MIT. The idea was to take a communication between two computers, chop it into small pieces termed 'packets', encapsulate these packets with addressing and other information, and then have these packets flow through a network of devices termed 'routers'. Routers were computers that had multiple communications interfaces, the sole purpose of which was to accept incoming packets, decode the addressing information, and send them on their way using the appropriate communications interface. Each router contained a map of the network topology - the specific manner in which routers were interconnected in the network - and using this map and the addressing information, it could determine exactly which communications interface it needed to use to get a packet travelling in the right direction to get to its intended destination.

The idea of a packet network was very powerful, and led to the definition of the 'catenet' model of networks, the basis of today's Internet. The US DoD generously funded this research, via DARPA, and this led to ARPANET, and eventually, the Internet we love and know.

The Internet, and all similar proprietary networking schemes (most of which have long become extinct), use a defacto hierarchical topological model for interconnecting individual nodes in the network, refer Figure 1. In this model, a multiplicity of computers at any given site are connected to a router, for instance via a Local Area Network or hardwired individual connections (the 1970s approach). If any of these computers intends to communicate with a peer at another site, the packets it sends must traverse multiple 'hops' between routers until they reach their intended destination.

An example might be a computer at Site B communicating with a computer at Site D. The packets travelling in either direction must traverse hops between Routers B, G, F and D. If we imagine a physical network, then the hop between B and G might be a high speed serial link over a copper cable, the hop between F and D might be a high speed optical fibre link, and the hop between G and F a microwave link with dish antennas on towers.

Land, sea and air forces depend upon effective networks for command and control, tactical manoeuvre and coordinated action. (Defence)

This specific - and trivial - example is hierarchical and highly structured. The uppermost level in the hierarchy is occupied by Routers G and F, the lowermost level by routers A, B, C, D and E. All connections forming the network topology are known beforehand. The Internet as we know it today follows this model, but immensely more complex with millions of computers and routers connected to form its topology.

This model was devised for fixed infrastructure networks, using mostly copper and optical fibre connections throughout. Wireless networks using the 802.11 protocols, now increasingly a ubiquitous feature in many portable computers and devices, emulate this model. In such wireless networks, a radio-frequency link in the 900 MHz, 2.45 GHz or 5.8 GHz band is used to connect the computer to a local network, and usually the lowermost router in the local hierarchy. In effect such wireless networks replace the last cabled connection between the computer and the network. Wireless it may well be, but otherwise it is little more than an extension of the fixed infrastructure network.

To better appreciate the unsuitability of the hierarchical model for highly mobile wireless networks, such as are required for military networks, it is necessary to delve a little deeper into the workings of the fixed infrastructure model. A key issue in any such network is that of how to tell every router in the network where it should send packets in order to reach a specific destination. Effectively this is the problem of addressing.

In the Internet scheme, we are familiar with network names in a text form, such as info@defencenews.com.au (email) or www.defencenews.com.au (WWW). If we wish to send an email or access a website, our computer must first query the network to get a network address, produced by translating the name into an Internet Protocol or IP address, in this instance 202.148.146.201. Next, our computer must open a stream connection to the machine addressed as 202.148.146.201. As the first packet is sent, each router along the way looks up this address, to determine which direction to send the packet in. Eventually, the packet arrives at 202.148.146.201, which responds accordingly. This is a gross simplification, but important to explain how the network functions.

The reality is more complex, as every query to discover a new name to address mapping must be directed to a 'Domain Name Server' - a computer running DNS software which maintains a directory of name to address mappings. As no single computer could realistically cope with such a number of queries, the Internet uses a very large number of redundant DNS servers, all organised in a virtual hierarchy. If a DNS server does not know a specific mapping, it queries the server above it in this hierarchy, and so on, until a server is found which knows the answer. Suffice to say, despite redundancy, DNS is a single point of failure. Several major Internet outages in Australia over the last decade resulted from failures in key DNS servers.

Once the DNS server provides the computer with the required IP address, it can initiate a connection via its nearest router. That router has to understand enough of the topology of the network to know where to direct the connection. Again, we see further hidden complexity in this 'route discovery' process. Routers maintain what are termed 'routing tables' which contain mappings between addresses and specific interfaces. Large routers may indeed have dozens of interfaces to other routers, and this information must be continuously maintained for the whole network.

