Multipath Routing Issues in Virtual Private Ad Hoc Networks

Peter Dedecker^{*†}, Jeroen Hoebeke^{*}, Ingrid Moerman^{*}, Joris Moreau^{*†}, Piet Demeester^{*}

^{*}Ghent University – IBBT – IBCN

Department of Information Technology (INTEC)

Gaston Crommenlaan 8 bus 201, 9050 Gent, Belgium

name.surname@intec.ugent.be

[†]University College Ghent

Department of Applied Engineering Sciences

Schoonmeersstraat 52, 9000 Gent, Belgium

name.surname@hogent.be

Abstract

In this paper we discuss the impact of proactive routing in comparison with a reactive routing approach in an environment where multiple communication technologies are used simultaneously. Therefore the Virtual Private Ad Hoc Networking platform will be introduced and used. A simulation environment is created and used to illustrate the outcomes with a simple scenario. Although many papers have been published comparing proactive and reactive routing protocols [1], this paper focusses on the impact on the VPAN platform with its two-level hierarchical architecture and nodes joining and leaving clusters.

1 Introduction

With growing broadband adoption rates and an unprecedented growth of wireless and mobile communications, including wireless LAN and cellular 3G networks, the Internet is evolving towards a large "network of networks", integrating a large number of different networking technologies and offering users broadband connectivity at any time and any place. Next to this, a convergence between traditional network devices and consumer electronics is taking place where devices that appear on the market are equipped with one or multiple network interfaces.

On the other hand, these large-scale communication networks overwhelm the user with available information, applications and services, increasing management and configuration complexity dramatically and introducing potential security risks. Moreover, a lot of applications require secure communication to take place only between a dynamic subset of distributed devices sharing a common context (i.e. communication between your personal devices, communication with colleagues or friends...), a characteristic not reflected by our current and future communication networks.

Therefore, it is expected that an evolution towards network virtualization will take place, imposing a logical struc-

ture onto this base network [3, 2]. These virtual networks will form a shielded and trusted environment for their participants, thereby logically structuring the network into small secure communities according to the needs of the end users, making the underlying base network invisible. Ad hoc protocols and techniques can be used to enable self-creating, self-organising and self-administering capabilities in these virtual networks, on top of existing network infrastructures, while service-discovery protocols enhance the use of upcoming service-oriented architectures and business models.

Precisely these trends led to the development of the Virtual Private Ad Hoc Network (VPAN) concept as described in [6] and [5]. In this VPAN concept, local and distributed nodes organise themselves in logical virtual networks, providing a secure and transparent overlay on available networks. Applications and services can be given access rights to such an overlay, thereby operating within this secure and confined environment. Security, network and management details are handled by the overlay in a self-organising way.

Figure 1. The VPAN concept.

All nodes that are allowed to join the overlay share a common cryptographic trust relationship. Locally, devices

that have link connectivity will automatically form clusters, i.e. groups of devices that establish secure and authenticated links and that do not rely on not-trusted nodes for their communication. The interfaces used therefore are called PAN-interfaces (Personal Area Networking) in this paper. Clusters on different geographical locations will be interconnected by establishing tunnels over an underlying IP network that acts as a carrier of the data, forming the complete VPAN. A graphical illustration is given in figure 1. For the communication within the VPAN, automatic private addressing and ad hoc routing techniques (intra-cluster and inter-cluster routing) are used. All communication is fully secured: only members of the VPAN can join the VPAN, all control information being exchanged is encrypted and completely shielded from the outside world or any other VPAN's and VPAN members can bind specific applications to their private VPAN IP address (application selectivity). Also, any communication medium (Bluetooth, WiFi, WiMax, UMTS, Ethernet,...) can be used to realize this intra and inter-cluster connectivity and the VPAN connectivity is maintained upon mobility.

In section 2 of this paper, some multipath issues that can arise in these kinds of networks are brought up. The current routing mechanisms with their associated shortcomings are discussed in section 3. With the simulation environment introduced in section 4, these shortcomings as well as the proposed improvements for low-cost routing will be evaluated in section 5, after which a conclusion is formulated in section 6.

2 Multipath issues in a VPAN topology

Figure 2. VPAN with multipath issues.

