Future Network & MobileSummit 2010 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2010
ISBN: 978-1-905824-18-2

Poster Paper

Boosting mobility performance with Multi-Path TCP

Marcelo BAGNULO[1], Phil EARDLEY[2], Alan FORD[3], Alberto GARCIA-MARTINEZ[1], Alexandros KOSTOPOULOS[4], Costin RAICIU[5], Francisco VALERA[1]

[1]*Universidad Carlos III de Madrid, Av. Uniersidad 30, Leganes, 28911, SPAIN*
[2]*BT Innovate, Adastral Park, Martlesham Heath, Ipswich, IP5 3RE, UK,*
[3]*Roke Manor Research*, Old Salisbury Lane, Romsey, Hampshire SO51 0ZN, UK,
[4]*Athens University of Economics and Business, 47A Evelpidon, Athens 11362, Greece*
[5]*University College London, Gower Street, London WC1E 6BT, UK*

Abstract: Fourth Generation mobile devices incorporate multiple interfaces with diverse access technologies. The current Mobile IP protocol fails to support the enhanced fault tolerance capabilities that are enabled by the availability of multiple interfaces. In particular, established Mobile IP communications cannot be preserved through outages affecting the Home Address. In this paper we describe an architecture for mobile host multihoming that enables transport layer survivability through multiple failure modes. The proposed approach relies on the cooperation between Mobile IP and Multi-Path TCP and aims to fully support multihoming and extend roaming capabilities of mobile devices.

Keywords: 4G, Mobile IP, Multi-Path TCP, Multihoming, Network Architecture.

1. Introduction

Mobile devices are emerging with multiple physical interfaces with different technologies in a single device, which greatly extends their roaming capabilities. Multihoming should bring several benefits. It enables a mobile node to preserve established communications as it moves through areas served by dissimilar access networks. It increases fault tolerance, including the preservation of established communications through different types of outages. It allows simultaneous use of multiple available paths towards the correspondent node, when multiple access technologies are available. However, currently available mobility protocols fail to support the aforementioned features, and new mechanisms are needed to fully support multihoming and so enable its benefits to be realised. Thus, it is essential to deal with these challenges of building the Future Internet, based on present mobile and wireless communications infrastructures.

In this paper, we present a mobile host multihoming solution. The Multi-Path TCP (MPTCP) protocol [1], which is currently under development at the Internet Engineering Task Force (IETF), splits a TCP connection over multiple paths. In this paper, we propose how to use a combination of MPTCP and Mobile IP [2, 3] in order to achieve better mobility support for multihomed nodes than MIP alone. Moreover, transparent support to existing applications using TCP is guaranteed, since MPTCP is presented as regular TCP to current applications.

In particular, the contributions of this paper are:

- To identify possible mobile host multihoming configurations and the limitations of

the Mobile IP protocol to support them. (Section 4)

- To propose and discuss a novel solution that integrates both MIP and MPTCP. (Section 5)

We also provide some essential background about the Mobile IP protocol (Section 2) and the MPTCP protocol (Section 3).

2. Mobile IP

Mobile IP (MIP) provides a scalable mechanism for mobility support in the Internet. By using MIP, a mobile node can change its attachment point to the Internet while preserving established communications. It has been defined both for IPv6 [2] and IPv4 [3]. While there are some differences between the two flavours of the protocol, they have a significant amount of commonalities. In this section we will provide a high level description of the Mobile IP approach and we will also note the key differences between MIPv4 and MIPv6.

The main components involved in MIP operation are: the *Mobile Node* (MN), originally located in the Home Network, that roams through different Visited Networks; the *Home Agent* (HA) located in the Home Network; and the *Correspondent Node* (CN). The MN has at least one stable address, called the Home Address (HoA), which is topologically meaningful as long as the MN is located at the Home Network. When the MN moves away to a Visited Network, it acquires at least one topologically meaningful address at its new location, the *Care-of Address* (CoA). However, independently of the MN location, packets addressed to the HoA are routed to the Home Network. As soon as the MN has left the Home Network, the MN uses a signalling message to inform the HA about its current location, i.e. its current CoA. This message is called *Binding Update* (BU) in MIPv6, and *Registration Request* in MIPv4. When the HA is aware of the MN location, it tunnels the packets addressed to the HoA to the MN at its present location i.e. CoA, preserving the communication.

