Mobile-IP Working Group
INTERNET-DRAFT
Expires: December 2003

Greg Daley
Monash University CTIE
JinHyoeck Choi
Samsung AIT
May 2003

Movement Detection Optimization in Mobile IPv6
draft-daley-mobileip-movedetect-01.txt

Status of this Memo

Abstract

This draft describes the state of the art techniques in movement detection and elaborates on their application to Mobile IPv6 networks. The aim of the draft is to describe the applicability of each mechanism and stimulate discussion on such techniques.

Table of Contents

1.0 Introduction

Seamless handoff systems aim at reducing the disruption caused by moving between IP networks. Movement Detection is a prerequisite task necessary for a Mobile IPv6 (MIPv6) Mobile Node (MN) to perform handover signaling[MIPv6-20]. As defined in the base MIPv6 specification, detection of movement is accomplished through reception of Router Advertisements (RAs), and comparison of these messages with previously received router information.

Other handover operations, such as Duplicate Address Detection(DAD) and movement binding signaling occur subsequently to movement detection. Since Movement Detection is a the first operation to occur in a non-predicted handover, Movement detection optimizations aim at reducing the time before the RA is received by the MN. This reduction of movement detection time reduces handover duration.

The MN inspects RA messages to determine if a router on that interface is providing routing information which supercedes or invalidates a current Care-of-Address (CoA). The CoA may be

invalidated because a router retracts a prefix (valid lifetime=0), timers expire which deprecate the existing CoA and Router entries or a heuristic is in place assumes movement if an RA is received with a new prefix, from a new router. In MIPv6, detection of movement causes configuration or selection of new Care of Addresses, and binding signaling is commenced.

2.0 The Current Status of Movement Detection

Movement detection is one of a series of stages in accomplishing MIPv6 handover. The section below indicates the steps required for handovers to occur.

2.1. Handoff Process

When an active MN moves to a new IP subnet, it changes its point of attachment to the network through the following handoff process.

First Link-layer handoff occurs to change the wireless AP to which MN is associated. After a new Link-layer connection is established, Network-layer handoff is performed, which broadly involves movement detection, IP address configuration and location update. Initially, the MN's IP protocol implementation may be unaware of the link change, or may have been informed of the arrival by a link- trigger. Alternatively, the network itself may be aware of the MN's movement and may make use of this information to aid movement detection.

The MN then checks the reachability of current AR. If the MN determines the AR is no longer reachable, the MN performs Router Discovery with new AR and subsequently a Router Advertisement with all options arrives from a new AR.

Once the MN discovers a new AR with necessary informations, stateless or stateful address configuration including DAD is performed[RFC-2462]. This entails waiting for address validation to complete, before further packets may be sent or received by the MN.

Mobility signaling procedures are then started, with the MN sending a Binding Update (BU) to its Home Agent. Additionally Return Routability (RR) procedures are started for route optimized conversations with Correspondent Nodes (CNs). As the Return Routability tests are completed, further BU messages are sent to CNs.

Since Movement Detection is prerequisite for other network-layer handoff operations, for seamless handoff, it is necessary to detect movement as fast as possible given the underlying link technology.

Proposals for predictive handovers change these assumptions and

ordering. The role played by movement detection in predictive handovers is defined in section 6.0.

2.2. Movement Detection Overview

The primary movement detection mechanism for Mobile IPv6 defined in [MIPv6-20] uses the facilities of IPv6 Neighbor Discovery, including Router Discovery (RD) and Neighbor Unreachability Detection (NUD).

In Movement Detection, an MN should check the (partial) reachability of the current AR and the validity of the current CoA. This allows the MN to distinguish between link-layer and network-layer handovers. In case of IP subnet change, an MN should also discover a new AR with necessary information, including on-link prefix to form a new CoA.

In brief, the Movement Detection process is as follows:

First there is some hint that MN may have moved to another subnet.
This hint itself doesn't confirm movement.

Then the MN probes the current AR to check its reachability and the validity of the current CoA.

If it turns out that MN has actually moved, it searches for a new AR with Router Discovery. After an RA from a new AR arrives with necessary information, the MN performs further operations, forming a CoA and sending Binding Updates.

There are 3 entities which may change in connection with MN movement, Access Point (Link-layer connection), Access Router and On-link Prefixes (IP Subnet).

These changes are indicated to MN with the following:
1) Link-layer trigger (Some information from lower layer which notifies link change)

2) a new IP address (in source address field of RA)

3) a new Subnet Prefix (in Prefix Information Option in RA)

To get the above indications, MN can perform NS/NA exchange, RS/RA exchange or just receive unsolicited RAs.

The source address of the RA is a Link-Local address, its uniqueness is not guaranteed outside a link. Hence the fact that MN receives

the RA with the same address doesn't guarantee that it comes from current AR.

ARs may omit a Prefix Information Option for efficiency. Hence the lack of the prefix of the current CoA in RA doesn't mean that current CoA is not valid.

The above defects may results in the problem like this. Assume MN moves from AR1 to AR2. If AR1 and AR2 have the same link-local address, the MN may not detect its movement even after several RAs.

