Intelligent Autonomous Handover in iMANETs

Ben McCarthy
Computing Department
Lancaster University
Lancaster, LA1 4WA, UK
mccarthb
@comp.lancs.ac.uk

Panagiotis Georgopoulos
Computing Department
Lancaster University
Lancaster, LA1 4WA, UK
panos
@comp.lancs.ac.uk

Chris Edwards
Computing Department
Lancaster University
Lancaster, LA1 4WA, UK
ce
@comp.lancs.ac.uk

ABSTRACT

Autonomously making good network handover decisions is a complicated process that is fundamentally important in many complex mobile scenarios. In many mobile scenarios it is often infeasible to assume that an end user can be required to intervene and manually perform or verify a network handover decision. For this reason, utilities are required that can specifically manage the network connectivity of mobile nodes, monitoring their constantly changing state and the changing environment around them in order to ensure that the most appropriate connection is utilised at any given time. In this paper we present our Handover Manager that we have developed specifically for use in rescue system mobile networking solutions, such as mountain and coastal rescue, based on our experiences from real use case deployments. In particular we describe the way our solution autonomously manages connections to multiple heterogeneous access network technologies, we provide results from a testbed based analysis we performed and finally we draw upon our experiences to highlight important areas for future consideration that are applicable to our Handover Manager approach and to the wider MANET community in general.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*

General Terms

Design, Experimentation, Performance

Keywords

Network Mobility, Intelligent Handover, iMANET, Implementation

1. INTRODUCTION

Using an experimental implementation of the Unified MANEMO Architecture (UMA), researchers at the Computing Department of Lancaster University have successfully been able to produce working Internet-connected Mobile Ad hoc Network (iMANET) deployments. Research into MANET protocols has been ongoing for many years, however still no definitive solution exists that allows clusters of MANET nodes to freely roam across the Internet whilst maintaining topologically correct, globally reachable IPv6 addresses. UMA solves this problem and allows MANET nodes to connect to any Internet access network either directly (and subsequently proliferate their connection to other MANET nodes by acting as a Gateway) or indirectly via a node that is acting as a Gateway. These capabilities are hugely beneficial for improving the performance, the resilience and the overall functionality of our mobile networking solutions but without the correct management of connectivity at the underlying layers, all these benefits can be completely lost.

In our solutions we use Mobile Routers (MRs) that handle the mobility of entire networks of devices (typically Personal Area and Vehicle Area Networks). At any given time the MR must have a clear understanding of the status and quality of its current connection to the Internet (if one is in place) and all immediately available alternative connections that it may utilise should its current connection become no longer satisfactory. To achieve this we have developed a Handover Manager that constantly monitors all of the MR's connection possibilities (with probing performed at both the Physical Layer and the Network Layer) including both Direct (via an access network) and Indirect (via an iMANET Gateway) connections. In this paper we present the design and implementation of the Handover Manager and present results from its use in a specially designed testbed environment.

In Section 2 we begin by providing a brief overview of the Unified MANEMO Architecture in order to facilitate understanding of the network layer interactions that make the functionality of our Handover Manager possible. In Section 3 we introduce the Handover Manager, we highlight important aspects of its design and implementation and present a typical scenario it may be required to operate in. In Section 4 we detail the testbed we have developed for the purpose of testing, in a real mobile environment, the Handover Manager and its interaction with our mobile networking approach. Finally in Section 5 we conclude the paper by discussing the shortcomings highlighted by this work and the handover related challenges that will face Mobile Ad Hoc Networks in general as deployments are increasingly incorporated into the Internet.

MobiOpp '10, February 22-23, 2010, Pisa, Italy.
Copyright 2010 ACM 978-1-60558-925-1/10/02 ...$10.00.

2. UNIFIED MANEMO ARCHITECTURE

Primarily, UMA is a networking solution that is designed to support Internet-connected MANETs (iMANET). This means that communication between MANET nodes and communication with nodes in the Internet must be supported irrespective of the iMANETs changing network topology or the use of changing heterogeneous Internet access networks. In this section we will provide a brief overview of how UMA achieves this and highlight in particular the key features that facilitate the operation of the Handover Manager. For a more comprehensive description of UMA please refer to our paper presented at [1].