While the MIPv4 protocol supports only the communication through the HA, the MIPv6 protocol has two operation modes: the *Bidirectional Tunnel* (BT) mode and the *Route Optimization* (RO) mode, as depicted in figure 1.

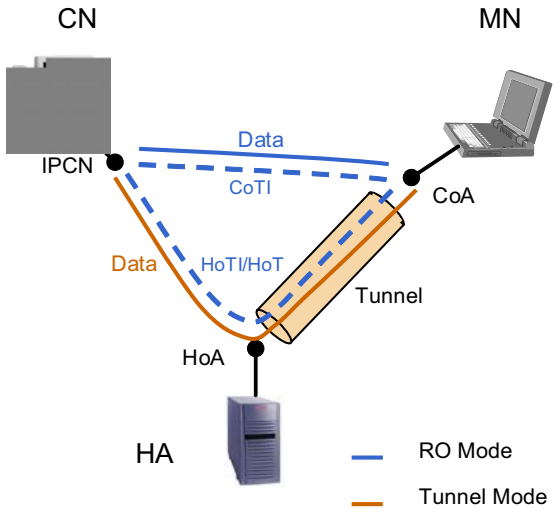

Figure 1. MIPv6 operation

In the BT mode, packets are routed through the HA as long as the MN is away from

home, as described above.

In the RO mode, the MN also informs the CN about its current location, sending it a BU message containing its current CoA. The result is that packets are exchanged directly between the MN and the CN without HA intervention. In order to protect these BU messages, a security mechanism called *Return Routability* (RR) check is used. The RR procedure consists on the CN exchanging with the MN two different nonces, one through the HoA (using a message exchange called HoTI/HoT) and another one through the CoA (using a message exchange called CoTI/CoT). If the MN can show that it has received both nonces, it can prove that the claimed HoA is co-located with the claimed CoA. In order to limit the scope on time of man-in-the-middle attacks, the bindings between a HoA and a CoA that are validated through the RR procedure have a maximum lifetime of 7 minutes. After such period, the RR procedure needs to be executed again to extend the lifetime of the binding.

3. Multi-Path TCP

The resource pooling principle [4] sets out a goal of making the network's resources behave like a single pooled resource (*i.e. separate resources appear to act as one large resource*). So, it implies that the load can be spread over several resources or relocated to a different resource. In the case of MPTCP, this load is data and the resource is links. The main objectives of resource pooling are to increase robustness against failures, provide better ability to handle localised surges, and increase resource utilisation. Combined, all these factors enhance user's experience and efficiency of resource usage.

MPTCP supports the use of multiple paths between source and destination, thus bringing the main benefits of reliability, flexibility and efficiency. When multiple paths are used, MPTCP will react to failures by diverting the traffic through paths that are still working and have available capacity. Old and new paths can be used simultaneously. Since MPTCP is aware of the congestion in each path, the traffic distribution can adapt to the available rate for each path. This provides resilience, higher throughput and handles more efficiently sudden increases in demand for bandwidth.

Although the current Internet's routing system only exposes a single path between a source-address pair, more and more hosts (especially mobile hosts) have multiple addresses that could be used to generate multiple paths to other hosts on the Internet. Thus, multipath transport could be deployed in the current Internet without requiring an upgrade to the routing system.

The MPTCP design [1] provides multipath TCP capability when both endpoints understand the necessary extensions to support MPTCP. This allows endpoints to negotiate additional features between themselves, and initiate new connections between pairs of addresses on multihomed endpoints.

3.1 – Basic operation

The goal of a multipath TCP flow is to take a byte-oriented stream of data from the application and splits it across multiple "subflows", i.e. one multipath TCP connection that consists of different TCP (sub)connections. To an application, an MPTCP session looks no different to a standard TCP connection – the addition and combination of new paths is handled at the transport layer.

A particularly notable issue is that regarding sequence numbers: TCP has a single sequence number for the data, and assigns some parts of it to the signalling mechanisms to give them reliability (SYN and FIN). MPTCP proposes that each subflow has its own sequence number space, and there is an additional data-level sequence number, presented in

a TCP option. The use of TCP options seems the most appropriate way of including metadata about packets without interfering with the payloads although it does carry a penalty in terms of increased packet overhead. Each TCP segment now needs to identify both the subflow and the data (flow) sequence number. The receiver uses the data sequence number to reorder the application data. The receiver only acknowledges subflow sequence numbers. The sender can infer from this exactly which data segments have arrived, and on which path, and can identify which data needs to be retransmitted upon packet loss. The inherent reliability of TCP exists at the subflow level.