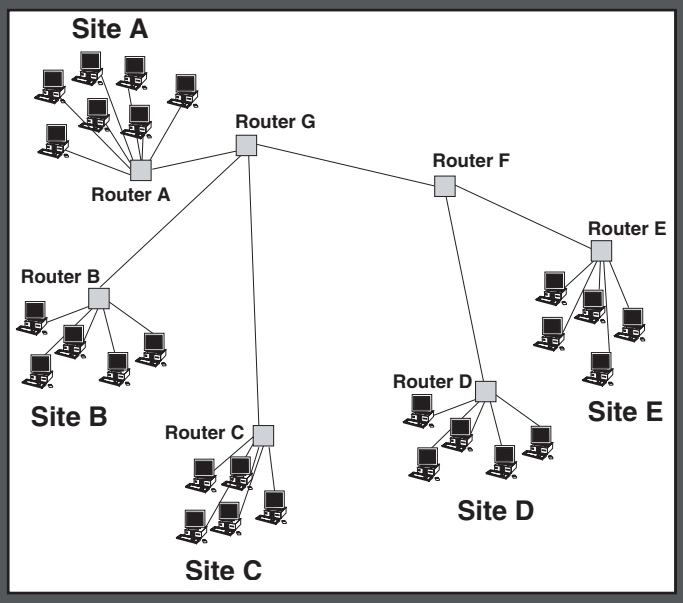

Figure 1

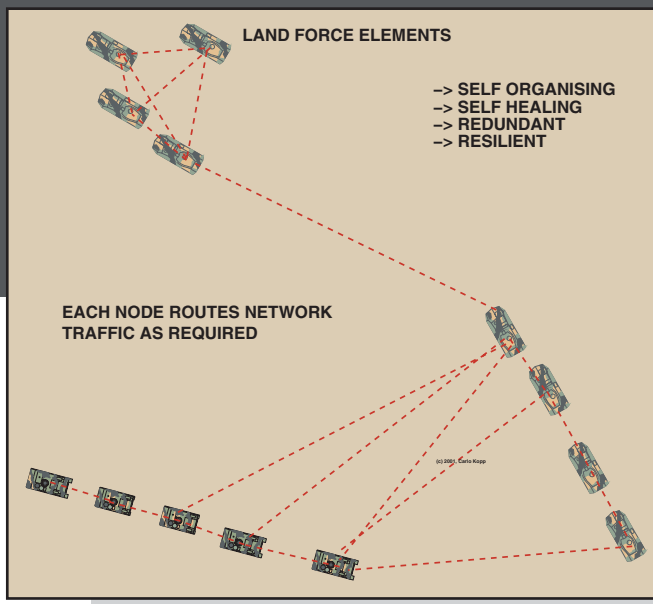

Figure 2 - Example of Ad Hoc Network

To distribute and manage this information, routers have to communicate with one another using specific protocols. The most widely used of the older protocols is Routing Information Protocol (RIP), with Open Shortest Path First (OSPF) being preferred in newer systems. When a router receives a packet, it uses the cached information gathered with the protocol to calculate what it believes is the best route to the destination. Suffice to say, there is further complexity in this mechanism that is too extensive to discuss here.

What we thus observe in the Internet is a very powerful networking medium, but one that is ill-adapted to an environment in which the connectivity through the network, and thus its topology, rapidly change in time. That is the reality of wireless military networks, and the reason Internet connectivity is not common in current military systems.

Modern ad hoc networking protocols emerged to bridge this gap during the 1990s. The central idea behind all ad hoc networks is that there is no fixed topology, or the topology is dynamic and rapidly changing. This is a model eminently suited to military networks. Refer Figure 2.

In an ad hoc network, every computer has a router attached to it - either embedded in software, or as a router/radio-frequency modem. We describe each such arrangement as a 'node' in the ad hoc network.

In such networks, every node routes if required traffic to and from its peers in the network, which is for all intents and purposes a cooperative network. This is an important distinction from conventional networks, where routers are specialised and most computers do not double up as routers.

The key to the function of all ad hoc networks is the performance of the route discovery protocol in use. Route discovery protocols for ad hoc networks differ considerably from route discovery protocols used in conventional fixed networks. In an ad hoc network, every time a connection is to be established, a node must send out a query which asks 'is a connection to my destination node available, and if so, what routes exist to get there?' In general, research indicates that 'reactive' route discovery techniques, which propagate a query across the network every time a connection is needed, work better than 'proactive' techniques that attempt to maintain a constantly updated table of possible connections. This is for a variety of reasons, but especially since traffic between nodes in close mutual proximity is often dominant.