UMA supports iMANET mobility by incorporating an innovative Home Agent and bi-directional tunneling approach similar to the one utilised by the NEMO Basic Support protocol (NEMO BS) [2] with an implementation of the Optimized Link state Routing protocol (OLSR) [3]. Principally the key aim of this technique is to ensure that MANET nodes are able to maintain persistent global reachability in the Internet whenever any one node in a MANET has Internet connectivity, and maintain that reachability irrespective of any subsequent roaming across heterogeneous Internet access networks that takes place.

Our previous research carried out in the area of Network Mobility led us to identify that it would be possible to leverage the approach adopted by the NEMO BS protocol to allow entire MANETs of nodes and mobile networks to legitimately transmit packets into the Internet without any cooperation from the access network. Essentially, the bi-directional tunnel approach imposed by NEMO BS prevents packets transmitted by nodes in a mobile network from being Ingress Filtered [4] by the access network that a MR is connected to. For packets travelling in the opposite direction up to nodes connected to mobile networks, this tunneling approach also ensures these are delivered to the MRs current location in the Internet. If this functionality is therefore incorporated into the Gateway node of a MANET it can then feasibly perform these operations on behalf of all of the MANET nodes that it provides an Internet connection for. This is the basic principal that has driven the design of our Unified MANEMO Architecture (UMA) solution and is illustrated in Figure 1.

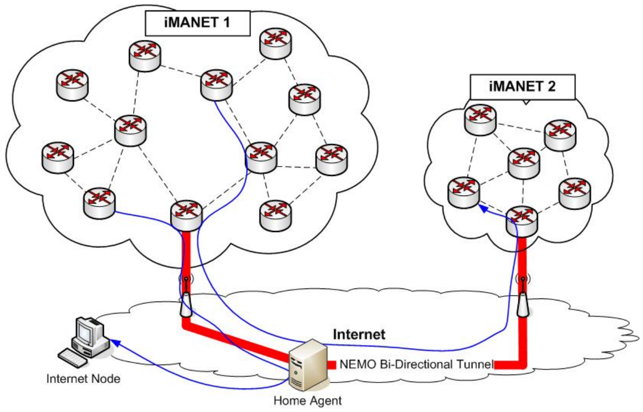

Figure 1: UMA Functionality Overview

UMA is able to form its bi-directional tunnel link with its Home Agent over any IPv6 Internet access network (and over most IPv4 Internet access networks with the use of an IPv4 to IPv6 transitioning technique such as 6-in-4 tunneling). This means that a UMA enabled MR can utilise numerous different heterogeneous access network connections, and simply designate which connection to use at any given time by establishing its bidirectional tunnel over that link. The complexity however is introduced when attempting to determine which connection is the best to utilise at any given time. UMA itself presents the MR with information related to the current characteristics observed in any iMANET it is attached to (these are discussed in more detail in the following section), but there is no centralised way to determine whether other connection options exist, compare the available options and then trigger handovers between networks, for this purpose we developed the Handover Manager.

3. HANDOVER MANAGER

In most scenarios it is highly beneficial for a MR to operate as an entirely autonomous entity, capable of operating continuously without any manual interaction from an end user. The primary role of a MR is to handle all aspects of mobility whilst presenting a standard IPv6 access network connection to any devices that connect to it. In this respect a MR is concealing the complexities of the underlying network in order to simplify the experience for the end user, and therefore constantly requesting user interaction to assist in handover decisions would effectively negate this functionality. To ensure that our MRs are capable of operating autonomously we developed an independent Handover Manager that interacts with the MR's network interface drivers and network layer mobility protocol to garner as much information about its surrounding connectivity options as possible and trigger handovers to the most appropriate network whenever necessary. Specifically, the task of our Handover Manager is particularly complex because it is required to simultaneously monitor multiple direct (Wifi, UMTS, etc) and multiple indirect (iMANET Gateways) network connectivity options and form handover decisions based on a wide range of varied criteria.

Figure 2 illustrates a typical scenario the Handover Manager may be presented with. In this diagram, the MR is connected to the Internet via the WiFi access point AP5. As an alternative, the MR could chose to establish its WiFi connection with another access point (AP1, AP2, etc) or it could choose to connect to the Internet via its UMTS connection with a cellular network, or via its ad hoc interface connection with MANET 1. During operation the MR will always transmit packets destined for nodes in MANET 1 directly into the MANET via its ad hoc interface, but for all other destinations the MR will route packets via its default route into the Internet. It is the management of this default route that is fundamentally the responsibility of the Handover Manager.

