

Decreasing Mobile IPv6 Signaling with XCAST

Thierry Ernst

Keio University, Jun Murai Lab *
5322 Endo, Fujisawa-shi, Kanagawa 252-8520, Japan
<http://www.sfc.wide.ad.jp/ernst>
ernst@sfc.wide.ad.jp

ICOIN, Jeju, Korea, Jan.31-Feb.02
LCNS 3391

<http://www.springerlink.com/index/U7JNGH7XK1GB06F2>

Abstract. Mobile IPv6 is the IETF proposition to support host mobility in the Internet. It provides routing optimization for packets sent to the mobile node at the expense of signaling messages that are periodically sent to the correspondent of the mobile node. We propose using XCAST, a multicast technique best designed for small groups, as a means to decrease this signaling. XCAST records the addresses of all the destinations in the packet itself. We evaluate the performance of our proposition against Mobile IPv6 route optimization. We obtain good results even when a very few routers, well located in the network, are able to process the extension.

1 Introduction

Fixed nodes are permanently attached to the same subnetwork and are identified by a permanent IP address which determines the subnetwork where they are attached to. Unlike fixed nodes, mobile nodes (MNs) change their point of attachment in the Internet topology. They are moving from subnetwork to subnetwork and are reachable at different locations in the Internet topology. As a result from this, MNs must change their addresses. However, the IP address has a dual semantic at the network layer (used both as a locator and as an identifier) and is also used at upper layers as a node identifier. There is a therefore a trade-off between retaining the IP address which fails routing and changing the address which breaks upper layer connections [1, 2]. Mobile IPv6 (MIPv6) [3] proposes *two-tier addressing* as the solution to this conflicting dual semantic and use of IP addresses. *Two-tier addressing* associates a MN with two addresses,

* The ideas and simulations presented in this paper were originally performed under the supervision of Hong-Yon Lach (Motorola Labs Paris, France) and Claude Castelluccia (INRIA Rhône-Alpes Grenoble, France) and co-financed by Motorola Labs Paris, France and the French government under a PhD program. The paper was later written at Keio University. The author is also grateful to Christophe Janneteau for his careful review on an earlier version of the present paper.

one is permanent and used as a location invariant identifier, and the other one is temporary and used for routing. An address translation mechanism offers migration transparency to upper layers and insures backward compatibility with transport protocols. Connections are not disrupted as a result of mobility. Interestingly, Mobile IPv6 allows routing optimization (direct routing) for packets sent to the MN. This saves bandwidth and preserves delays, at the expense of signaling messages to update the current location of the MN.

Table 1. Traditional Multicast vs Explicit Multicast

Metric	Traditional Multicast	Explicit Multicast
Group Members	Unknown to the source	Known to the source
Identifier	Group identified by single address	No group identifier
Membership Management	Protocol needed	No group membership management
Distribution Tree	Multicast routing protocol responsive for distribution tree establishment	No need to build a tree, use standard unicast routing table
Multicast Address	Multicast address discovery	No multicast address discovery
Packet Overhead	No overhead	List of group members self contained in the transmitted data packet
Packet Processing	Check forwarding interface for each branch of the tree	Check forwarding interface for each destination
Signaling	Large amount of signaling	No signaling
Memory	State in each multicast router	No state in routers
Scalability	Large number of group members	Large number of small groups

Multicast aims at minimizing bandwidth consumption when there are multiple destinations for a given packet. Hence, the aim of multicast is to avoid duplicate information flowing over the same link. The packet is only duplicated when all destinations are not reachable via the same next hop. In the traditional multicast model, as defined in [4, 5], a *multicast address* is assigned to a collection of nodes that form a *multicast group*. A *multicast routing protocol* constructs a *multicast delivery tree*. We commonly distinguish two kinds of multicast delivery tree, the *Shortest Path Tree* (SPT), and the *Shared Tree*, or *Core-Based Tree* (CBT). The SPT is a *minimum spanning tree* rooted at the source. Each source in the group has its own SPT. The CBT is a single delivery tree built per multicast group, and is shared by all senders in this group. This tree is rooted at a single *core* router. Packets are thereafter sent to the multicast address and forwarded along the delivery tree.