One of the most popular techniques used is the Dynamic Source Routing (DSR) protocol, proposed by Johnson, which extends the source routing

model used in the current Internet protocol suite, and is a purely reactive protocol. Every packet contains an ordered list of intermediate routing nodes, every node maintains a route cache (a table in memory), 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 along the path add their address along the way, and update their own 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.

Suffice to say, DSR is one of many techniques, which include Zone Routing Protocol (ZRP), Destination Sequenced Distance Vector (DSDV), Temporally Ordered Routing Algorithm (TORA), Ad hoc On Demand Distance Vector (AODV) and others, including various hybrids.

The Zone Routing Protocol (ZRP), proposed by Haas, employs a proactive route discovery scheme within a local 'zone' in close proximity 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 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 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 that 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.

This menagerie of choices in protocols reflects the ugly reality that ad hoc networks rely on what are essentially transient connections between nodes forming the network. Since nothing is permanent, the network must be designed from the outset to cope with continual changes in possible connections, and thus possible or available topology.

For any given collection of nodes forming a network, there is no guarantee that every node will be connected to the network at all times, or that any network itself will at any time not split into more than one smaller networks.

This is a consequence of the realities of wireless link propagation through the atmosphere and low altitude propagation behaviour - termed fading in the radio frequency engineering world - resulting from complex terrain.

In the broadest terms we can divide ad hoc networks into airborne ad hoc networks, refer Figure 3, between nodes which are airborne, and surface based ad hoc networks, for instances between vehicles or warships.

In airborne ad hoc networks, possible connections between nodes, each carried in an aircraft, are limited by the range of the radio (or laser) datalink used, radio propagation for the wavelength in use through weather, and the curvature of the earth. For aircraft at the tropopause, ranges can be as great as hundreds of kilometres.

This contrasts starkly with the reality observed in surface based ad hoc networks, especially land based. Since radio (or laser) links mostly do not perform once line of sight is lost between two communicating systems, the rate of topology change in such networks is much higher than in the airborne equivalent.

Radio propagation for land based or low altitude airborne links is always fraught with difficulties, due to signal reflecting and scattering off terrain, resulting in fading effects. In this respect an ad hoc network suffers the same impediments as other networks, and resistance to multipath fading will be largely driven by th type of radio modulation used, and the operating wavelength of the radio link.

Where ad hoc networks differ in this environment from more traditional alternatives is that fading effects are manifested as link dropouts between specific nodes, and will thus be reacted to by the route discovery protocol used. As a result if fading causes a specific connection to be lost, if any other path to the node in question still exists, via other connections, the network will find it and attempt to maintain connectivity. This is quite different from a more conventional scheme where loss of connectivity to routing node isolates the platform from the network. In this sense ad hoc networks offer potentially much greater resilience.