For any given subflow sequence number, a single data sequence number can be mapped on it (the mapping is an injective function). This is to ensure compatibility with traffic normalisers, which may replay cached data if they see the same (subflow) sequence number, whilst ensuring the sender knows which data has been received.

This also means that, if a data segment is lost on subflow A and retransmitted on subflow B, subflow A must keep retransmitting the segment until it is acknowledged, in the normal way, even though the data segment may have arrived already at the destination (as it cannot map new data on the same sequence number).

3.2 – Connection management and path discovery

A multipath connection starts as a normal TCP connection with an additional option that indicates that the sending host is Multi-path TCP capable and it includes a token that the sender uses to locally uniquely identify this multipath connection (which can be seen as being analogous with a port number). If one of the endpoints does not support multipath, or if middleboxes drop new options, the connection falls back to a standard TCP connection.

Once the connection has been opened, and thus identifying tokens for each endpoint are exchanged, either endpoint can create additional subflows by initiating the three-way handshake on a different path. These SYNs will contain the token the endpoint received upon connection initiation, which is included to allow correct association of the subflow to the overall Multipath TCP connection. The source and destination ports of these subflows no longer have the same meaning as for regular single path TCP, as the token is also used for demultiplexing (it is likely that they will use the same destination port as the initial connection to ease NAT and firewall traversal, and to assist with traffic engineering).

Subflows can be added or removed during the lifetime of the connection, perhaps to accommodate new available paths. Removal is normally done using the FIN/FIN ACK exchange. However, in order to ensure graceful connection shutdown even when some paths have failed, we need an additional DATA FIN which acts as a connection level FIN, indicating that the sender has no more data to send.

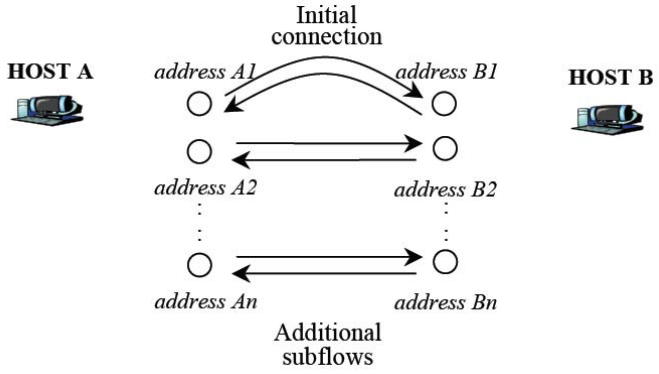

Figure 2. An MPTCP connection establishment

Two-ended multipath TCP discovers paths through the presence of multiple routable IP addresses on a host (figure 2). A host is aware of its IP addresses, and can attempt to initiate new connections to the other endpoint using these new addresses. An endpoint can also inform its peer of alternative addresses it is reachable on, in the event that it cannot itself initiate connections due to the peer being behind a NAT or firewall.

4. Limitations of Multihoming Support in Mobile IP

In this section we first present different multihoming configurations and then we identify the limitations of the MIP protocol for multihoming support.

We consider that a MN is multihomed when it has more than one HoA and/or more than one CoA. We can identify the following scenarios where a MN is multihomed:

- A 4G MN that has multiple physical interfaces, presumably with different access technologies. In this case, the MN may have multiple HoAs, since each interface may have a different Home Network; and it may also have different CoAs, since each physical interface may be located at a different Visited Network.
- A MN with a multihomed Home Network. In this case, the Home Network is connected to multiple ISPs, each of which delegates a prefix, resulting in multiple HoAs, one per prefix.
- A MN that is roaming in a multihomed Visited Network. As in the previous case, when a Visited Network is multihomed, multiple prefixes are available. Thus, a MN visiting the multihomed network has the possibility of configuring multiple CoAs.

A feature that is common to all the identified multihoming configurations is the availability of multiple paths between the MN and the CN. The existence of multiple paths enables both extended fault tolerance (since in the case of a failure the communication can be preserved by using an alternative path) and enhanced performance (since the MN can use multiple paths simultaneously in order to exchange packets). However, as we describe next, current MIP protocol fails to provide full support for these features.