If the connection that the MR has in place with AP5 is satisfactory then it will continue to utilise it, however if the MR's connection with this access point becomes weak and its throughput drops below a satisfactory level then the Handover Manager must recognise this event and intervene by establishing a connection with a more suitable alternative. As the loss of an available connection can happen at any time, it is imperative that the Handover Manager has an up-to-date understanding of the available access networks available and their relative quality in comparison to one another.

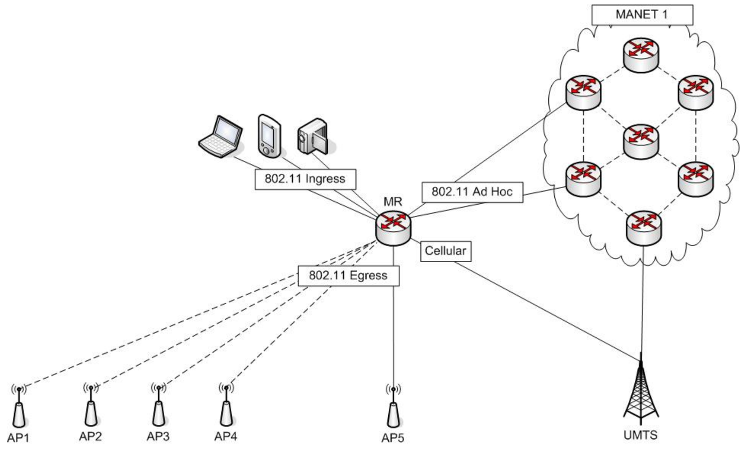

Figure 2: Example of Typical Handover Options Available During Operation

3.1 Design Considerations

In order to accurately gather information about the surrounding networking environment the Handover Manager simultaneously monitors a number of parameters, such as the signal strength of current WiFi connection and cellular connection, as well as the signal strength of all alternative WiFi APs. It also monitors the layer 3 connectivity of all available connections, as well as the network characteristics of any available iMANET connection.

The process of monitoring each of these parameters is handled by separate threads that each update a centrally stored database which holds relevant information about every available connection. Of the network parameters monitored, the signal strength of the MR's current WiFi connection and of its cellular network interface are the most straightforward pieces of information to ascertain as both values can be requested from the respective interface drivers. To ascertain the signal strength of all other available access points the MR periodically performs a scan of its surrounding environment and stores the results it gathers along with the previous n results in a sub table (where n is a value that is configurable at run time via the configuration file).

This allows the Handover Manager to derive the average signal strength and ignore potentially erroneous information caused by temporary obscurities in the radio environment.

In addition to monitoring the physical signal strength of the direct access network connections, the Handover Manager also monitors the network layer availability of its direct connections (if a connection is present) by periodically performing ICMP echo request/reply transactions with the access router of the access network.

Obviously the most important information this reveals is whether network layer connectivity is actually available over a given connection, since many situations can arise where an access network maybe physically visible at the radio layer but it is not possible to actually establish Internet access over the connection. Another important network characteristic this process exposes is the network latency times experienced. This is particularly useful in current deployments because as IPv6 adoption is not massively prevalent yet, IPv6 transitioning mechanisms which can add additional latency to the end-to-end path are routinely utilised to provide access over existing IPv4 networks. In addition, the latency experienced over a cellular data connection can vary considerably based on the type of technology supported (e.g. GPRS, EDGE, HSDPA, etc) and the available signal quality.

As well as monitoring the interfaces that the MR can establish its own direct connection with, the Handover Manager also monitors the network characteristics of the available paths to the Internet achievable via any iMANET that the MR is currently connected to. This comprises of two steps: 1) The Handover Manager takes into account the latency that each available iMANET Gateway advertises that it is currently experiencing. 2) The Handover Manager considers the link quality metrics associated with the multihop path between itself and any specific Gateway. In order to achieve the first step the OLSR component of UMA had to be augmented to accept the latency information about the Gateway's Internet access and then include it into one of the existing OLSR messages. The ability to carry information related to a Gateway's connection is not possible in OLSR so to introduce this capability we included it into the OLSR Host and Network Association (HNA) message format. HNA messages are used by Gateway's in OLSR to advertise their ability to reach the Internet and so by including the connection latency information of the Gateways within their actual Gateway advertisement message it means that this related information is collocated in one message. Upon receiving these messages OLSR then forwards their content (Gateway address and its current latency experienced) on to the Handover Manager for it to store and process. On the other hand, OLSR does inherently provide the MR with information related to the quality of the multihop path up to any particular Gateway, so this too is provided to the Handover Manager. In both of these cases the information flow is different for the Handover Manager than it is with the direct access network connections it forms. Rather than having to proactively probe or transmit traffic into a connection, in the indirect case of iMANET communication, information about the available network connections is sent to the Handover Manger.