Explicit Multicast (Xcast) [6–9] (also known as *Small Group Multicast* or *List-Based Multicast*), is an orthogonal and recent multicast technique, designed to complement traditional multicast. The basic idea is to carry the list of recipients of a packet in the packet itself. Intermediate routers must read the list of destinations to check if they have distinct next hops. Compared to traditional multicast, Xcast is very simple. There is no multicast routing protocol, no multicast address, and no group membership protocol. Both techniques are indeed complementary to one another since a "one size fits all" protocol seems unable to meet the requirements of all applications. As shown on tab.1, Xcast seems more appropriate for a large number of multicast groups with a small number of members, whereas *traditional multicast* is more appropriate for a large number of group members. Applications of Xcast include *narrowcast-like* (*few-to-few*) applications (e.g. IP telephony, collaborative applications, etc), whereas tradi-

tional multicast is targeted to *broadcast-like (one-to-many)* applications (e.g. TV and radio programs, weather forecast, etc).

In this paper, we propose Xcast as a means to carry Mobile IPv6 signaling while maintaining routing optimization. Mobile IPv6 is outlined in section 2 whereas the Xcast extensions brought to Mobile IPv6 are described in section 3. The performance of this solution is evaluated in section 4.

2 Mobile IPv6 and its Shortcomings

Mobile IPv6 associates a MN with two addresses. The Home Address MN_{HoA} is permanent and obtained on a link in the home network (home link). In addition, the MN obtains a new temporary Care-of Address MN_{CoA} on each subsequent visited link. This terminology is illustrated on fig.1. The MN may own several Care-of Addresses at anytime, one of which is selected as the primary MN_{CoA} . The binding between MN_{HoA} and the primary MN_{CoA} is registered with the Home Agent (HA), a special node on the home link. This registration is performed by means of a *Binding Update* (BU) message containing both addresses. Once it receives a BU, the HA adds or update an entry in its *Binding Cache*. The Home Address is used as the key for searching the *Binding Cache*. As a result of this registration, the HA is able to intercept all packets intended for the MN and to encapsulate them to the current Care-of Address MN_{CoA} .




Fig. 1. Mobile IP Terminology

A correspondent node (CN) willing to communicate with a MN first calls the *DNS* which returns the MN's Home Address. Packets are then routed to the home link where they are intercepted and encapsulated by the HA to MN_{CoA} . The packet is decapsulated by the MN and the CN is inserted in the *Binding List*. At this point, the MN may also send a BU containing its primary MN_{CoA} to some or all CNs recorded in its *Binding List* to avoid triangle routing via the HA (fig.2.a). The CN authenticates the packet and adds an entry in the *Binding Cache* like the HA. Forthcoming packets are then directly sent to the MN_{CoA} using an *IPv6 Routing Extension Header*.

BUs are sent alone in separate packets containing no payload, but according to [3], BUs may be piggybacked in payload datagrams. However, this has so far not been specified. Typically, the MN sends 5 consecutive BUs with a 1-second interval just after forming a new primary Care-of Address. After this, the MN keeps sending BUs at a lower rate, typically every 10 seconds, in order to refresh *Binding Cache* entries.

Routing optimization from the CN to the MN is one of the most interesting Mobile IPv6 features, particularly when MNs are not located in their native administrative domain. In this situation, simulation results in [10] have shown that the mean distance, expressed in number of hops, when routing optimization is used between the CN and the MN, is half the distance via the HA when routing optimization is not used. Not only the *transmission cost* (total number of bytes transmitted over the network) increases, but also the mean delay. Simulations have also outlined that the distance varies much more on the triangle route than on the direct route, thus further exhibiting the need for routing optimization. When the transmission cost is compared against the *mobility cost* (total number of bytes consumed by Mobile IPv6 signaling), the simulations show that the aggregated cost remains less important over the direct route with routing optimization than over the triangle route with no routing optimization, even for a very low data rate between the CN and the MN.