In terms of fault tolerance capabilities, the fate of ongoing communications in current MIP specifications is tied to the availability of the path between the MN and the CN through the HA. In the case of MIPv4, this is obvious since there is no Route Optimization mode that would allow direct communication between the MN and the CN without passing through the HA. In the case of MIPv6, even though there is a Route Optimization mode, because of security reasons, failures along the path though the HA affect ongoing communications even though alternative working paths are available, as it occurs in figure 3c.

In the case of a multihomed MIP (either MIPv4 or MIPv6) node with multiple HoAs and multiple CoAs that is communicating in BT mode (figure 3a), it is trivial to see that a failure affecting the reachability to the HoA would break the established communication. This is true even in the case that other reachable HoAs are available, because MIP does not provide support for changing the HoA used for an established communication. In summary, in BT mode, a failure in the path between the CN and the MN through the HA affects any communication established using the corresponding HoA, even if there are other working HoAs available.

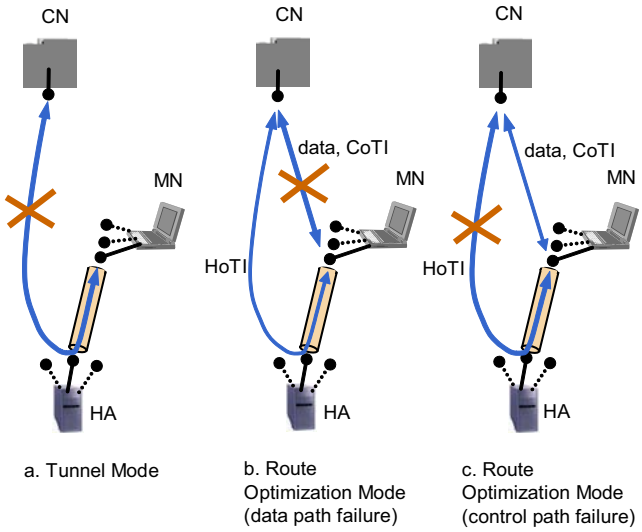

Figure 3. Failure scenarios

In MIPv6, in the case of a MN with multiple HoAs and multiple CoAs that is communicating with a CN in RO mode (figure 3b) it is also clear that a path failure between the CN and the CoA used for the communication would affect the ongoing communication. The MIPv6 protocol could provide means to survive this outage if it could detect it and try to use an alternative CoA, or fall back to the path through the HoA. Since MIPv6 does not have any mechanism to detect this type of outage, the communication will be interrupted in this case. In addition, the established communication is not only vulnerable to outages in the path used to exchange data packets, but it is also vulnerable to path failures between the CN and the MN through the HA (figure 3c). This is so because the path through the HA is used to periodically exchange HoT/HoTI packets. In case that an outage affects this path, the HoT/HoTI packet exchange would be interrupted. The result is that the binding between the HoA and the CoA in the CN will expire [2], and the communication will fall back to the path through the HA, which is not working. So, in RO mode, ongoing communications are not only vulnerable to path failures between the CN and the currently used CoA, but they are also vulnerable to outages in the path between the CN and the MN through the HA.

In terms of performance, current MIP protocols support the so-called flow-binding feature [5][1] that allows the MN to bind different flows to different CoAs. While the definition of a flow is very general, and the same flow can potentially be bound to multiple CoAs, it is not possible to perform a distribution of the traffic that maximizes the usage of the different paths, since MIP does not have the knowledge of the available capacity in each path. The result is that load distribution in MIP is done blindly with respect to the available capacity in each path, resulting in suboptimal performance.

In the next section we will propose an architecture that relies in MPTCP to overcome the aforementioned limitations.

[1] The flow binding capability has been defined for MIPv4 and it is currently under definition for MIPv4

5. A Mobile Host-Multihoming Architecture

5.1 – Description of the proposed Architecture

In this section we describe how to integrate the MIP and MPTCP protocols, without requiring any modification to either of their protocol messages, in order to provide full multihoming capabilities to mobile devices.

The MN and the CN include in their stacks both a MPTCP and a MIP module. The MIP mechanisms are placed underneath the MPTCP layer. The proposed approach separates established session mobility from initial contact when the mobile is roaming. MPTCP is used for the former, whilst MIP is used for mostly the latter. MIP is also used to ensure reachability in the case simultaneous movement.