The entities may change together or separately. Therfore any indications can't represent subnet movement precisely by itself.

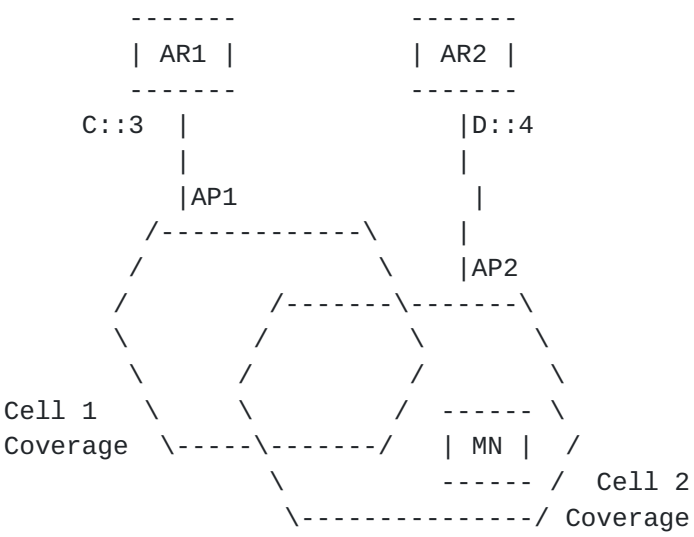

Figure 1. MN moving between two networks

For example, in Figure 1, MN moves from AP1 to AP2 and all three entities change.

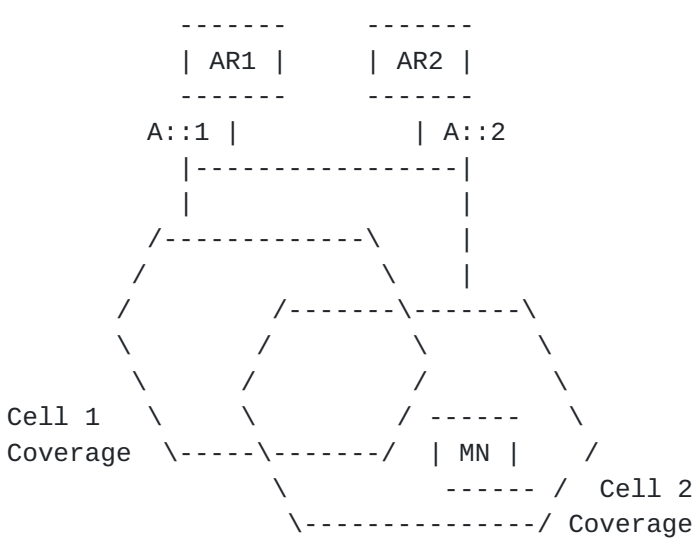

Figure 2. MN sees two access routers on the same subnet

In Figure 2, AR1 and AR2 are routers on the link. Each advertise a prefix A:: to hosts. When the MN moves from AP1 to AP2, it changes AP and AR but remains in same IP subnet.

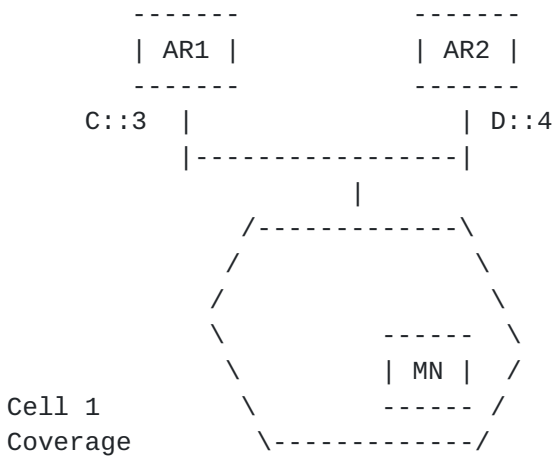

Figure 3. Two Access routers on the same subnet

In Figure 3, AR1 and AR2 are routers on the link which share a common AP. Each advertise a prefix C:: or D:: to hosts. Assume AR1 is MN's current default router. Then the arrival of a RA from AR2 doesn't mean MN has moved.

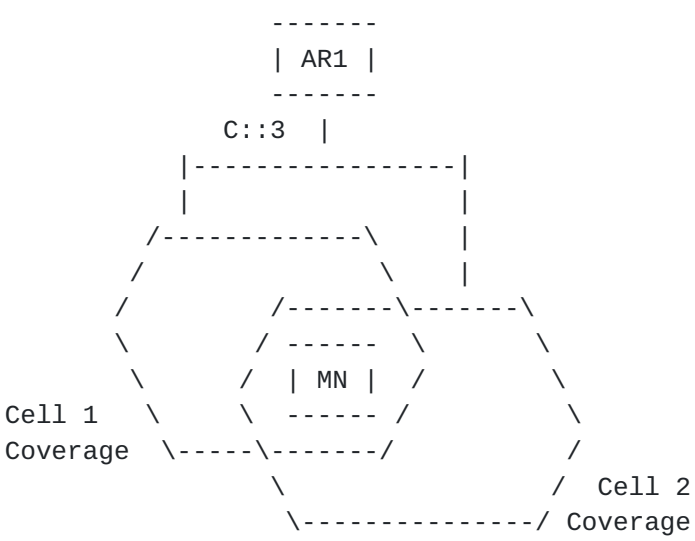

Figure 4. MN moves between wireless links in the same subnet

In Figure 4, If MN moves AP1 to AP2, it still remains at same IP subnet even though it receives link-layer change notification from lower layer.