However, since BUs are sent periodically, we have also observed that the MN sends short packet bursts, separated by silent periods on the order of several seconds. This periodic burst is propagated to the entire network and consumes a significant amount of the available bandwidth in situations where the MN has a large number of CNs. This is particularly true on the wireless link between the MN and its Access Router when the MN has formed a new Care-of Address (5 consecutive BUs are sent to each CN). This problem is referred to as a *binding update explosion* in [10] or a *binding update storm*.

We conclude that while there is no doubt that Mobile IPv6 provides for optimal routing advantageously, the transmission cost has to be balanced against the mobility cost which increases linearly with the number of CNs. Solutions to minimize this mobility cost must thus be looked into.

3 XCAST Delivery of Binding Updates

In this section, we propose extending Mobile IPv6 with Xcast to deliver BUs as a means to minimize signaling. The addresses of several CNs are recorded in the BU, instead of sending individual BUs. This only requires minor extensions to Mobile IPv6. For doing so, we define a new *Xcast Header*, implemented as an *Hop-by-Hop Extension Header*. As such, it should be processed by all routers that understand this option. Its length varies according to the number of CNs, number that may only be limited for performance considerations. Each destination recorded in the header corresponds to a rank in a *bitmap* field. When set, a bit indicates that the corresponding destination still remains undelivered. This field must be updated by each duplicating router. Its purpose is to prevent

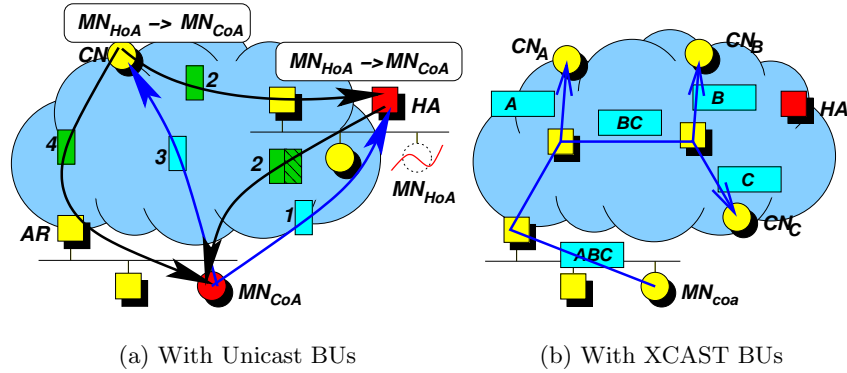


Fig. 2. Routing Optimization with Mobile IPv6

loops and to avoid the processing of destinations to which a duplicate packet was already transmitted. A *remove flag* specifies if destinations not reachable from the interface where a BU is going to be forwarded must be removed. The MN must be able to fill the *Xcast Header* and is provided a basic decision mechanism to decide whether BUs should be sent individually or by means of Xcast. In the latter case, the *Xcast Header* is filled with the corresponding CNs as found in the *Binding List*. The *header length* and *bitmap* fields are filled appropriately. CNs may also be split into separate Xcast BUs.