Figure 2 gives us a simple example of multipath issues. We observe two clusters, where the left cluster contains two gateway nodes resulting in two different tunnels to the right cluster. However, these two tunnels can have different characteristics, for example in terms of bandwidth, delay, jitter, packet loss, financial cost or reliability. If traffic is being sent from node 3 to node 2, the direct tunnel can be used. Another approach is to send traffic to node 1 which forwards it to node 2. This can be useful if the tunnel between

nodes 3 and 2 is set up over an (expensive and slow) UMTS interface while the tunnel between node 1 and 2 can be a cheap and fast xDSL connection which compensates the extra hop. Such path decisions can be made based on available network and traffic qualities, depending on user preferences or business strategies. In this paper we analyse the current routing approach of the VPAN platform and propose some handling methods to reduce financial cost. This is part of a bigger scope where the aim is to provide end-to-end Quality of Service (QoS) where possible and provide an indication for a Type of Service (ToS) in other circumstances.

The typical topology of the VPAN-based network, with distributed clusters interconnected by gateways, introduces multipath issues on three different levels. First of all, we have the intra-cluster level where source and destination are nodes within the same cluster. Normally, within a cluster, we don't have any financial cost associated with link usage.

If source and destination nodes are located in different clusters the problem gets more complex. To get packets to the other cluster, they must be forwarded over a secure tunnel by a gateway node. Multiple gateways in a cluster as well as multiple WAN-interfaces (Wide Area Network interfaces, providing internet connectivity eg xDSL, WiFi, UMTS,...) with different capabilities on a single gateway node increase the choice and complexity of selecting the best route. When a gateway is chosen, we fall back to the intra cluster routing problem to get the packets to this selected gateway. As these packets reach their gateway node, a tunnel to a gateway in the remote cluster must be chosen where multiple gateways in the remote cluster also increase choice and complexity. Routing between the two gateways is done by the underlying network in a non-transparent way: the VPAN platform has no influence on these decisions. When arrived at the remote cluster, the intra cluster issue arises again to deliver the packet to the right node.

3 Current routing in the VPAN platform

These three levels are directly visible in the VPAN routing software: intra-cluster routing, gateway selection and inter-cluster routing are separated as shown in figure 3.

For intra-cluster routing, proactive as well as reactive protocols are available with performances depending on different parameters. The proactive intra-cluster routing protocol is an implementation of WRP (Wireless Routing Protocol [10]) adapted to the VPAN environment: interaction with the cluster formation process and the propagation of gateway information has been taken into account. Using this proactive intra-cluster routing protocol, an arriving packet can be classified immediately as intra-cluster or inter-cluster. Using the reactive intra-cluster routing protocol, which is an AODV-based [12] protocol only implementing the core functionalities to establish a route, a

Figure 3. Current routing decisions in the VPAN platform

route request message needs to be broadcasted in the cluster. Packets with a destination address corresponding with another node in the same cluster are sent to the next hop as determined by the selected intra cluster routing protocol, while packets with destination addresses corresponding with the IP of the current node are delivered to the kernel. For destination addresses of nodes in other clusters and using the proactive intra-cluster protocol, the packet is forwarded to the gateway selection element which forwards it to the inter-cluster protocol if the node has access to the internet. Both proactive as well as reactive inter-cluster protocols, providing always-on (proactive) as well as on-request (reactive) tunnels to other clusters and there corresponding gateways, are available.

If the node isn't a gateway itself, another node in the cluster providing gateway services (which are broadcasted in the cluster in the proactive approach) is selected. The destination annotation of the packet is changed to this chosen node and the packet is handed over again to the intra-cluster forwarding element.

All nodes perform this tasks independently. The gateway selection element can encapsulate the packet in a brand new IP-packet with destination IP set to the chosen gateway preventing packets from being forwarded in a loop. Currently, gateway selection is done purely on a hopcount base.

When intra-cluster as well as inter-cluster protocol are set to reactive, the previous decisions come together. Then the source node sends a route request (RREQ) which is broadcasted in the cluster. When no route reply (RREP) arrives, a new RREQ with inter-cluster scope is sent. These requests are forwarded by all nodes to their neighbours and by all gateways on all there tunnels. After reception by the destination, a (unicast) RREP is sent back to the source node. Upon reception of this RREP, a complete path is set up at once and can be used immediately.

The used protocols are defined in a VPAN profile that is installed on all nodes in advance.

4 Simulation environment

To analyse these current routing approaches and suggest some improvements, a simulation platform has been developed.