We consider a simple scenario, in order to illustrate the key points. The connection is initiated between the MN's HoA and the address of the (static) CN (IPCN), with Mobile IP playing its standard role when the CN initiates the connection and the MN is away from home.

If the MN is already away from home, or as soon as it moves away, the MIP layer of the MN creates a binding between HoA and one of the CoAs available at the visited network. The MIP layer also notifies the MPTCP layer of the host's stack so that, as soon as communication is established, MPTCP can add the CoAs as an additional addresses for the MPTCP connection, and then distribute the traffic using the coupled congestion control proposed for MPTCP [6]. This fact would result in distribution of the load, based on the available capacity for each path/address-pair. Note that the application protocols located on top of the MPTCP layer use the address pair that was used to initiate the communication.

In fact, since the MN has multiple interfaces, it may have multiple CoAs and even multiple HoAs. These can all be added to the MPTCP connection and traffic distributed across them appropriately. Also, when the MN moves, MPTCP uses the new CoA and MIP (as normal) updates the HoA-CoA binding.

Also, MPTCP could close the subflow associated to the HoA once a CoA is available, so that packets are not exchanged through the path through the HA, but are sent directly. This is likely to be a sensible optimisation.

The resulting state is the following:

- The address pair used to initiate the connection are HoA_i and IPCN
- The additional addresses that MPTCP has for each endpoint are:
  - IPCN for the CN
  - HoA_i, CoA_1,…, CoA_n for the MN
- The MIP layer has a binding for HoA_i to CoA_p. It may also have other bindings for other HoAs.

5.2 – Benefits

The combined use of MIP and MPTCP in the manner described above provides several benefits.

First of all, a node can actually realise the theoretical benefits that come with multihoming: *improved resilience* and *increased throughput*. MPTCP naturally detects if one of the interfaces stops working and allows communications to continue over the other path(s) without a break. Besides, MPTCP naturally distributes traffic over the paths to make more efficient use of the available capacity on the several paths accessed through the different interfaces.

Moreover, a mobile node maintains continuous communications as it moves through areas served by dissimilar access networks. Handover is “smooth” since both (or all) the interfaces are used simultaneously, i.e. make-before-break handovers. In fact, better performance will be achieved when the handovers are “slow” i.e. when the both the old and the new CoAs are maintained for the longest period. This is radically opposed to traditional “fast” handover approaches, where the goal is to move to the new CoA as soon as possible, once it has been decided to handover.

Additionally, an IPv4 mobile node can effectively do route optimisation, which is not possible only with MIPv4.

Finally, on-going communications are not affected by a failure along the path through the home agent. By contrast, as discussed in Section 4, such a failure is a problem with MIP, even if this path is not being used for data.

The last two benefits arise because MPTCP is developing a security mechanism, for adding new paths (addresses) to an MPTCP connection that will not rely on a MIP-style Return Routability test.

6. Conclusions

In this paper we have proposed an architecture for the provision of multihoming support to 4G mobile nodes that uses a novel combination of Multipath TCP and Mobile IP. Such architecture enables the preservation of established communications through outages and a significant increase in the performance in the communication though the concurrent usage of multiple paths. The established connection will be able to survive as long there is at least one path available. Moreover, the performance of the communication will be improved as multiple paths will be used concurrently, and will be used according to the capacity available in each path.

Acknowledgments

This research was supported by Trilogy (http://www.trilogy-project.org), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme.

References

[1] A. Ford, C. Raiciu, M.Handley, TCP Extensions for Multipath Operation with Multiple Addresses, draft-ford-mptcp-multiaddressed-02.txt, 2009.
[2] D. Johnson, C. Perkins, J. Arkko, Mobility Support in IPv6, IETF RFC 3775, June 2004.
[3] C. Perkins, IP Mobility Support for IPv4, RFC 3344, August 2002.
[4] D. Wischik, M. Handley and M. Bagnulo, The Resource Pooling Principle, ACM/SIGCOMM CCR, 2009.
[5] H. Soliman, G. Tsirtsis, N. Montavont, G. Giaretta, K. Kuladinithi, Flow Bindings in Mobile IPv6 and NEMO Basic Support, draft-ietf-mext-flow-binding-04.txt, 2009.
[6] D. Wischik, M. Handley and C. Raiciu, Control of multipath TCP and optimization of multipath routing in the Internet, Proc. NetCOOP 2009.