All of the aforementioned information is gathered and periodically updated for the entire duration that the Handover

Manager is operational. After completing all of its monitoring processes once, the Handover Manger has effectively developed a network connectivity map of its local environment which it then strives to keep up-to-date as the nodes mobility causes network infrastructure to come into and out of range. This network map is then utilised by the Handover Manager's decision engine, which constantly cycles through all the available information in order to maintain an accurate and prioritised list of connection alternatives.

3.2 Handover Preference

As the Handover Manager's decision engine cycles through all of the collected data it adheres to a connection preference model in order to determine which connection to utilise at any given time. This preference model (illustrated in Figure 3 is based on the performance capabilities of the underlying networks, and therefore directly reachable WiFi access networks are given highest priority because throughput and latency times can be expected to be good over these networks. Data connections established directly via a cellular network are then prioritised next, largely because of the relative stability of the connection, afforded by the widespread coverage of these networks. At present our preference model gives the Internet access offered by an iMANET Gateway the lowest rank of the available connections. If neither of the direct connection options are available then the Handover Manager will begin to consider the best possible iMANET Gateway to establish its Internet connection via. At this point the Handover Manager compares the link quality metrics of the path between itself and the available Gateways and the Gateways' current latency times to determine the best available option. It is important to remember that this decision only affects the routing of packets destined for nodes in the Internet, all packets destined for nodes located in the iMANET that the MR maybe connected to will always be transmitted directly into the iMANET. This is because specific routes to the nodes in the iMANET will be known by the MR and therefore packets destined for those nodes won't be transmitted according to the default route. This ensures that a MR with a connection to an iMANET never wrongly transmits packets destined for another node in the same iMANET via the Internet first. Whilst at present we give iMANET Gateway connections the lowest ranking, we are continuing to analyse the real stability and throughput capabilities we can achieve from connections established indirectly via iMANET Gateways with the intention of determining whether there is a threshold where connections via iMANET Gateways can be considered to be more preferable than a direct connection to the Internet. Specifically, we are investigating whether iMANET Gateways with WiFi connections and good MANET link quality metrics (i.e. only 2 or 3 hops away) can be observed to be consistently better options than a direct cellular connection.

Once the Handover Manager has performed a network handover and the resulting connection is confirmed to be operational it records that the connection's status is operational. If the connection is a direct connection to a WiFi access network then the Handover Manager will not attempt to handover to any other network, unless the signal strength of the connection falls below a specific preconfigurable bound or network layer connectivity is lost. This measure is implemented in order to prevent the MR's network connection "flapping" back and forth between similarly high quality WiFi access network connections. However, if the Handover Manager established a connection to the Internet directly via a cellular network or indirectly via an iMANET Gateway then it will consider any new WiFi access network that becomes available for immediate handover. In this case, because the Station mode WiFi interface is not in use as the main connection to the Internet it will attempt to establish a connection with the Access Point it has identified and then perform a network layer check. If the connection is successful then the Handover Manager will immediately prefer the newly acquired WiFi network connection over the existing cellular or iMANET connection and therefore trigger UMA to carry out a new Binding Update over this connection.