An MN's Movement Detection scheme should combine available information to detect movement correctly. It should not mistake some hint as movement while the MN hasn't moved. That may result in continual handoff, and hence excessive mobility signaling. If the MN moves, it needs to detect movement sufficiently fast so that it can complete handover signaling without significantly degrading application performance.

On the other hand, if the MN doesn't move though it receives some hints (like Figure 3,4), it is not imperative to detect its non-movement so fast. It will not degrade performance even if MN can't quickly confirm that it still remains at the same subnet.

A movement detection scheme should not result in excessive signaling traffic. It should not flood the network with unnecessary RS/RA or NS/NA messages.

2.3. Movement Detection Mechanism

Movement Detection Mechanism consists of following steps.
Step 1. Hint
Step 2. Checking the reachability of the current AR
Step 3. Checking the validity of the current CoA
Step 4. Discovering new AR with necessary information.

2.3.1 Receiving Movement Hints

There is a set of hints which can indicate that Network-layer movement has occurred. The hints can be Link-layer trigger, the receipt of a new RA or the lack of RA from current AR. This hint itself doesn't confirm L3 movement.

2.3.2 Checking the reachability of the current AR

An MN checks whether current AR is still reachable by probing. The MN may probe with NS or RS.

In case of NS probing similar to NUD, an MN sends unicast NS messages to the current AR. If the current AR replies with a NA, the MN can be sure that it is still reachable. If an NA doesn't arrive after 1 sec, the MN resends the NS. After 3 probe attempts, the MN decides that the AR is no longer reachable.

If an MN actually has moved to new IP subnet, it will take 3 seconds to detect that the current AR is not reachable (sending 3 NS probes, plus waiting 1 second for each). With NUD, we can detect a node's presence very quickly. Conversely, it takes substantial time (3 Sec) to detect that node is NOT there.

In order to reduce the time taken to detect a router's non-presence, the MN may use a timeout. Instead of retransmitting, the MN just sends one NS, and waits for a reply for fixed time. If the MN times out before receiving a reply, it assumes that it has moved.

Attempts at avoiding the cost of NUD without resorting to eager bindings or NS/NA heuristics are discussed in section 4.3.

2.3.3. Checking the validity of the current CoA

When the NS/NA exchange is for the AR's link-local address, the MN can't be sure that it still remains at the same IP subnet since link-local scoped addresses uniqueness is only guaranteed on the link. The MN may have moved to a new AR which happens to have the same link-local address as the current AR.

Hence MN also has to confirm the validity of the current CoA by checking on-link prefixes. The MN sends a unicast RA to current AR. When a solicited RA with all options arrives, the MN checks whether

it contains the prefix of current CoA in any of its Prefix Information Options.

As an alternative solution, the MN may send the NS to a globally unique router address, if it is carried in a MIPv6 modified Prefix Information Option advertised by the current router. An NA response from this router address can uniquely provide reachability confirmation for the router, since only the current router may have this address.

2.3.4. Discovering a new AR

If it turns out that MN has actually moved, it has to find a new AR using Router Discovery[RFC-2461]. The MN sends a Router Solicitation to the All-routers multicast address. When a solicited RA with all options arrives, the MN selects a new AR, forms a new CoA and perform further operations.

We can perform Movement Detection Steps 2,3,4, with only one RS/RA exchange as illustrated below.

To check the reachability of current AR, instead of using NS, the MN sends an RS to the All-Routers multicast address.

If current AR is still reachable, MN will receive an RA with all options within roughly 1.5 Sec (1 Sec random RS delay and 0.5 sec random RA delay). Since RA messages do not explicitly indicate if they are solicited, we can't say that the current AR is reachable if we receive an RA. We can say though, if we don't receive a RA in time, it's highly probable that the current AR is not reachable.

If a solicited RA with all options arrives from current AR, the MN can confirm that current AR can still reach MN and current CoA is still valid.

If no RA arrives from the current AR, but the MN receives several RAs from new ARs, the MN chooses a new default AR.

Though the MN can't confirm reachability of the new AR, if its RA contains a Source Link-Layer Address option, the MN will gain a stale Neighbor Cache entry for the router. This means the MN can start sending packets. Moreover the solicited RA from the new AR contains all the necessary information for IP configuration. Hence the MN can perform further operations immediately without additional signaling messages or delay. After the handoff process is completed, the MN

can perform NUD with the new AR to confirm the reachability at its leisure.

Below is a comparison of the comparitive merits of RS/RA and NS/NA probing

The benefits and drawbacks of NS/NA probing are:

1) Since there is no Random Delay, MN and AR can send NS and NA immediately.