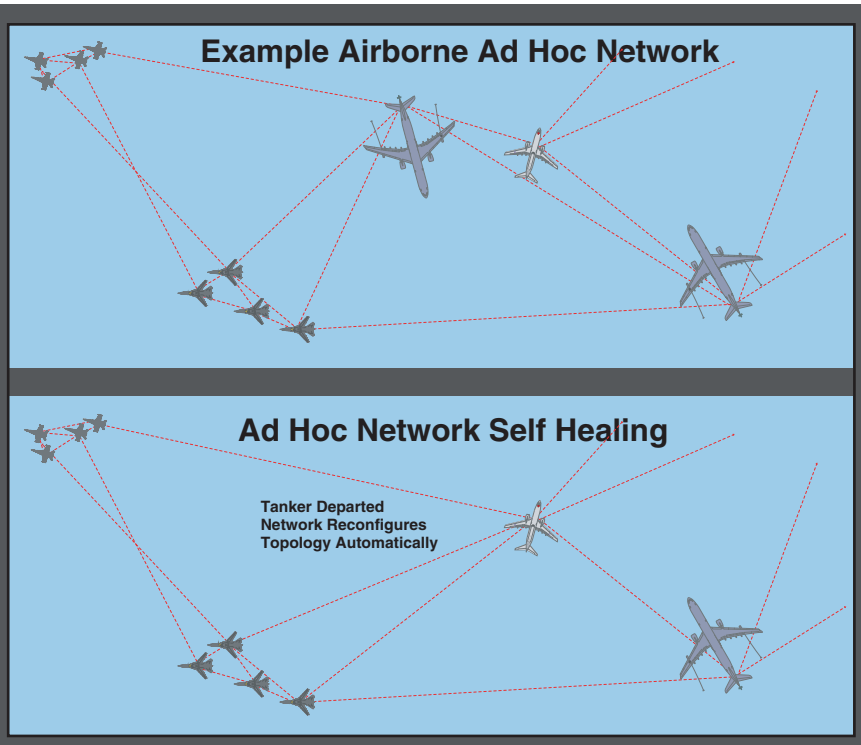

Figure 3

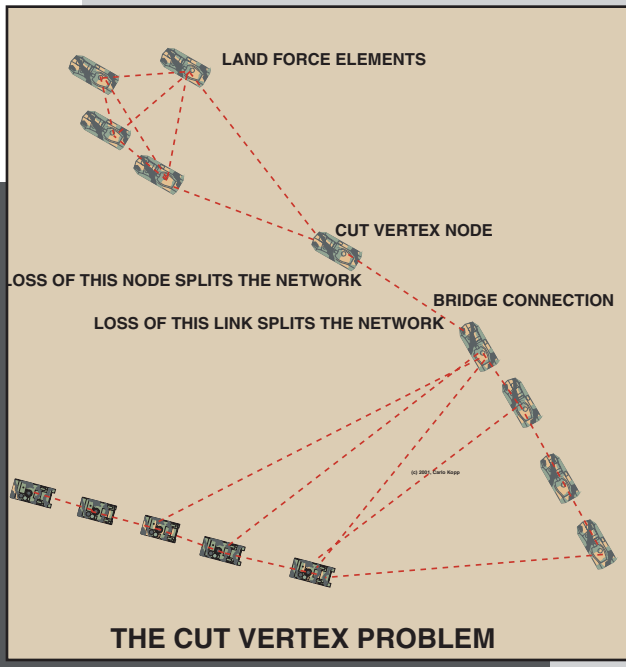

Figure 4

Like all networks, ad hoc networks are subject to the same constraints of graph theory. The particular problem we are interested in is that of the 'cut vertex' or 'bridge', where a single node or connection is the only path between two parts of the network. If it is lost, the network is separated into two smaller networks. This has proven to be an ongoing issue with commercial networks, since the location of nodes cannot be easily controlled. In military networks, this changes since it is possible to command specific platforms to loiter if feasible to maintain connectivity. An example might be the placement of a smart tanker aircraft orbit, or the spacing between convoys of land vehicles, to ensure that connectivity is maintained continuously.

The biggest problems will however arise with land force components, due to the difficulties in propagation, especially for a rapidly advancing manoeuvre force. Areas well behind the FEBA, where saturated with convoys providing resupply of fuel, ammunition and other consumables, provide good opportunities for connectivity if each convoy has several vehicles equipped as networking nodes. In effect the convoys become an advancing chain of network routing nodes, between the staging area and FEBA.

The greater headache arises with manoeuvre elements comprising rapidly advancing armour and nap of the earth helicopters, as these may advance well ahead of the supporting assets and lose clear line of sight, especially if terrain is complex or not amenable to good radio propagation.

The US approach to this, embodied in the JTRS model, is to use common and compatible equipment on Air Force, Army and Marine Corps platforms. As a result, when a direct connection is lost, the network can find routes via the airborne platforms flying overhead. Where for instance a fighter is tasked with killbox interdiction or close air support, it will be orbiting in proximity to the advance ground force elements and provide a path to route traffic over terrain obstacles to defeat radio propagation limitations.

UAVs are also good candidates for coverage extension, with the important caveat that the networking equipment and associated antennas may displace much of the payload of a smaller UAV. Piggybacking such equipment on larger UAVs should be actively encouraged, both to exploit the potential of the UAV as an airborne repeater, and to make the output from its sensors more readily accessible to ground based assets.

What does fall out of this is that for best effect, any system using an ad hoc networking scheme must have as many platforms as possible equipped with networking equipment. The additional resilience of the ad hoc network comes at the price of needing to provide more platforms with networking equipment.

The current state of modern ad hoc networking has advanced considerably since its conception during the mid 1990s.

The largest recent single package of research funding for ad hoc protocol development was provided by DARPA during the mid 1990s, under the GLOMO (GLObal MObility) program. The aim of GLOMO was to provide 'blue sky' research funding to investigate a wide range of theoretic and practical issues. Many of the simulation tools funded by GLOMO are now used globally for university and defence research and development in networking.