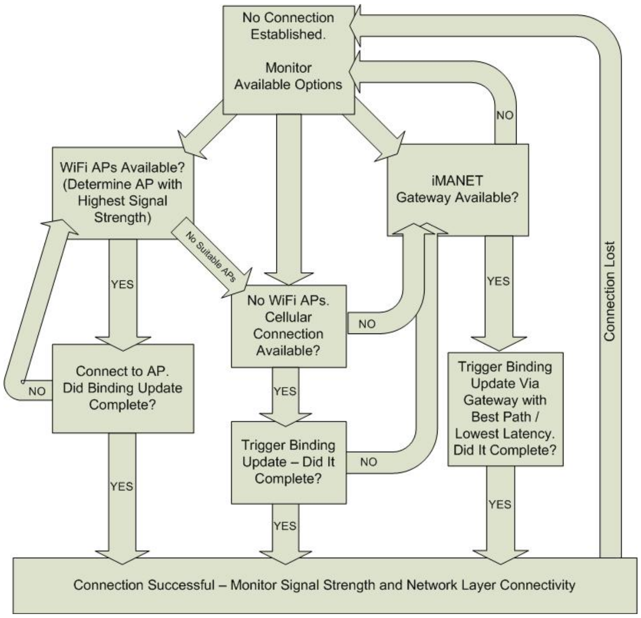
No Connection Established.
Monitor Available Options
NO
WiFi APs Available? (Determine AP with Highest Signal Strength)
iMANET Gateway Available?
No Suitable APs
YES
No WiFi APs. Cellular Connection Available?
NO
YES
Connect to AP. Did Binding Update Complete?
NO
Trigger Binding Update Via Gateway with Best Path / Lowest Latency. Did It Complete?
Connection Lost
YES
Trigger Binding Update – Did It Complete?
NO
YES
YES
YES
Connection Successful – Monitor Signal Strength and Network Layer Connectivity

Figure 3: Handover Manager connection Prefernece Model

In addition to performance and reliability considerations the Handover Manager can also act as a first line of security by preventing the MR from connecting to potentially malicious networks. The Handover Manager supports coarse or fine grained specification of policies related to which access points to connect to and their characteristics. This can range from allowing the MR to connect to any openly available access point, to only permitting it to connect to access points that match certain SSIDs or even specific individual access points identified by their MAC address. If an access point that matches the outlined specifications comes into range then the Handover Manager creates a record for it and stores information about it for the remaining time it is visible, otherwise it is ignored.

4. TESTING

For the purpose of evaluating our Handover Manager we developed a testbed that would allow us to analyse the accuracy of the network map that the manager developed. In particular we were most interested in determining whether the Handover Manager recognised changes in the underlying network topology of the available Internet access networks around it, and how long it took to determine that a change had occurred. In this section we detail the testbed setup we configured and the subsequent testing of the Handover Manager, which we performed in our labs using our custom built Mobile Router platform.

4.1 Mobile Router Hardware Platform

Physical deployments of our mobile networking solution are achieved using a custom built, battery powered mobile router that is designed to be carried in a backpack or mounted in a vehicle.

The mobile router hardware consists of an embedded system board that has a 680 Mhz CPU and 3 mini-pci slots which we installed three 802.11a/b/g (WiFi) cards into. One WiFi interface is configured in Station mode in order to establish connections with available Internet access networks. Another of the WiFi interfaces is configured in Ad hoc mode, this interface is used to establish Mobile Ad Hoc Networks with other nodes within range. The third WiFi interface in the Mobile Router is configured in Access Point mode to provide a wireless hotspot for end user devices to connect to and access the network. Finally, we also affixed a USB port to the board which we used to connect a 3G GSM modem. The board is powered by a standard COTS rechargeable Li-ion battery and it is all enclosed in a lightweight and waterproof plastic enclosure.

4.2 Test Design

To determine the capabilities of our approach we defined a number of criteria that would help us identify the overall performance and responsiveness of the Handover Manager. Those criteria were as follows:

- Length of time taken to establish full initial connectivity map.
- Length of time taken to recognise a change in the existing connectivity map.
  - Change to WiFi access network availability.
  - Change to Cellular Network availability.
  - Change to iMANET Gateway availability.
- Length of time taken to recognise a lost connection.
  - Loss of connection via WiFi access network.
  - Loss of connection via Cellular Network.
  - Loss of connection via iMANET Gateway.

To record the length of time taken for the Handover Manager to establish information about its surrounding network environment we began by producing a testbed which represented a typical networking landscape that a Mobile Router could be presented with. Illustrated in Figure 4, the setup comprised of 4 WiFi access points (Each connected to different IPv6 subnets of the Lancaster University network), 1 UMTS connection (provided by the UK cellular network operator O2) and a iMANET consisting of 3 UMA enabled MRs (2 of which were acting as Gateways). From the moment it was activated we then recorded the length of time the Handover Manager took to establish one complete connectivity map detailing all of the relevant information about each of the different possible access networks.