2) The solicited flag confirms bi-directional reachability.

3) The MN needs to perform at least one RS/RA exchange afterwards. (Unless a globally unique router address is probed).

The benefits and drawbacks of RS/RA probing are:

1) With only one RS/RA exchange, MN can check the (partial) reachability of a current AR, validity of current CoA and receive all the necessary information from a new AR.

2) There are two random delays of up to 1 Sec for RS and 0.5 Sec for RA. Hence it take more time to RS/RA exchange with RA timouts being set appropriately. This may be solved with FastRA (Section 3.5).

3) It may cause excessive multicast RA traffic.

4) MN can't confirm reachability of AR.

We may shorten the time taken to detect movement by performing multiple operations in parallel. For example, by sending NS to current AR and RS to all-routers multicast address at the same time, we can perform Step 2, 3, 4 simultaneously. If there are many wrong movement hints, this may cause excessive multicast traffic.

There is still much to investigate about the necessary steps of Movement Detection, their order of performance order, efficient (NS and RS) probing mechanisms and the trade-off between Movement Detection time and signaling traffic.

2.4. Movement Detection Performance with Neighbor Discovery

Movement detection algorithms are based on Neighbor and Router Discovery mechanisms[RFC-2461]. Neighbor Discovery allows solicitation of NA in order to confirm reachability. Router Discovery allows the periodic multicast of RAs to nodes on a (fixed) IPv6 network. [RFC-2461] additionally allows solicitation of RAs in order to confirm network identity, or speed device configuration.

Neighbor Discovery protocol constants were sized for networks of many nodes, where it was sufficient to provide configuration within a few seconds. This has caused significant delays to be built into Neighbor and Router Discovery[RFC-2461]. Networks supporting MIPv6 MNs need to be able to check (partial) reachability and receive RAs in shorter time intervals than are available for standard Neighbor Discovery. This is important if Mobile Nodes have existing higher-layer sessions when Movement Detection is performed, which may be affected by slow handover times.

Movement detection performance is measured from the time when the new Link-layer connection is established until Movement Detection is completed through suitable RA reception from new AR.

At first MN receives a hint that it may have moved. The time taken to receive for this hint varies. With Link-layer trigger support, it can be done instantly.

Alternatively an MN can monitor periodic RA beacons. The base MIPv6 document uses RA Interval Timer expiry as a hint. An MN may implement its own policy to determining the number of missing RAs which it will interprets as hint for possible movement.

With payload traffic tracking, we may get a hint earlier. An MN may implement its own policy to determining the interval of idle time with no traffic which it will interprets as a hint for possible movement. Care should be taken in this case to ensure that spurious hints do not cause unnecessary probing of the network.

Without schemes such as those above to provide hints, the MN must wait to receive an RA from a new AR before undertaking Movement Detection. Hence the detection delay depends on the frequency of Router Advertisements.

Periodic RA beaconing transmits packets within an interval varying randomly between MinRtrAdvInterval to MaxRtrAdvInterval seconds. In [RFC-2461] minimums for these values are 3 and 4 seconds, respectively. Since MN movement is unrelated to the advertisement

time on the new network, the MN is expected to arrive on average half way through the interval. This is about 1.75 seconds with [RFC-2461] advertisement rates. Worst case performance (without packet loss) is when the MN arrives just after an RA, and the next RA is scheduled close to MaxRtrAdvInterval. If [MIPv6-20] advertisement intervals are in use, these values drop to 0.025 and 0.70 seconds respectively.

Next the MN probes the current AR to check its reachability and CoA validity. There is no single agreed way mandated for this.

For example, assume the MN uses NS probing. If the MN probes with NS using NUD-like retries, the AR's unreachability will be detected after 3 sec for the case when the MN actually has moved. If the MN uses one NS and a timeout, the duration depends on the timeout value. We may set timeout value as 2 * RTT time from MN to AR. There is no consensus on timeout values yet and the RTT time in wireless environment may be highly volatile.

Afterwards, an MN should perform one RS/RA exchange, whether the current AR is reachable or not. This is subject to routers delaying 0-500 ms before responding to the RS, and the advice that the MN delays 0-1000 ms before sending an RS [RFC-2461]. Additionally, if there is no verified link-address available for Router Solicitation, the router must respond with a multicast RA.

Movement detection optimizations seek to lower the time taken to perform Neighbor and Router Discovery, either through MN solicitation, or timely unsolicited advertisement of router information. To achieve this aim, modifications to NUD and Router Discovery are made on mobile-supporting networks and MNs.

3.0 Movement Detection Schemes

3.1 Periodic Router Advertisement Beaconing

Beaconing is a term used to refer to multicasting of network identification information at regular intervals. Mobile IPv6 reduces the lower bound of the MinRtrAdvInterval and MaxRtrAdvIntervals to 30 and 70 ms respectively[MIPv6-20]. With these settings, beacons will be sent no more closely than 30 ms apart, and with no greater separation than 70ms. Routers are required to send the beacons at random times within this interval. This means that an MN will receive an RA within 70ms of arriving on the link, and may expect to receive an RA within 25ms, if we assume MN entry into the network to be randomly distributed in the interval.