The current VPAN platform is a result of our work on Personal Networking (PN) and Personal Network-Federations (PN-F) technology in the IST MAGNET and IST MAGNET Beyond projects [7, 4] in the European Commission's 6th Framework Program (FP6). This led to a proof of concept implementation used for real-life performance measurements on a testbed as discussed in [5]. The simulation platform however enables us to test scalability with a large number of nodes and real mobility, and leads to reproducible test results. Therefore we have chosen the NS-Click [11] platform, built upon the NS-2 Network Simulator [9] and the Click Modular Router [8], and made the VPAN proof-of-concept software compatible with it. Now an unlimited number of nodes and a VPAN Agent can be put in a simulation environment together for analysis in chosen circumstances.

As we want to analyse and improve gateway selection algorithms with financial cost optimisations, a UMTS-emulation is necessary and obtained by incorporating UMTS characteristics (bandwidth, delay) in the Click nodes.

5 Low-cost routing and gateway selection: implementation and analysis

Gateway selection decisions can be based on different metrics. Currently, only hopcount is available while new metrics (eg bandwidth, delay,...) will be implemented in the future. In this section, gateway selection that favours paths with low financial cost will be implemented and analysed versus the current available hopcount based gateway selection.

5.1 Scenario

To evaluate the suggested implementations, a scenario as seen in figure 4 will be used. This is a reduced topology of figure 2 where only the numbered nodes and the Agent remain. The scenario consists of two clusters with nodes 1 and 2 as fixed nodes while node 3 is a mobile node that travels between these two clusters. Direct (PAN) communication between node 3 and node 1 or 2 is done by the WiFi-interface, while node 3 is also equipped with an UMTS interface to allow communication when not in the environment of another WiFi node. 3s after starting the simulation, a client process on node 3 will start sending UDP-traffic to a server process on destination node 2. Two streams are sent:

Figure 4. Simulation scenario.

one for low-cost optimised routing and another one for hopcount based routing. This scenario with a travelling node leaving and joining clusters occurs frequently, for example with personal devices like the popular netbooks which use the WiFi network at home, a UMTS interface while travelling, and again a WiFi network at the office.

5.2 Proactive routing

In these simulations, the proactive routing protocol is used for intra-cluster as well as inter-cluster routing while tunnels are set up in a proactive way: they are always on and can be used immediately. Results of these simulations can be found in figure 5 and below. TUN-traffic is the application-level traffic (UDP) that is received by node 2. PAN-traffic is traffic on the PAN interface, corresponding here to the WiFi-interface of node 3. WAN-traffic is traffic sent over the WAN-interface, in this case the UMTS-interface of node 3.

Figure 5. Proactive routing; Node 3 moves from node 1 towards node 2.

In figure 5, TUN-traffic is twice as high as the WAN-

and PAN-traffic as two equal UDP-streams are sent. This is clear as the low-cost approach sends traffic to node 1 on the PAN-interface, which forwards it through the tunnel towards node 3 while simultaneously the hopcount based approach sends the second stream directly by the UMTS-interface (WAN) through the tunnel towards node2. On the graph, the WAN-flow seems slightly more than the PAN-flow, which can be explained as VPAN-traffic is encapsulated in OSI L2 packets for intra-cluster traffic and in IP packets for inter-cluster traffic through tunnels.

At the end of the first phase in figure 5, a falldown of TUN-traffic occurs. Node 3 can't reach node 1 anymore, but the linkbreak is not yet detected. Link break detection is done by the neighbourdiscovery module which sends beacons at a regular time interval. When X beacons (in our case: 2 beacons) are lost, the link is marked as broken and the routing table is adapted. Here, the beacon interval is set at 200 ms so it can take up to 600 ms until the linkbreak is detected. During this time, all packets sent on the PAN-interface are lost.

Notice the trade-off between packet loss and overhead due to beacon interval settings: the shorter the beacon interval, the faster link breaks (and new links) are detected and fewer packets get lost, but the more overhead can be observed. On pages 152-156 of [5] one can find a theoretical analysis of the impact of the beacon interval on overall performance. In our testcase, further research is necessary to optimise and dynamically adapt the beacon interval when travelling speed changes. Input parameters for a beacon interval adaption algorithm can be a GPS-signal or changes in measured link quality. But even without these optimisations no application level connections are lost as the VPAN uses private addressing and link breaks are detected before these connections time out.

In the second phase of figure 5 there is no direct connection anymore between node 3 and node 1 so all traffic is sent over the UMTS-interface. When node 3 enters the neighbourhood of node 2, both nodes detect each other depending on the previously mentioned beacon interval. When a link is set up, all traffic is handled by the intra-cluster routing. The small top in the TUN-traffic at the moment where sending of packets over the PAN-interface starts can be explained by packets still travelling through the slow UMTS-tunnel (delay = 200ms) and arriving after the first packet is received on the PAN-interface.