Once the Handover Manager had established its connectivity map and begun periodically updating its information, we analysed its responsiveness to alterations in the surrounding environment by then subsequently shutting down network infrastructure (and in the case of the cellular connection, reducing the 3G modem's signal quality). For this phase of testing the main connection to the Internet was established via AP4 which remained untouched and always available. This in turn meant that all changes to other networks were peripheral to the main Internet connection and would not therefore cause the MR to handover at anytime, but instead would only affect the potential ordering of networks in the Handover Manager's prioritised list of connection alternatives. To record the responsiveness to changes in the WiFi availability we shut down AP1 and AP2 simultaneously, this ultimately left only AP3 as a viable alternative and therefore we recorded the point from when the APs were shutdown to the point at which AP3 became the highest ranked (and only available WiFi connection alternative). To test the Handover Manager's responsiveness to change on the cellular interface we chose to manipulate the connection by reducing the modem's received signal strength. We did this by detaching the modem's dedicated external antenna and then placing the modem into a Faraday cage. Using this approach was not enough to remove all trace of signal quality but it was enough to dramatically reduce the reported signal strength and subsequently disrupted the network layer data connection. For the test of responsiveness to change in the iMANET network conditions we simply forced the iMANET node to stop advertising it was a Gateway in its HNA messages, as this is what would happen in a real scenario if an iMANET Gateway lost its own direct connection to the Internet.

After performing the tests to measure the responsiveness of the Handover Manager to changes in the surrounding network environment we also wanted to see how quickly it responded to losing its main point of connection to the Internet (and therefore subsequently utilising the best possible alternative connection available at that time). For these tests WiFi AP1, AP2 and AP3 remained permanently deactivated and AP4 was only activated for the first test where the WiFi connection was established and then lost. This was to prevent the Handover Manager from immediately preferring the other WiFi APs over the cellular or iMANET Gateway connection as it is designed to do. This meant that for each of the three connection types tested we first setup the main connection to the Internet via the appropriate network and then upon dropping the connection a subsequent handover was performed via the cellular network (expect when we tested the loss of the cellular network connection, in which case a handover was performed via the iMANET as no WiFi AP was available).

Finally, for each test performed we set the rate at which each monitoring process performed its sampling to one of four settings (0.5 seconds, 2 Seconds, 5 Seconds and 10 Seconds). All of the separate individual scanning and monitoring processes that the Handover Manager performs are carried out periodically, and therefore it is obvious that the frequency that these processes are carried out will have a significant effect on its responsiveness. Setting the interval between the time each process is run to be relatively high (e.g. 10 seconds) will result in a lot of dead time where the Handover Manager will be oblivious to potentially large scale changes to the surrounding network connectivity options. However, setting the sampling rates to be extremely low (e.g 0.01 seconds) will ultimately utilise more power and CPU resources.

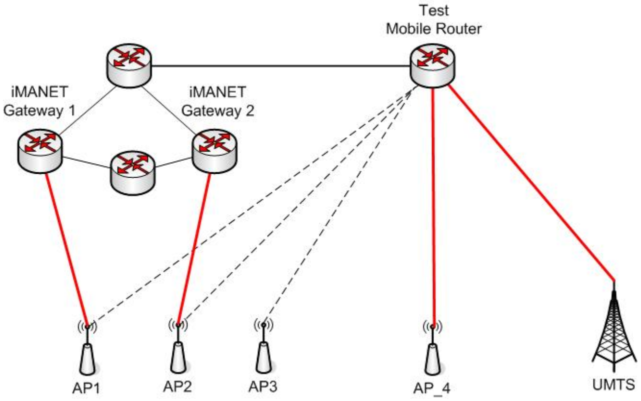

Figure 4: Handover Manager Testbed Configuration

4.3 Results

The results from our lab based testing are presented in Table 1. One of the most obvious trends we observed and that is evident in our results was the obvious significance of the periodic interval utilised by the monitoring processes. However the length of the interval period did not have any effect on the first set of results because the Handover Manager immediately starts its monitoring processes when it is activated in order to populate its network connectivity map as quickly as possible. Once completed, the interval period is then only adhered to as the monitoring processes begin to update the network map information with the changing status of access networks around them. One significant factor that we observed to have added to the time the Handover Manager took to complete its tasks or respond to a change was the network layer probing it performs and the overhead that it introduces. This was evident to us during testing because our output data would show that all of the signal strength scans (of both the directly connect WiFi AP and cellular connection and the peripheral WiFi APs) were completed relatively quickly (depending on the interval rate). However, the subsequent information gathered from determining the latency experienced at the network layer would often arrive significantly later. This is mainly because of the number of processes that must complete beforehand in order for the network layer information to be gathered. Specifically the Handover Manager must wait firstly for the connection with the appropriate network to be established, then once in place it must then wait for the IPv6 Neighbor Discovery process [5] to complete (namely a Router Advertisement / Neighbor Advertisement transaction must take place). This then allows the Handover Manager to ascertain the IPv6 Link Local address of the network access router and finally an ICMP echo request can be sent.

Once the Handover Manager has established its connections and built its initial network connectivity map these network layer delays are no longer a problem because the appropriate addresses have been established and recorded. In the tests measuring the Handover Manager's responsiveness to change in the surrounding network environment, the deactivation of two of the WiFi APs was quickly detected because the AP signal strength scanning process is very fast. With the iMANET connection, the loss of a Gateway causes UMA to proactively alert the Handover Manager, whilst this means that the issues imposed by the sampling interval do not affect this test, the responsiveness is still somewhat sluggish because OLSR must first timeout the Gateway to ensure it has definitely stopped transmitting HNA messages, this process in itself was observed to take over 3 seconds.

The Handover Manager's ability to detect lost connections is of particular importance since the sooner a completely unusable connection is detected the sooner a usable alternative can be supplied. With the Handover Manager alot of emphasis is placed on a dramatic loss in signal strength, or the network layer connectivity being reported as lost when detecting connection losses. In the first tests where the WiFi and cellular connection were lost the handover was triggered relatively quickly because the loss of the network was detected from the immediate loss of signal strength reported by the interface (although the reaction with the cellular link was a little more sluggish because our method of dropping the signal strength was not as definitive as deactivating the WiFi AP). For the iMANET Gateway, the bottleneck is again the time taken by OLSR to determine that the Gateway is no longer present. In this case the MR waits for a set period of time (i.e. this is another potentially configurable interval value) after it received its last advertisement from an iMANET Gateway before it decides that the Gateway is no longer present. This process could potentially be speeded up by forcing the Gateway to proactively advertise the fact that it is no longer a Gateway. This approach would make sense because Gateways can often still be connected to the same MANET but they lose their own direct connection to the Internet. In which case they must no longer advertise that they are a Gateway, but at the same situation should not necessarily be treated in the same that a MANET node disconnection is (as it is treated now).

One important advantage that the use of the Handover Manager provides is the ability to benefit from the use of "Make-before-break" handovers in Vertical Handover situations. This is because the Handover Manager is able to establish a connection simultaneously with a different heterogeneous access network to the one that is currently utilised as the main connection to the Internet, and then crucially the Handover Manager can step in at any moment and force UMA to perform a Binding Update over any interface it specifies. This ability means that the large overhead attributed to establishing connections at Layer 2 can be avoided in some circumstances. For instance if a MR has a connection to the Internet via its cellular interface and a WiFi network subsequently moves into range, the Handover Manager can establish a connection to the WiFi network and first check it for Network Layer connectivity before triggering UMA to switch over its connection to the Internet. By using this approach the Handover Manager ensures that the only delay in setting up a session is imposed by the time it takes to communicate with a MR's HA over the Internet. This means that in low latency networks the whole registration process can be achieved in under 0.5 seconds.

5. FUTURE CHALLENGES

In this paper we have detailed our design and implementation of an autonomous Handover Manager for use by mobile nodes that can utilise heterogeneous Internet access

Interval Period	Establish Full	Recognition of Change in Connectivity Map			Recognition of Change in Connection State		
(Seconds)	Connectivity Map	WiFi	Cellular	iMANET	WiFi	Cellular	iMANET
0.5	5.4 s	0.6 s	1.9 s	3.3 s	1.5 s	2.9 s	4.1 s
2.0	5.8 s	1.0 s	2.3 s	3.4 s	2.8 s	4.9 s	4.2 s
5.0	5.7 s	2.8 s	4.1 s	3.3 s	4.4 s	6.7 s	3.9 s
10.0	5.5 s	5.5 s	5.9 s	3.2 s	6.7 s	8.3 s	4.4 s

Table 1: Testing Results Summary

networks both directly and indirectly via iMANET Gateways. From the results gathered in our testing phase we have demonstrated that our approach is entirely feasible for use in real world deployments and it achieves its main purpose of entirely automating handovers in complex mobile networking scenarios. This initial work we have completed has highlighted interesting areas of consideration for future work specifically related to the improvement of our Handover Manager and its functionality. However, in addition it has also highlighted some broader ranging concerns related to the challenges of incorporating Mobile Ad Hoc Networks into the Internet and the shortcomings that need to be tackled to better support intelligent handover in this domain.