This technique requires no action on the part of the MN other than listening to RA multicasts. The bandwidth consumption by multicast beacons is 14 kbps when RAs only include one Prefix Information option. Addition of a Source-Link-layer-Address option and a MIPv6 Advertisement Interval option typically increase this to 16.6 kbps.

On some networks, such overhead (~20 multicast packets per second) causes a serious burden on network bandwidth. In these cases, [RFC-2461] specified intervals SHOULD be used, if other movement detection mechanisms are available.

Additionally, the reduced interval between messages may have side effects for non-MIPv6 nodes on the same networks. The AdvDefaultLifetime value is used to set the lifetime of the default router in seconds, as advertised in the Router Lifetime field of the RA. The minimum value specified in [RFC-2461] for this value is MaxRtrAdvInterval. This value is less than one second when using MIPv6 specified advertisement intervals. Even if default router lifetimes are rounded up to the nearest second, nodes which assume MaxRtrAdvInterval is at [RFC-2461] values could be confused about the lifetime of their default router. Routers SHOULD ensure that AdvDefaultLifetime is greater than or equal to 4 seconds, in order to avoid this confusion.

3.2 RA caching in Link-layer Access Points (Fast Router Discovery [FRD-00])

One method which requires no solicitation from the MN is network triggering of RA. Router advertisements are sent to the MN when it attaches to an access point (AP) associated with this network.

In network deployments, the router may not be the link-layer device which the MN connects to, and therefore may be unaware of MN link-connection. Only in the case where the the AP advises the router of connection or AAA state, can the router send (unscheduled) unsolicited RAs before receiving packets from the MN.

The Fast Router Discovery (FRD) draft[FRD-00] places the responsibility of sending triggered RA messages upon APs. The Access Points cache RA's recently sent from the router, and deliver a frame to the MN when it connects. This frame is datalink-unicast to the MN and contains the most recent unsolicited RA.

In this case, less frequent transmission of unsolicited multicast RA messages may be used. At the same time, the first frame which is queued for the MN is the RA required for movement detection.

Deployment of FRD requires each of the APs for the network to be

capable of both the caching and triggered sending operations. Analysis and experimental results indicate that this is potentially the fastest network-layer based movement detection optimization, dependent on AP processing capacity.

3.3 Solicitation on Interval Timeout

As specified in MIPv6, the Advertisement Interval Option describes the maximum time between unsolicited RA messages (MaxRtrAdvInterval). This option is used in movement detection as a packet loss management system, where after elapse of one or more Advertisement Intervals without RA reception, the MN can send a router solicitation.

In the case where MNs send RSs after the loss of multicast RA messages, all MN's which have not received the RAs will time out at the same instant. [RFC-2461] specifies an additional delay of 0-1000 ms required for the purposes of desynchronizing RS messages sent from many hosts. An MN which does not have other confirmation that it has moved SHOULD follow this policy. Further details of this issue, especially with regard to simultaneous movement, are presented in section 4.1.

In the case where the MN moves by itself, this algorithm provides best performance when the MN connects to a new network just before an interval passes. If the MN is acting upon one packet's loss then this provides an RS immediately. If the timer elapses when there is no connection to a link, it is implementation dependent whether the RS packet is lost. Typically though, the MN will leave the previous access network on average half way through the mean RA interval. The expected time left until the Advertisement Interval elapses upon the MN leaving the network is:

```
ExpIval = 0.75*MaxRtrAdvInterval - 0.25*MinRtrAdvInterval
```

This does not account for the link-layer handover time, which may be tens or hundreds of milliseconds. When these values are the minimums described by [RFC-2461], this value is 2.25 seconds and therefore does not seem to provide significant benefit unless faster (MIPv6) beacons are being sent. At the MN specified beacon rate, the expected residual interval is 45 ms.

Any of the methods which require responses to Router Solicitations as specified by [RFC-2461] also incur an additional delay of between 0 and 500 ms before a response is sent by the router. Section 3.5 describes a mechanism to cope with this issue.

3.4 Link-up Triggers on the Mobile Node

For situations where RA packets have been transmitted correctly, the Advertisement Interval's best case occurs when the timeout occurs just after the MN joins a new link. In environments where the MN can receive an indication a link has been joined, this information can be used to trigger an RS immediately[MIPv6-20], although once again a random delay before sending RS's is advised by [RFC-2461].

Additionally, even if the MN doesn't send an RS upon receiving a link-up trigger, it can use the trigger to validate received RA messages for movement detection with eager-binding. The MN may be able to enter an 'eager-binding' state until it receives its first RA on the new link. If it receives an RA from a previously unseen router at this time, it may be useful to confirm bidirectional reachability with this outer, and then undertake movement signaling.

Upon entering a new subnet, there is a small chance that the MN will have a duplicate address collision with another device's Link-Local address[RFC-2462]. When an MN solicits an RA, it typically sends from the Link-Local address, unless this address is tentative[RFC-2461]. If the MN sends from the Link-Local address, unicast responses are allowed, and in this case, rate limiting of multicast RA messages is avoided.

If the MN joins a link, then until it knows the identity of the link (has received an RA), it MUST assume that it is on a new link, where its Link-Local address is tentative. This means that RS messages will either be deferred until DAD operations have been performed on the Link-Local address or the RS MUST be sent with an unspecified address. A multicast response will be scheduled no sooner than:

```
Max(LastMcastRATime + MinDelayBetweenRAs,
                                now + 0-500 ms RS Response Delay)
```

Where MinDelayBetweenRAs could be as high as 3 seconds. Even if the response is not multicast, the RS response delay is still incurred.

3.5 Fast Router Advertisement [FASTRA-02]

Fast Router Advertisement (FastRA) removes the random delay required of a router before it responds to RS messages. It relies upon only one router on a subnet being configured for FastRA, so that responses are not simultaneous. FastRA principally aims at delivery of unicast RA messages, since the rate limiting of multicast RA messages specified in [RFC-2461] specifies that RA messages may not be sent within 3 seconds of each other.

Similar action could be performed with multicast RA responses if FastRA adopted the MinDelayBetweenRAs as in [MIPv6-20]. In this case, the response would only be delayed if the last multicast RA occurred more recently than MinDelayBetweenRAs ago (or MaxFastRAs has been consumed). This could work well if the arrival of mobile nodes occurred much less frequently than the unsolicited multicast RA interval.

FastRA incorporates a rate limiting feature aimed at diminishing the potential effect of FastRA traffic on nodes which are already connected to the network. Routers may transmit no more than MaxFastRAs advertisements in an interval before discarding solicitations until the next unsolicited multicast RA.

Either of the solicitation mechanisms may be used to get FastRA response from a router, although Advertisement Interval timeout will only be invoked on packet loss if Link-triggers are available.

Movement detections times are bounded only by the time to send the Multicast RS message and send the unicast RA response. Recent testing has indicated 95% of RAs were received within 15 ms of sending an RS on 802.11b networks, when Neighbor Discovery was being performed on the MN's Link-Local address. If RS messages include Source link-layer Address options[RFC-2461] or are multicast responses with no timer delays, movement detection time will be lower.

4.0 Performance Considerations

4.1 Effects of Solicitation Delays

[RFC-2461] specifies that a node SHOULD delay a random interval of between 0 and 1000 (MAX_RTR_SOLICITATION_DELAY) ms before sending an RS if it is the first packet the node sends on the link [RFC-2461]. A similar delay is stipulated for DAD packets in [RFC-2462], for the same circumstances.

These delays are provided for the case where many devices are configuring on the link at the same time. In a mobility environment, this may occur if many MNs are traveling together, for example on a train, or at peak hours on a freeway. For this environment, there is some possibility that the MNs' simultaneous transmission of multicast RS or DAD packets will will cause interference or backoff and retransmission.

The effect of such simultaneous movement and subsequent multicast transmission is the topic of current research. On several wireless

technologies, the effect is thought to be minimal, especially where discrete codes or data channels are provided to each subscriber.

There are certainly other environments where many devices simultaneously transmitting have a detrimental effect though, and in these cases, the configuration by MNs SHOULD be serialized. The serialization provided by a random timer is one mechanism by which simultaneous transmission may be avoided. Other methods, are reliant upon serializing effects in the link-layer, such as AAA operations. These effects are link-dependent and where they provide protection, MNs SHOULD take advantage of them to avoid random timer delays.

It should be noted that multicast bombing may occur even in when no RS is performed, if many nodes simultaneously receive an RA beacon from a new router. These MNs' first operation is to undertake DAD procedures. Link procedures are unlikely to provide serialization in this case, since all MNs will receive the multicast at approximately the same time.

In any case, the MN SHOULD undertake the RFC-2461 and RFC-2462 prescribed delays if any of the following is true:

* The MN has no upper-layer sessions

* The MN has no sessions which have sent or received data within UPPER_LAYER_ACTIVITY seconds. The value of UPPER_LAYER_ACTIVITY is implementation specific, but defaults to 120 seconds.

* The MN has more highly preferred interfaces which have the currently bound CoAs configured for all current sessions. Also, these CoAs are known to be successfully receiving and sending data.

This limits the effect of link contention to active devices requiring an expedited handover service.

4.2 Performance Comparisons

A table is provided which indicates the relative performance of several movement detection schemes.

These handovers do not include delays due to DAD for unicast responses, nor do they include RS/DAD delays to avoid multicast link bombing. Additionally, the cost of determining reachability with the current AR is ignored.

Presented times are on a wireless link of ~2 Mbps, which is capable

of multicast at 1 Mbps.

	Uni/Multicast Advertisement	overhead (kbps)	Move Detection Time (ms) Avg	Max
Beacon	Multicast	>14	25	70
FRD	Multicast L3 (unicast L2?)	<1	<10	<20
Timer(a)	Unicast	<1	ExpIval +250	MaxRtrAdvInterval +500
LinkRS tentative	Multicast	<1	250	530
LinkRS unicast	Unicast	<1	250	500
FastRA tentative	Multicast	<1	<10	60
FastRA unicast	Unicast	<1	<10	<30

Notes:

(a) These duration values include link-layer handover time, where no other row does.

4.3 Avoiding NUD without eager binding

The cost of performing NUD to the current AR in order to check whether movement has occurred is expensive in the case that it has. Proposals to avoid using NUD with the current access router have been made in the mobile-ip working group.

One set of proposals which rely upon information in received router advertisements to guarantee that a previously configured router is uncontactable.

It is also possible to make use of link-layer information which indicates a network domain change. Link-layer triggers pass information to the network-layer which either explicitly provide movement detection[WLANPIO-00], or disambiguate subsequently received RA information.

Further reference will be made to drafts elaborating on these ideas as they become available.

4.4 Effects of Packet Loss

Packet loss from (network and MN) triggered systems can be a significant factor MIPv6 handovers performance. When a quickly delivered RA message is lost, then the MN may wait until the next unsolicited multicast RA message, which can take up to four seconds (RFC-2461 MaxRtrAdvInterval).

In MN triggered systems, if RS messages are lost and have to be resent, movement detection times are increased by up to four seconds. This timeout is governed by the value of the Neighbor Discovery constant RTR_SOLICITATION_INTERVAL.

In solicited RA environments, the exchange of RS/RA messages is susceptible to loss of either packet, as is the case with NS/NA if the router needs to perform Neighbor Discovery on the MN before sending a unicast RA response.

Beacon systems which transmit RAs at high rate are less susceptible to the adverse affects of packet loss, since replacement packets are transmitted quickly after the packet loss event.

Backup effects from Advertisement Interval timers may play a part in the solicitation of replacement RA messages, although unless link-layer handover times are considered, these provide worse performance than beacon based systems.

5.0 Combining Movement Detection Optimizations

When arriving on a network, an MN is unaware which movement detection optimizations are in place on the network, or whether any are in use. The MN may therefore choose to send an RS before it receives an unsolicited RA, even though either FRD or RA beaconing are in place. In all likelihood, both RA delivery mechanisms will be activated.

In this case, the movement detection time will typically not be affected, except that more packets will be sent to the wireless medium. Therefore, if an MN has received a link-trigger, and the MN subsequently receives an RA before it has scheduled an RS packet to be sent, the MN SHOULD NOT send the RS.

In most cases without packet loss, the presence of fast beacons will not significantly affect the performance of FastRA or FRD systems.

As mentioned in section 4.4 though, the harm caused by packet loss is significantly lowered if beacons are received within a short period. If the overhead of running beaconing systems is sufficiently low for the wireless link type, then beacons MAY be used with either of FRD and FastRA.

Networks with either of FRD or FastRA capability are unlikely to also use the other technology on their systems, since FRD and FastRA are closely matched in performance and have low latency times. If the network has both capabilities, there is some chance that RA messages from AP and Router attempt to be delivered simultaneously to the MN. Since there is no way for the MN to know that FRD is in use when soliciting, it MAY send an RS in any case, if it has not received an RA already from this link.

6.0 Predictive Handover Effects

Movement detection optimizations' applicability is principally in non-predictive movement environments, although there may be some benefit for anticipated/fast handover systems as movement confirmation and correction mechanisms.

Handover prediction allows MNs to select the network to which they move, and perform handover signaling in anticipation of this event. This allows the tunneling and buffering of MN traffic within network routers while the link-layer handover is occurring.

With some link environments, it may not be possible to guarantee that the MN will arrive on the selected link, nor that the MN has indeed arrived on that link. In these cases, it is still necessary to confirm that the MN is arrived on the link through movement detection algorithms. This allows the MN to send corrective binding signaling in the case that the network is a different one than was the anticipated destination.

7.0 IANA Considerations

No new message formats or services are defined in this document.

8.0 Security Considerations

Movement detection optimizations are reliant upon reception of a Router Advertisement with properly configured Prefix Information Options [RFC-2461].

Since Movement Detection is based on Neighbor Discovery, its trust

models and threats are similar to the ones presented in [SEND-psreq-01]. The attacks described in 4.0 of [SEND-psreq-01] can be applied to Movement Detection too. Movement Detection schemes SHOULD incorporate the solutions developed in IETF SEND Working Group if available, where threat assessment indicates such procedures are required.

Moreover the threats described in 4.2 of [SEND-psreq-01] may cause more serious problems. When there is an indication that the current IP connection has changed (Link down, New Router Advertisement et cetera), non-mobile nodes will first perform NUD[RFC-2461]. The MIPv6 handoff process (including Movement Detection) is time sensitive. So mobile nodes may start an eager handoff without Neighbor Unreachability Detection.

If higher layer notification of connectivity is not available, and eager handoff strategies are in place, any node or router which advertises an RA with a false prefix will cause MNs to perform spurious handover signaling and DAD operations.