When node 3 starts in the neighbourhood of node 2 and moves towards node 1, we obtain results as shown in figure 6. Traffic is sent over the WiFi-interface first (intra-cluster). When the link break occurs, a falldown of TUN-traffic occurs. Upon link break detection, all traffic is routed through the WAN-interface over UMTS. When node 3 enters the neighbourhood of node 1, both nodes detect each other and set up a direct link and traffic with low-cost priorities is sent

Figure 6. Proactive routing; Node 3 moves from node 2 towards node 1.

over the WiFi-interface to node 1 which forwards it over the tunnel towards destination node 2. Traffic routed on a hopcount base keeps being sent over the WAN-interface at a higher financial cost.

5.3 Reactive routing

In these simulations, the reactive routing protocol is used for intra-cluster as well as inter-cluster routing while tunnels are set up on a reactive way: they are only set up when packets need to be sent. Node 3 still sends UDP-traffic to node 2. Results of these simulations can be found in figure 7 and below.

Figure 7. Reactive routing; Node 3 moves from node 1 towards node 2.

With reactive routing, the protocol doesn't have a overview of all nodes in the VPAN or the cluster. Routes are set up on request, as described in section 3. During path

setup time, all data packets are buffered so they can all be sent when a path comes available and no packets get lost. This can be seen in figure 7: the first traffic occurs at +/- 8 seconds from start, while in figure 5 the first packets are sent out at +/- 3 s, immediately after starting the client's traffic generation process. The difference in start time is due to path setup. After start, a peak can be observed after which traffic stabilises at a lower level. These are the buffered data packets sent out immediately at maximum rate once a path is set up.

When a link break occurs, a route error message (RERR) is being sent by the nodes that lost their connection and a new path needs to be set up. This explains the gap between seconds 24 and 28 in figure 7 corresponding to the time to set up this new path.

An established path is used until it breaks. So in figure 7 node 3 still uses its WAN interface instead of the direct link between sending node 3 and destination node 2. The route is not updated due to the reactive nature of the protocol. Also here, the low but longer peak is due to packets in the buffer which are sent after the new route setup, but here at a smaller rate than the first peak, as the UMTS interface is slower.

Unlike the proactive approach, we see (in the left of figure 7) node 3 using the path over node 1 to node 3 instead of the direct WAN-link between both. Where the proactive protocol uses the best available path (in terms of hopcount, or in the modified version in terms of cost), reactive routing uses the first path that can be established and keeps using this path, as long as it is available. As the UMTS-link is slow and all other links have almost no delay, the RREQ following the path from node 3 over node 1 to node 2 (see scenario topology in figure 4) arrives as first in node 2 which means the RREP is sent along this path which is also used then.

Here, both traffic streams are treated equally: the first available path is used. An optimisation could be made to start sending when a route is set up, but keep listening for possible better route replies and switch to these new routes when in favour. Multiple routes for multiple types of traffic should be available. Note however that with such an optimisation, node 3 will still keep using its WAN-interface for sending traffic to node 2 in this scenario, even when they reside in the same cluster and have a direct link.

When node 3 moves the opposite way, we come to similar observations, shown in figure 8. Here the start of the traffic flow occurs earlier, as both nodes are in the same cluster and we don't have to wait until the first RREQ with intra-cluster scope times out. After the link break, the new path is set up using the WAN-interface and buffered packets are sent out. When entering the cluster of node 1, this path is still being used however a more cost-effective solution is available.

Figure 8. Reactive routing; Node 3 moves from node 2 towards node 1.

6 Conclusion

In this paper multipath issues and their effect on a proactive and reactive routing protocol are evaluated using a developed simulation platform and a small scenario. Even with the limited scope of this small scenario, some general conclusions can be made.

First of all, in this kind of scenarios with mobile nodes with multiple interfaces and simultaneous used connections, reactive protocols are not suited to deal with possible better available paths due to the nature of this reactive approach. One should consider the trade off between overhead due to the proactive routing against the possible losses due to keep using less optimal paths.

As a second remark, with some slight optimisations, proactive protocols can be enhanced to deal with this kind of topology changes more effectively. A suggestion has been made to exploit the connection between neighbour and link break detection, and node mobility.