From the perspective of our Handover Manager, one of the key observations we made is the importance of the periodic sampling rate to each of the individual monitoring processes. These monitoring processes are integral to the functionality of the Handover Manager and thus it is extremely important to configure the rate at which they sample the network environment correctly. Manipulating these individual values during an initial configuration phase to adapt the performance of the Handover Manager for different situations and scenarios could be a difficult and labourious task for an end user. Instead we intend to research the possibility of using a collection of generic profiles that cover the requirements of a diverse range of scenarios, allowing the end user to simply select one of the profiles to be adopted and thus simplifying the process of configuring the Handover Manager greatly. However, utilising static profiles in this manner still implies that a predefined sampling interval is appropriate for all scenarios, but as mobile networking scenarios are known to be inherently transient this is unlikely to always to be ideal. Instead we intend to also research the potential benefits of implementing an adaptive periodic scanning and monitoring scheme that allows the Handover Manager to alter the rate at which it samples certain network characteristics based on the current network environment. For this purpose we intend to carry out detailed further analysis of all the post operational data we attain from the use of our Mobile Routers. This will hopefully help us to determine whether obvious situations and scenarios arise where sampling rates could be greatly reduced without detriment to the overall functionality of the Handover Manager (and thus preserve battery life) or greatly increased at other times to bring about marked improvements in performance at key moments.

Supporting Internet access for MANET nodes and allowing MANET deployments to proliferate the connectivity of multiple heterogeneous access networks to all of their constituent nodes can potentially provide highly robust networking solutions for many important scenarios. In most MANET scenarios, incorporating the ability to communicate unrestrictively via the Internet can bring added advantages and therefore it can be seen as a key goal for the future. Our UMA networking approach demonstrates one comprehensive technique for realising this type of network, and research work into other possible techniques continues to take place. However, what is made apparent by the handover research in this paper is the need to incorporate consideration for providing the information to carry out intelligent handovers, at as early a stage as possible in the protocol design efforts that are taking place. From our experience we recommend that consideration is taken to ensure that any protocol solution is designed to facilitate the dissemination of an iMANET Gateway's access network connection capabilities. Disseminating accurate information about the capabilities of each of the Gateway's different potential links creates a number of issues. Firstly, it is important to determine what type of network capability information the Gateway should record (latency, packet loss, jitter, etc) for different scenarios different information may be important. If this is the case and the data collected can differ significantly, then the approach to disseminating this information must be adaptable enough to deal with these changes also. In addition, determining the best way to then express that information to other nodes throughout an iMANET is also a non-trivial task (should a fine or course grained approach be used, can generic profiles be ascribed to different collections of characteristics?) Finally it is then also important to consider the best way to actually disseminate this information, to ensure interoperability and standardised use across future deployments. For this reason, part of our future work effort will be devoted to developing a protocol independent approach for iMANET Gateway's to disseminate adaptable information about their access network connections in order to improve the ability for other nodes to make intelligent handover decisions. Hopefully, by making a comprehensive solution to iMANET Gateway information dissemination available to the MANET community in general, we will ensure that as more MANET deployments are incorporated into the Internet, many of their requirements needed to support intelligent handovers will already have been fulfilled.

6. REFERENCES

[1] B. McCarthy and C. Edwards, and M. Dunmore. "Using NEMO to Support the Global Reachability of MANET Node". In Proceedings of Twenty Eighth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2009), 19-24 April 2009, Rio de Janeiro, Brazil.

[2] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. "NEMO Basic Support Protocol". IETF RFC 3963, January 2005.

[3] T. Clausen and P. Jacquet. "Optimized Link State Routing Protocol (OLSR)". IETF RFC 3626, October 2003.

[4] P. Ferguson and D. Senie. "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing" IETF RFC 2827, May 2000.

[5] T. Narten, E. Nordmark, W. Simpson, and H. Soliman. "Neighbor Discovery for IP version 6 (IPv6)". IETF RFC 4861, September 2007.