For non-mobile case, if a node receives a bogus RA which doesn't include the prefix of its current address, it doesn't assume that its current prefix becomes off-link. In Neighbor Discovery, the only way to cancel a previous on-link indication is to advertise that prefix with the L-bit set and the Lifetime set to zero [RFC-2461]. Hence the node keeps using the current address and not so much harm is done.

In the mobile case, the threat is more serious. Assume an attacker sends an RA which includes only false Prefix Information Options. If a MN receives a bogus RA which doesn't include the prefix of the current CoA, it will assume that movement has occurred. The MN will start DAD (with the bogus prefix) and send BUs (with a false CoA). Hence MN will be disconnected, or its packets will be intercepted and subject to man-in-the-middle attack[SEND-psreq-01].

Moreover if we configure MNs to send RS and DAD without delay, this bogus RA attack may cause multicast bombing too. An attacker can send a bogus RA without source link-layer Address option. Then all MNs will receive the bogus RA at the same time and start Neighbor Discovery simultaneously. They will send RS without delay at the same time and cause RS congestion. An attacker can deceive all MNs to believe they have moved simultaneously by sending a suitable bogus RA. In this case, DAD would be performed by all nodes on the link immediately on multicast RA reception.

The security issues described above are not specific to any Movement Detection scheme presented in this draft but are inherent in any mechanism which uses only Router Discovery for movement detection.

Information from lower layers will be useful to mitigate the above threats. Assume there is an MN for which link-layer trigger is provided to notify link-layer change. If the link-layer trigger precedes a new RA, it is likely that the RA is valid and the MN has actually moved.

When the MN moves to a new IP subnet, link-layer change usually precede movement. So first link-layer change is notified to the MN and it anticipates movement. Hence when a new RA arrives, the MN can reasonably believe it.

On the other hand, if the MN receives a new RA without the notification of link layer change, it is likely that the RA is bogus. In this case, the MN SHOULD be suspicious:

Before initiating the handoff process, it SHOULD perform Neighbor Unreachability Detection to request a RA from its default router. Also, without link-layer information, RFC-2461 and 2462 delays before sending RS and DAD messages SHOULD be performed, until NUD has completed.

This document references several other documents, each of which defines its own security considerations. Readers are referred to these documents for further information.

Normative References

[RFC-2119] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. Request for Comments (Best Current Practice) 2119 (BCP 14), Internet Engineering Task Force, March 1997

[RFC-2461] T. Narten, E.Nordmark, W. Simpson. Neighbor Discovery for IP Version 6 (IPv6). Request for Comments (Draft Standard) 2461, Internet Engineering Task Force, December 1998.

[RFC-2462] S. Thomson, T. Narten. IPv6 Stateless Address Autoconfiguration. Request for Comments (Draft Standard) 2462, Internet Engineering Task Force, December 1998.

[FASTRA-02] M. Khalil, J. Kempf, B. Pentland. IPv6 Fast Router Advertisement (FastRA), Internet Draft (work in progress), October 2002.

[FRD-00] JinHyoeck Choi, DongYun Shin. Fast Router Discovery with RA Caching in AP. Internet Draft (work in progress), Feb 2003.

[MIPv6-20] D. Johnson, C. Perkins, J. Arkko. Mobility Support in IPv6. Internet Draft (work in progress), January 2003.

[SEND-psreq-01] P. Nikander (Ed.), J. Kempf, E. Nordmark. IPv6 Neighbor Discovery trust models and threats. Internet Draft (work in progress), January 2003.

Non-Normative References

[CGA-00] Tuomas Aura. Cryptographically Generated Addresses (CGA). Internet Draft (work in progress), February 2003.

[WLANPIO-00] Paul Tan. Recommendations for Achieving Seamless IPv6 Handover in IEEE 802.11 Networks. Internet Draft (work in progress), February 2003.

Acknowledgements

Thanks to the authors and editors of the MIPv6 (David Johnson, Charles Perkins Jari Arkko), FastRA (Mohammed Khalil, James Kempf and Brett Pentland), and FastRD (JinHyeock Choi {thx from Greg} and DongYun Shin) drafts. We have relied heavily upon their work and aim only to illuminate their good ideas. Additionally, we thank Ed Remmell and Erik Nordmark for their contributions in the working group. We're sure they'll recognise some of their ideas presented here.

Authors' Addresses:

Greg Daley
E-mail: greg.daley@eng.monash.edu.au
Phone: +61-3-9905-4655

Address:
Centre for Telecommunications and Information Engineering
Department of Electrical and Computer Systems Engineering
Monash University
Clayton 3800 Victoria
Australia

JinHyoeck Choi
E-mail: athene@sait.samsung.co.kr
Phone: +82-31-280-9233
Address:
i-Networking Lab, Samsung AIT (SAIT)

Appendix:

..

Changes Since Last Revision:

Since 00:

More diagrams indicating MD issues.

Stronger elaboration of MD mechanisms(NUD/NS/NA/RD)

Added stronger guidance for avoiding multicast bombing.

Added section on avoiding NUD safely (w/placeholder for potential references to LinkIDs draft &etc).

Added eager-binding heuristic after link-up trigger (EN).

This document expires in December 2003.