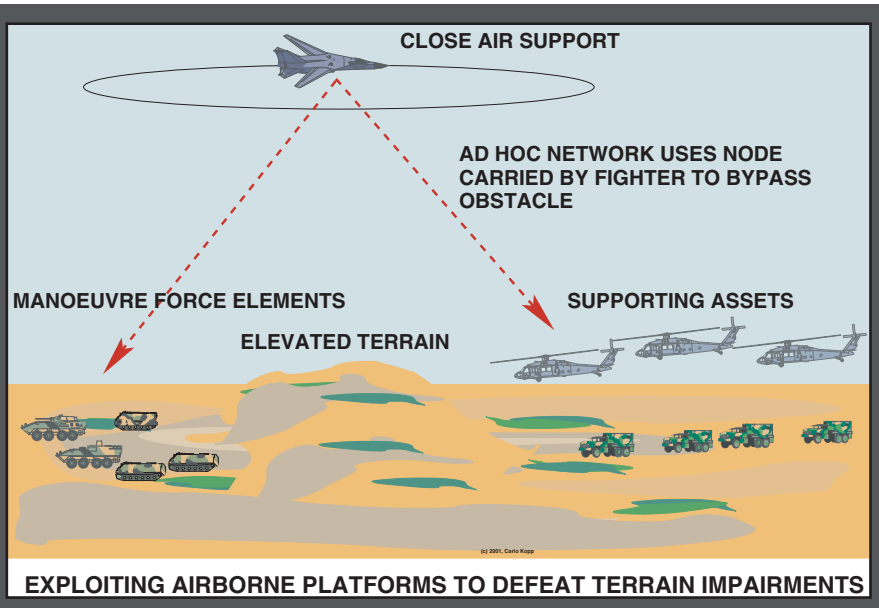

Figure 5

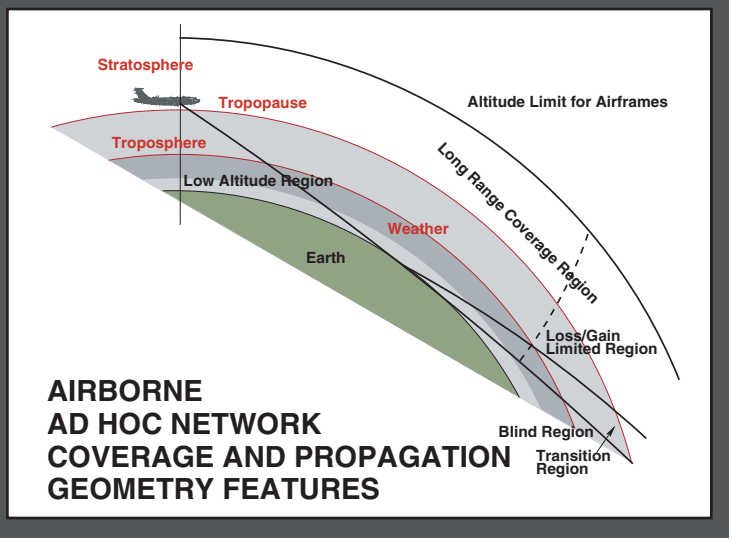

Figure 6

In the commercial domain, the Internet Engineering Task Force (IETF) has been running the Mobile Ad-hoc Networks (Manet at http://www.ietf.org/html.charters/manet-charter.html) effort to research and develop protocols intended to expand the Internet protocol suite into the ad hoc network domain. The effort is well advanced, with route discovery and management protocols now defined as draft RFC standards.

In the military domain, JTRS remains the flagship in ad hoc networking protocols, and is now approaching deployment and operational test of early network terminals. The experimental Near Term Digital Radio (NTDR) program was devised to precede JTRS and fill a capability gap - NTDR is reported deployed with three Stryker Brigade Combat Teams.

In Australia, only two universities are highly active in ad hoc networking research. Monash University in Melbourne launched its research in ad hoc networking in 1997, and recently reached a milestone with its first prototype implementation of an ad hoc networking protocol, designed to provide broadband services for low density suburban and rural environments. The protocol is architected to provide for high levels of cryptographic security, resilience to denial of service attacks, very tight access controls, and has potential applications other than commercial. The University of Wollongong is developing a much less ambitious system for use in outback communities. The University of South Australia's ITR group is also active in this area, as are a number of other universities. DSTO have provided some funding to ITR, but do not appear to have a major commitment to this area.

In perspective ad hoc networking techniques will eventually dominate military networking, but we are at least a decade away from seeing significant penetration in defence markets. The big issue will be the performance of the US JTRS system in coming years as it is deployed to operational units. As ad hoc networks add considerable functional complexity, compared to all legacy networks, it will take a number of years before JTRS matures and achieves its full potential in robustness.

Further reading:
http://jtrs.army.mil/
http://www.csse.monash.edu.au/~carlo/adhoc.html
http://www.csse.monash.edu.au/research/san/