The *Xcast Header* must be processed by a number of routers and the CNs, as illustrated on fig.2.b. A single BU is sent to *A*, *B* and *C*. On receiving a packet with an *Xcast Header*, a Xcast-enabled router checks if there is still destinations to which the packet remains undelivered. If all bits in the *bitmap* are unset, no duplication is required anymore. For each bit set, the Xcast-enabled router reads the corresponding address and interrogates its routing table for ascertaining the next hop towards this address. As the next hop to *A* diverges from the next hop toward *B* and *C*, the packet is duplicated as many times as there are distinct next hops towards the destinations. Useless destinations are also removed to decrease packet overhead, and the *bitmap* is set appropriately. A copy is transmitted on the interface toward *A* and one on the interface toward *B* and *C*. In circumstances where Xcast-enabled routers are not widely deployed, a CN may receive a BU with a number of CNs remaining in the header. In such a case, processing the *Xcast Header* at the CN ensures BUs' delivery to all CNs. On receiving such a BU, the CN first checks the *bitmap*. If there is more than one bit set, the BU is duplicated, the bit corresponding to the CN is removed, and the packet is forwarded.

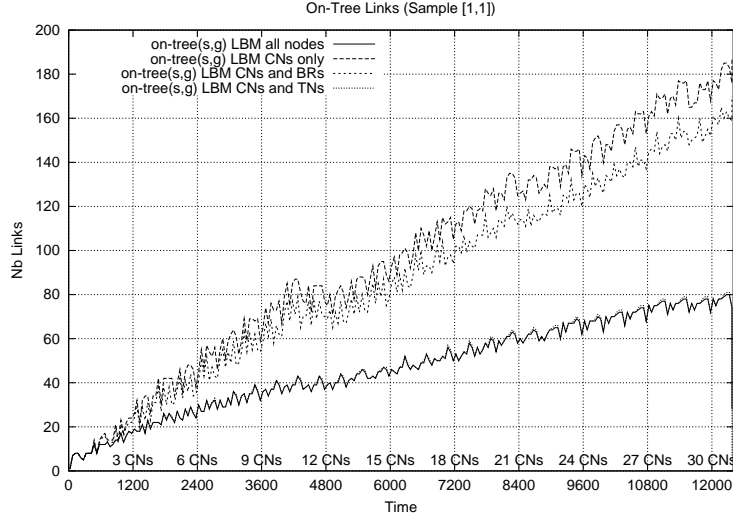


Fig. 3. On-Tree Links

4 Performance Evaluation

In this section, we compare the Xcast delivery of BU against the standard unicast delivery. Simulations were conducted using *NS-2* which has been extended on purpose [11]. A 1050-routers hierarchical topology was generated with *GT-ITM*. This topology comprises 10 backbones and a total of 100 sites. Border routers (BRs in figures) are routers connecting sites to a backbone and transit routers (TNs in figures) are routers in a backbone. CNs are randomly attached to a site router. Given the size of sites, CNs are located 1 to 3 hops away from a border router.

The MN is displaced for 50 seconds in a specific site and is communicating with a number of x randomly selected CNs varying from 1 to 32. The MN doesn't move between subnetworks in a given site. In order to obtain uniform results, independently from the location of the MN and the selected CNs, the simulation is performed 8 times for each set of CNs, positioning the MN in a different randomly selected site. We show all the 8 positions on figures to emphasize that the location of the MN and CNs has no significant impact on the results. This explains why curves are always irregular between any two numbers of CNs. This indeed corresponds to a scenario where the MN is performing wide-area mobility (i.e. topologically distant displacements between sites or access networks), i.e. changing site every 50 seconds and increasing the number of CNs every 8×50 seconds.

The same seed numbers are always retained for selecting CNs and generating displacements of the MN. The performance of Xcast is evaluated under four situations, when all routers and CNs are Xcast-enabled; when only CNs are

Xcast-enabled; when both CNs and BRs are Xcast-enabled; and when both CNs TNs are Xcast-enabled. The *remove flag* is always set.

More details about the NS-2 extensions and the simulation results can be found in [10]. In fig. 3 to 7, Xcast is referred to as *LBM (List Based Multicast)* whereas Mobile IPv6 is referred to as *unicast*.