It is also possible to see the VPAN as one big ad hoc network in which an existing ad hoc (reactive, proactive, adaptive...) routing protocol can be used. This, however, has a number of drawbacks since the hierarchical structure of the VPAN network topology is not exploited: these protocols will have a significant higher overhead than protocols that explicitly take into account this hierarchical network topology. In addition, in order to be able to use a standard ad hoc routing protocol, some additional measures need to be taken such as the fact that tunnels should be seen by the protocol as links.

In the hierarchical VPAN topology and its routing protocols, a reactive inter-cluster protocol can be combined with a proactive intra-cluster protocol and reactive tunnel establishment. This avoids the delay when setting up a path to a node in a remote cluster: using a proactive intra-cluster

protocol, we don't have to wait for the time-out of the intra-cluster try as it occurs with the reactive protocol. Also the overhead of keeping up tunnels is eliminated. In scenarios with expensive links (e.g. UMTS-links), always a good path is used when available. This needs the gateway nodes to broadcast the cost of their tunnels (or more exactly the cost of using their WAN-interfaces as one node can have multiple WAN-interfaces) within the cluster. In this case, topology changes between the clusters are not taken into account for path optimisations, but one can expect these to be quite rare and have almost no impact.

7 Further work

First of all, the assumptions from this rather small experiment should be tested in bigger scenarios, with more nodes and a variable mobility, to evaluate the effectiveness of proactive intra-cluster routing combined with reactive inter-cluster routing in more scenarios.

Secondly, the proactive protocol used here is only effective when the sending node is the moving (gateway) node and a gateway in the home cluster must be chosen. When for example node 2 sends traffic to node 3, this will be done directly to the (expensive) WAN-interface of node 3. An optimisation in the inter cluster routing protocols where tunnel costs need to be incorporated, is necessary.

Finally: this paper is only the beginning of a whole story. Next to hopcount and monetary cost, more parameters can be taken into account (like delay, available bandwidth, reliability,...) and traffic can be routed along multiple paths according to its needs. Research to more optimal routing solutions according to traffic characteristics is necessary.

8 Acknowledgements

This research is supported in part by the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT) project Usenet and the Interdisciplinary Institute for Broadband Technology (IBBT) project SPAMM. Peter Dedecker is research assistant at University College Ghent and affiliated researcher at Ghent University.

References

- [1] M. Abolhasan, T. Wysocki, and E. Dutkiewicz. A review of routing protocols for mobile ad hoc networks. *Ad Hoc Networks*, 2(1):1–22, 2004.
- [2] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the internet impasse through virtualization. *Computer*, 38(4):34–41, April 2005.
- [3] K. P. Birman. The next-generation internet: Unsafe at any speed? *Computer*, 33(8):54–60, 2000.

- [4] D. Calin, A. McGee, U. Chandrashekhara, and R. Prasad. MAGNET: An approach for secure personal networking in beyond 3G wireless networks. *Bell Labs Technical Journal*, 11(1):79–98, SPR 2006.
- [5] J. Hoebeke. *Adaptive ad hoc routing and its application to virtual private ad hoc networks*. PhD thesis, Ghent University, 2007.
- [6] J. Hoebeke, G. Holderbeke, I. Moerman, B. Dhoedt, and P. Demeester. Virtual private ad hoc networking. *Wireless Personal Communications*, 38:125–141, 2006.
- [7] IST-MAGNET. My personal Adaptive Global NET. <http://www.ist-magnet.org>.
- [8] E. Kohler. The Click Modular Router Project. <http://www.read.cs.ucla.edu/click/>.
- [9] S. McCanne and S. Floyd. ns—Network Simulator. <http://nsnam.isi.edu/nsnam/>.
- [10] S. Murthy and J. J. Garcia-Luna-Aceves. An efficient routing protocol for wireless networks. *Mob. Netw. Appl.*, 1(2):183–197, 1996.
- [11] M. Neufeld, A. Jain, and D. Grunwald. Nsclick:: bridging network simulation and deployment. In *Proceedings of the 5th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems*, pages 74–81. ACM Press, 2002.
- [12] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc On-Demand Distance Vector (AODV) Routing. RFC 3561 (Experimental), July 2003.

10-13 January 2009

Las Vegas, Nevada, USA

Empowering the Connected Consumer

Getting Started

Welcome

Conference Information

Tracks

Authors

Search

IEEE COMMUNICATIONS SOCIETY

www.ieee-cnc.org

Celebrating 125 Years
of Engineering the Future
for all purposes or for creating new
ones that will be obtained from the IEEE.

©IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE.