

Performance Study of the Mobile IPv6 Protocol and its Variations

George Aristomenopoulos*, Christos Bouras** and Kostas Stamos**

*Computer Engineering and Informatics Department, University of Patras, Greece

**Research Academic Computer Technology Institute, N.Kazantzaki Str., Patras Univ. GR-26500 Rion, Patras, Greece
and Computer Engineering and Informatics Department, University of Patras, Greece
aristome@ceid.upatras.gr, bouras@cti.gr, stamos@cti.gr

Abstract - In this paper, we use the OMNET++ simulator in order to evaluate the performance of the basic Mobile IPv6 protocol and some of its proposed variations. The most important metric we are interested in is the handover latency, which we measure for various combinations of the proposed Mobile IPv6 variations and then use this metric, combined with factors such as the complexity of the implementation, in order to evaluate and identify the best possible configuration for the operation of the protocol.

I. INTRODUCTION

Mobile IPv4 [1] has been considered as the de facto standard in providing mobility on the Internet. The rapid growth of Internet however has created needs for more efficient addressing, minimized handover latency and Quality of Service, that IPv4 cannot face. IPv6 promises to cope with these matters.

Although many platforms already support IPv6 and although it guarantees increased security, IPv6 and mainly Mobile IPv6 (MIPv6) [2] still remains a “work in progress” by the Internet’s standardization body, the IETF. One of the major challenges it faces are the so called handovers, or handoffs. The most common way for a handover to occur is the transition of a Mobile Node (MN) from one’s cell coverage area to another’s. During that time, MN loses connectivity with its cell and is unable to communicate with its peer. The longer the handover process, the greater the number of packets dropped. Various techniques have been proposed in the scope of MIPv6 standardization, promising minimal handover latency and consequently fewer lost packets.

In this paper, we use the OMNET++ [3] simulator in order to experiment with the variations on the standard MIPv6 protocol and their possible combinations. Our goal is to evaluate these proposals via simulation and to find a lower limit for the handover interval.

The rest of this paper is structured as follows: Section II gives an overview of the core Mobile IPv6 protocol and its proposed improvements, which are currently at various stages of standardization. Section III introduces the simulation environment we have used for our experiments, while section IV presents the results from our tests. Finally, section V sums up the paper with our conclusions and suggestions for future work.

II. MOBILE IPv6

According to MIPv6 [2] every mobile node (MN) is always identified by a Home Address (HoA), independently from its location on the network. When away from home, the MN is assigned a Care-of-Address (CoA), which indicates its current point of attachment to

the Internet. IP protocol offers binding mechanisms to register the Care-of-Address with the Home Agent (HA). From this point on any packets destined to the MN, are intercepted by HA and tunneled to the CoA. The time that lapses until the MN becomes aware of its movement, creates a new address and registers it to the HA is defined as handover latency. It is the primary cause of packet loss, resulting in performance degradation. As a result, numerous methods of minimizing the handover latency have been proposed, and a selection of them are presented below.

A. Layer 2 Trigger for Mobile IPv6

The technique of Layer 2 Triggers suggests the simple notification from Layer 2 to Layer 3 every time a change in the L2 link occurs. Since link up triggers are commonly implemented in all wired and wireless devices, the deployment of the Layer 2 Triggers technique is simple and inexpensive. So this simple notification can reduce the handover latency since it triggers the MN to start the handover procedures sooner.

However, a L2 handover is not always associated with a L3 handover, and therefore the L2 trigger risks being unnecessary. We expect this not to be a serious error, since the only cost is a RS and the corresponding RA which notifies the MN that it has not really changed its L3 network and no L3 handover needs to take place.

B. Fast Solicited Router Advertisements

During router discovery a router must delay the reply to a Router Solicitation message for a random period of time between 0 and MAX_RA_DELAY_TIME seconds (the default is defined at 500ms). The purpose is to prevent collisions and flooding of Router Advertisements (RA) when more than one router exist in the area. In order to achieve faster reaction times we allow at most one router in every connection to act as a fast router and respond instantly with unicast RA. A RA that is transmitted instantly is called a *Fast RA*.

C. Fast RA Beacons

The Fast RA beacons technique is based on the standard MIPv6 and does not propose any changes to the MIPv6 stack. MIPv6 defines that MIN_DELAY_BETWEEN_RAS seconds (default 3 seconds) should pass between successively multicast RAs. So an idea to reduce the handover latency is to force routers to send RAs more frequently than 3 seconds. Using the absolute minimum permissible values for the variables MinRtrAdvInterval (0.03 seconds) and MaxRtrAdvInterval (0.07 seconds) - which override MIN_DELAY_BETWEEN_RAS - allows routers to

send multicast RAs more often and as a result mobile nodes speed up the handover procedure.

D. Optimistic Duplicate Address Detection

The Optimistic Duplicate Address Detection [4] allows nodes to have and use a *tentative* address, i.e. an address whose uniqueness has not been verified yet. This is not a problem as long as all addresses are uniformly distributed, and so the DAD procedure almost always succeeds. (The possibility of an address collision is infinitely small). As a result all nodes with ODAD are capable of continuing their communications sooner.

E. Early Binding Updates

Mobile IPv6 uses a "return routability" procedure to verify a binding update with respect to authenticity and validity, and concisely the process holds two tests. A home-address test (HoT) which authenticates the mobile node, and a care-of-address test (CoT) which checks the validity of the new care-of address. A shortfall, however, is that the two address tests, though typically performed in parallel, constitute a considerable fraction of the binding-update latency. Both tests are potentially run over very long distances. Early Binding Updates [5] move these tests to a time when they do not hurt the overall protocol efficiency. So the MN runs the home-address test before every handover, if there is an anticipation mechanism, or periodically otherwise. In every case the mobile node has a new *Home Keygen Token* every time it changes a point of attachment and does not need to rerun it. Correspondingly a care-of-address test can be executed in parallel with sending data to and from the new care-of address.

F. Hierarchical Mobile IPv6

HMIPv6 [6] introduces a new type of node, the *Mobility Anchor Point* (MAP), which acts as a HA at the local network. The MAP is a router that maintains a binding between itself and a mobile node currently in its domain. It reduces signaling outside the local network and is recognized by the flag MAP in RAs. A MN entering a MAP domain will receive RAs containing information on one or more local MAPs. The MN can bind its current position (LCoA) with an address on the MAP's subnet (RCoA). Acting as a HA, MAP will intercept all packets on behalf of the MN and forward them to MN's current address. When a MN changes its point of attachment within a MAP domain, it needs to register only the LCoA with the MAP since the global address (RCoA) has not changed, making node movement transparent to correspondent nodes.

III. SIMULATION MODEL

The OMNeT++ simulator is an open-source, object oriented simulation environment. Combined with the IPv6Suite extension [7], which implements the core Mobile IPv6 protocol and its variations described in the previous sections, we are capable of studying the previous techniques. For our research we have created a topology which is realistic enough but also not very complicated, which consists of 10 different subnetworks located next to each other with small overlaps. The mobile node starts from its home network and moves with various speeds, crossing each subnetwork, one every 150 meters. In every experiment the mobile node *client* crosses all 10 access points, while transmitting numbered

Internet Control Message Protocol (ICMP) ping messages to the correspondent node *server* every 10 msec. The *server* doesn't reply, but only records the incoming messages. Using these records we can measure the packet loss and therefore the handover latency.

Specifically as dropped packets we count packets $p_2 - p_1$, where p_1 is the sequence number of the last arrived packet to the CN before the handover and p_2 is the sequence number of the first packet that arrives after the handover. Packets with wrong numbering are ignored.

As handover latency we define the minimum between the interval of dropped packets, and the time between the L2 handover and the receiving of a new Binding Acknowledgment (BA) message confirming the new care-of-address. Exceptions are the cases of HMIPv6 where we stop timing on receiving the LBA (Local Binding Acknowledgment), and the case of ODAD as we stop on receiving the BU, since the NCoA is used instantly.

During the simulation we ignore the first handover and we only study the rest of the handovers (9 in total). The reason is that we want to isolate various unpredictable factors and initializations allowing us to be more precise with every technique.

IV. EXPERIMENTS

A. The effect of the mobile node's speed

We will examine initially if and how much the handover latency and packet loss are influenced by the speed of the mobile node. We run our model for every technique adjusting the mobile node's speed from 1 m/sec to 20m/sec and recorded the results. For greater accuracy we run each experiment 10 times with a different seed for the random functions, recording a total of $10 \times 9 = 90$ handovers for each case.

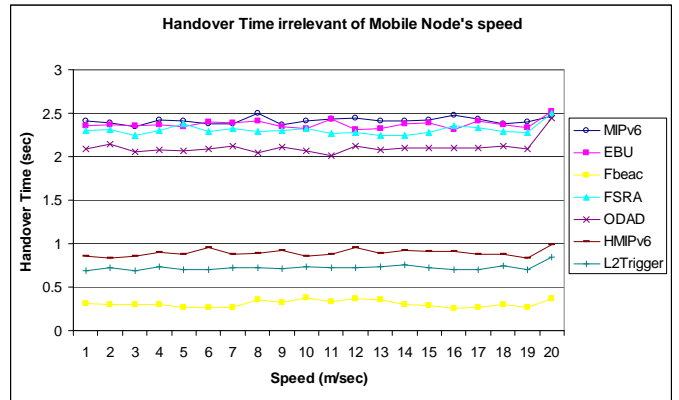


Figure 1. Relation between speed of MN and handover latency

The graph of Figure 1 shows the relation between the mobile node's speed and the handover latency. It appears that according to our results the speed of a mobile node is mostly irrelevant of the handover latency for every tested technique. This observation should not be surprising if we take into account of the processes that take place in order for a L3 handover to be completed. More specifically, in order for one node to register to a new network, messages have to be exchanged between the mobile node and the corresponding router, which are not influenced by small speeds such as those of the mobile node. These messages are exchanged in orders of

magnitude faster than the mobile node is speeding between subnetworks, and therefore its speed can not be fast enough in order to make a difference.

B. Basic Extensions

Given that the handover latency does not depend on the speed of the mobile node, we have fixed the speed of the mobile node for the following experiments at 10 m/sec (36 km/h), a very realistic speed, which could represent a vehicle driving through a city. We first try to evaluate each one of our techniques independently, by studying packet loss and handover latency.

As mentioned before, we run every experiment 10 times with different seed every time and record the results. Table 1 shows the average handover latency and packet loss per handover that we measured. Below we shall study the behaviour of each MIPv6 extension in detail.

TABLE 1. SUMMARY OF RESULTS

Technique	Avg (sec)	Std Dev	Packet loss
Fbeac	0.306906	0.039364	3.1
L2Trigger	0.722721	0.033074	7.2
HMIPv6	0.893267	0.040732	8.9
ODAD	2.107279	0.084238	21.0
FSRA	2.305996	0.057003	23.0
EBU	2.371481	0.050416	23.7
MIPv6	2.414414	0.038788	24.2

C. Optimistic Duplicate Address Detection

The ODAD technique offers a 310ms improvement in handover latency compared to the basic MIPv6, as we can observe from our results. We would expect greater improvement given the faster CoA shaping, eliminating the time DAD needs, which is one second by definition. This happens because ODAD is being used whenever a new network prefix is advertised by an access router, i.e. whenever a new RA with a different network address is received. But the mobile node is unable to understand the presence of a new subnet unless the missed RA movement detection mechanism has been triggered. This means that for our model $(\text{MaxConsecMissedRtrAdv}+1) \cdot \text{MIPv6MaxRtrAdvInterval} = 2 \cdot 1.5 = 3$ seconds have to pass, before the MN becomes aware of its movement. This is enough time for non-ODAD nodes to receive a new RA and complete the DAD procedures before the movement detection mechanism has been triggered.

A change that might help would be the decrease the variable $\text{MaxConsecMissedRtrAdv}$ to 0. By this the movement detection mechanism would be triggered as soon as one RA is not received within the expected time interval.

D. Fast Solicited Router Advertisements

As it appears from Table 1, the use of FSRA brings an improvement of roughly 90 ms at the handover latency compared to basic MIPv6, while the expected value should be around 250ms (the average mean of 0 and MAX_RA_DELAY_TIME), given that the waiting time before responding to a RS is 0. However FSRA works only when the MN sends a RS, this is when the MN detects movement to a new subnet. The small improvement is due to the fact that sending a RS is random since the missed RA movement detection

mechanism depends on the random unsolicited RA interval, which varies from 1.5 to 3 seconds. This interval is big, and in many cases overlaps the time needed by the standard procedures of MIPv6. This explains the increased packet loss per handover for the FSRA technique.

E. Fast RA Beacons

Observing the handover latency and packet loss for the Fast RA beacons technique, it becomes evident that a reduced interval between unsolicited RA can introduce remarkable results (1.8 seconds). As explained before an unsolicited RA will trigger the movement detection mechanism. Therefore by sending RAs in average every 50ms (the average neam of $\text{MIPv6MinRtrAdvInterval}$ which is 30ms and $\text{MIPv6MaxRtrAdvInterval}$ which is 50ms), Layer 3 is capable of detecting movement really fast and start the procedure for forming a new CoA.

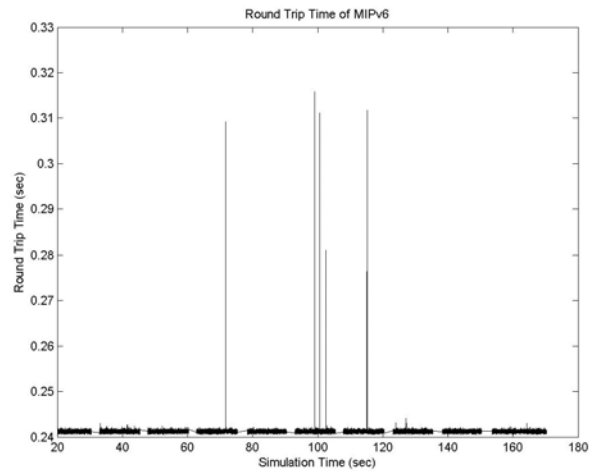


Figure 2. Round Trip time for the basic MIPv6

To send RAs every 50 ms, means that each second 20 RAs exist on the medium, 25 times more RAs concerning the 0.8 RA that are sent under regular conditions. Something like that could cause increased collisions between the ping and RA packets. We perform further experiments recording the Round Trip Time of ping packets for both MIPv6 and Fast Ra Beacons cases. In Figure 2 we see our result for a random run of our model when using MIPv6 and in Figure 3 when using MIPv6 with Fast RA Beacons.

We observe in Figure 2 that the distribution is almost uniform, with minimal variance and that generally the RTT is constant. The few spikes in our graph are a result of the not-ideal specifications of the Physical Layer that is simulated in our model. Examining now the Fast RA beacons case, we notice that reducing the interval between successively RAs from 1-1.5sec to 30-70ms produces many spikes, caused by the collisions of ping and RA packets. This confirms our initial assumption that the probability of packet collisions has increased due to the increased frequency of transmitting RAs.

We can therefore conclude on the Fast RA Beacons technique that it is not suitable for large networks with many moving nodes as the conflicts are increased, as also it is not suitable when transmitting streaming multimedia content since the requirements in the variance should be

constant and specified, something that is not possible under these circumstances.




Figure 3. Round Trip time for MIPv6 with Fast RA beacons

F. Early Binding Updates

The Early Binding Updates technique promises at least a RTT improvement in handover latency compared to MIPv6. Our results failed to verify this promise, and gave a slight improvement of 33 ms compared to the expected 240 ms. In order to understand this result we have to consider the way the EBU technique works.

According to EBU, a mobile node should perform a Home-Address Test periodically every 3.5 minutes, or every time it discovers a new subnet, in order to reduce the time needed for the new CoA binding. In our scenario EBU fails in every case. The detection of a new subnet and so the beginning of a Home-Address Test depends on when the movement detection mechanism is triggered, i.e. 1.5-3 seconds, which offers no major advantage regarding standard MIPv6. On the other hand the 3.5 minute period is too large for our model, in which a new handover occurs every 15 seconds on average.

A solution that could improve the outcome of the EBU technique in micro-cellular networks, such as ours, could be the execution of much more frequent Home-Address Tests, in expense of additional overhead in the network.

G. L2 Triggers

The performance of the L2 Trigger technique is also remarkable. The Layer 3 notification of a potential handover that triggers the MN to send a RS, can spare 1.7 seconds from the handover latency as our experiments show.

Apart from the fact that it drastically reduces the handover latency, the L2 Trigger technique has another big advantage. Its import and use from a mobile node is simple and economically inexpensive. It only demands an update of the MIPv6 stack, in order to forward the “L2 Link Up” message to the upper layer 3.

As we already discussed in section II.A, it is possible for the L2 triggers to be misleading. This happens when the mobile node moves between different Access Points which are connected to the same subnet interface and therefore no L3 handover is going to take place. As a result the mobile node would send unnecessary RSs, increasing traffic on the network. Our results however

verify our assumption that the benefit of the technique overcomes the cost of the unnecessary network traffic created.

H. Hierarchical Mobile IPv6

As explained before, in this case a MAP acts as a HA on the local network. Therefore the mobile node addresses its BU to the local MAP and not to HA and CNs who are probably further away. As a result, the handover latency is reduced by at least 1.5 round-trip time seconds, since the return routability process is not needed for every CN. The theoretical reduction of the handover latency for our model is $1.5 \cdot 0.24 = 0.36$ seconds. But our results show a much higher improvement of 1.5 seconds, just 190ms more than the L2 Trigger technique.

Analyzing our results for the first handover, i.e. the handover during which the MN enters the MAP domain, the handover latency is 1.8 seconds on average. We have to note here that we exclude the first handover from our calculations since we study the behavior of every technique after initialization. According to HMIPv6 the mobile node should bind its CoA with an RCoA of the MAP domain and advertise that address to CNs and HA.

But while the MN moves inside the MAP domain all is needed is to announce its new LCoA to the MAP. The RCoA remains unchangeable, the movement is transparent to the CNs and no more bindings are required. For this reason handover latency and packet loss are significantly reduced.

Hierarchical Mobile IPv6 achieves a considerable reduction of the handover latency, similar to the L2 Trigger technique, but in the expense of great complexity. The use of HMIPv6 requires several changes at the MIPv6 stack, like adding new message types, as and manual configuration of every access router that would act as a MAP.

I. Combination of Techniques

TABLE 2. COMBINATION OF TECHNIQUES RESULTS (SORTED IN ASCENDING PERFORMANCE)

Technique	Handover Latency	Packet loss per handover
L2Trig Fbeac FSRA ODAD	0.3109	31
Fbeac	0.3108	31
L2Trig EBU Fbeac FSRA ODAD	0.3105	31
L2Trig Fbeac ODAD	0.3101	31
Fbeac ODAD	0.3095	30.9
EBU Fbeac FSRA ODAD	0.3092	30.9
L2Trig EBU Fbeac FSRA	0.3073	30.7
Fbeac FSRA ODAD	0.3071	30.7
EBU Fbeac ODAD	0.3064	30.6
L2Trig EBU Fbeac ODAD	0.3055	30.5

After we have analyzed the performance, the advantages and the disadvantages of each technique separately, now we will try to find those combinations that constitute the better solution for the minimization of the handover latency and packet loss, keeping always in mind the complexity and the usage cost of each technique. We repeat the already discussed experiments, using the same specifications (a mobile node crossing the

topology with a speed of 10m/s sending ping packets every 10ms), for every possible combination of the available techniques. We run our tests for all 48 possible combinations and we present the top 10 in Table 2.

First of all, we observe that the overall reduction of handover latency is not necessarily equal to the sum of every technical improvement. Indeed, there are cases where the combination of certain techniques has decreased performance compared to each technique on its own (for example, the combination of FSRA and HMIPv6 has worse performance than HMIPv6 on its own).

According to our results, the best performance/cost ratio is achieved by the Fast RA beacons technique. We note that even when it is used alone has remarkable results. As mentioned before though, the use of Fast RA beacons results in increased variance of RTT due to RA and ping packet collisions. So we end up to a dilemma. We have to either put up with the higher jitter for the sake of better handover latency, or we have to use the next best technique that makes no use of Fast RA Beacons and has double latency. (The combination of L2 Triggers and HMIPv6 results in 0.62 seconds handover latency).

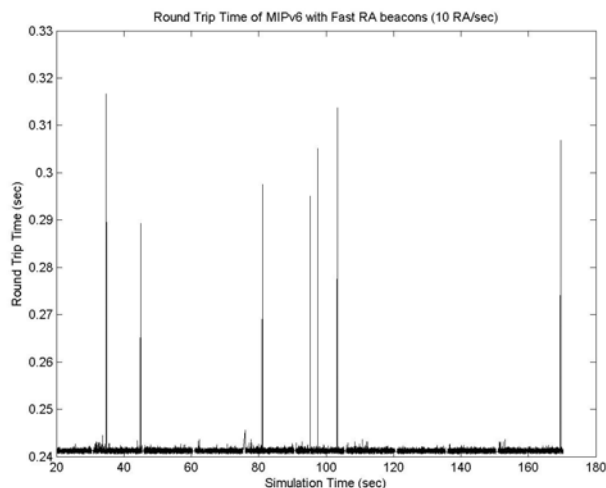


Figure 4. Round Trip time for MIPv6 with 10 RAs per second

This dilemma is crucial and depends on many factors, such as the kind of traffic through the medium, the available bandwidth etc. Consequently it is tricky to come to a forthright conclusion. So we perform another experiment using Fast RA Beacons. We record the RTT of the ping packets, altering the values of the variables `MIPv6MinRtrAdvInterval` and `MIPv6MaxRtrAdvInterval`.

For the results of Figure 4 we have set the variables `MIPv6MinRtrAdvInterval` and `MIPv6MaxRtrAdvInterval` to values 60ms and 100ms respectively. So 10 RAs per second are transmitted on average. We observe that the spikes on the graph have visibly decreased in relation to Figure 3, i.e. when transmitting 20 RAs per second. Furthermore we notice that the graph is almost identical to the one on Figure 2, concerning the standard MIPv6 protocol. We have managed to reduce the jitter of the traffic by reducing the

number of transmitted RAs per second, and accordingly we have boosted the performance of our network.

The major question however still remains; which is the measured handover latency? The results are encouraging, since the mean handover latency was calculated to 0.336 seconds, 26ms more than the Fast RA Beacon experiment with 20 RA/sec. This is a very small overhead considering the associated advantages.

Given these encouraging results we perform another experiment, reducing still further the RA rate, at 5 RA/sec. In this case, although the variation is similar to the previous experiment, the handover latency is however increased by 130ms, at 0.46 seconds. This result means that we can safely ignore this protocol configuration: Compared with the previous configuration, it offers no improvement, and instead it has worse performance.

Therefore, according to our simulations, in order to have the minimum handover latency in combination with the smaller possible jitter, we are led to use the Fast RA Beacons technique, altered so as to transmit 10 RAs per second instead of 20 which is the lower allowable value. This solution has no cost of use as long as no new elements are imported to the MIPv6 stack; all it requires is the appropriate router configurations.

V. CONCLUSIONS AND FUTURE WORK

Using a simulator we measured the performance of various techniques and concluded that in order to attain the smaller handover latency, the Fast RA Beacons technique is a very effective improvement. It is capable of reducing the handover latency to 300ms in comparison to the 2.5 seconds that standard MIPv6 offers. Its use in combination with any other technique produces evenly good results, but not that good so as to justify any added cost or complexity. Furthermore we showed that the extra unwanted jitter that Fast RA Beacons import can be faced by reducing the number of transmitted RAs per second to half, without losing in performance.

Further work in this area should examine the results of this paper in a real-world setting, where other factors related to the quality of the implementation might also be significant. The experimentation with realistic traffic patterns of traffic such as real-time applications and the effect of handover latency on their quality is also a field of future study related with our presented work.

VI. REFERENCES

- [1] C. Perkins "RFC 2002 IP Mobility Support", Internet Engineering Task Force, October 1996
- [2] D. Johnson, C. Perkins, J. Arkko "RFC 3775 Mobility Support in IPv6", Internet Engineering Task Force, June 2004
- [3] A. Varga. The OMNeT++ discrete event simulation system. In Proceedings of the European Simulation Multiconference (ESM'2001) Soc, 2001.
- [4] N. Moore, "RFC 4429 Optimistic Duplicate Address Detection (DAD) for IPv6", Internet Engineering Task Force, April 2006
- [5] C. Vogt, R. Bless, M. Doll, T. K'fner, "Early Binding Updates for Mobile IPv6", Internet draft, Internet Engineering Task Force, February 2004
- [6] H. Soliman, C. Castelluccia, K. El Malki, L. Bellier, "RFC 4140 Hierarchical Mobile IPv6 Mobility Management (HMIPv6)", Internet Engineering Task Force, August 2005
- [7] J. Lai, E. Wu, IPv6Suite Simulation Framework, Monash Univ., <http://ctiware.eng.monash.edu.au/twiki/bin/view/Simulation/IPv6Suite>