4.1 Which routers should be XCAST-enabled




Fig. 4. Mobility Cost: BUs to CNs only

As shown on fig.3, the number of *on-tree links* (total number of links consumed by periodic BUs transmitted over the entire network from a position s of the MN to the group g of x CNs during interval of time $t=50$) when all routers are Xcast-enabled is equivalent to the situation where only transit routers and CNs are Xcast-enabled (lower two curves). Similarly, this number is nearly equal when only CNs are Xcast-enabled and when both border routers and CNs are Xcast-enabled (upper two curves). This is due to the uniform distribution of CNs in the topology (the probability that two CNs are in the same site is low). We also see that the slope of the upper curves is higher than the lower curves. The former two increase linearly whereas the latter two look like $\log(\text{nb CNs})$.

The *mobility cost* (total number of bytes consumed by periodic BUs transmitted over the entire network from a position s of the MN to the group g of x CNs during interval of time $t=50$) as shown on fig.4, is not proportional to the number of *on-tree links* since the packet length increases with the number of CNs. Xcast performs badly when the feature is not well deployed in the network since BUs are bounced from one Xcast speaker to another.

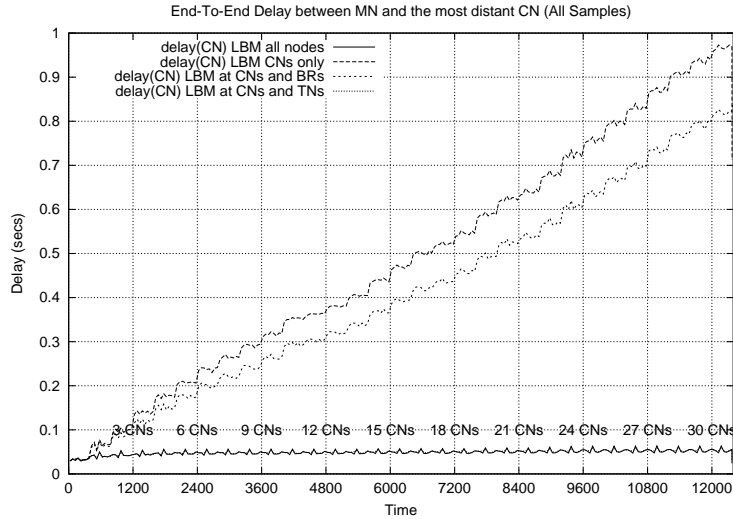


Fig. 5. End-To-End Delay

We note that the slope of all curves is not constant and slows down between 9 and 12 CNs, as for the *on-tree links*, particularly the upper two curves. On the *on-tree links* curves, the slope is higher before 9 CNs, and lower after 12 CNs; on the *mobility cost* curves, the slope is lower before 9 CNs and higher after 12 CNs. This seems to indicate first an optimal value when a BU is most likely to be duplicated, and second an optimal value when the list of CNs is most likely to be split into two equivalent number of CNs. As the BU progresses in the network, the combination of these two events is most unlikely to happen.

The *maximum end-to-end delay* between the MN and the most distant CN is shown on fig.5. This delay has an upper limit when all routers or at least TNs are Xcast-enabled (lower two curves). It obviously tends to increase rapidly to a very large *end-to-end delay* when only CNs are Xcast-enabled (upper curve). In this case, the packet traverses all CNs before reaching its ultimate destination.

4.2 XCAST vs Unicast

When we compare the number of *on-tree links* for the Xcast delivery of BUs against standard Mobile IPv6 (fig.6), we see that Xcast-enabled at both transit routers and CNs (lower curve) is more appropriate than unicast from a number of CNs turning around 5 (upper curve). In terms of *mobility cost*, fig.7, shows that Xcast only enabled at CNs (upper curve) compares very poorly against standard Mobile IPv6 (second curve from top). However, the benefit of Xcast over unicast is obvious when there exists a minimum number of Xcast-enabled routers well located in the network, i.e. at CNs and TNs (third curve from top).

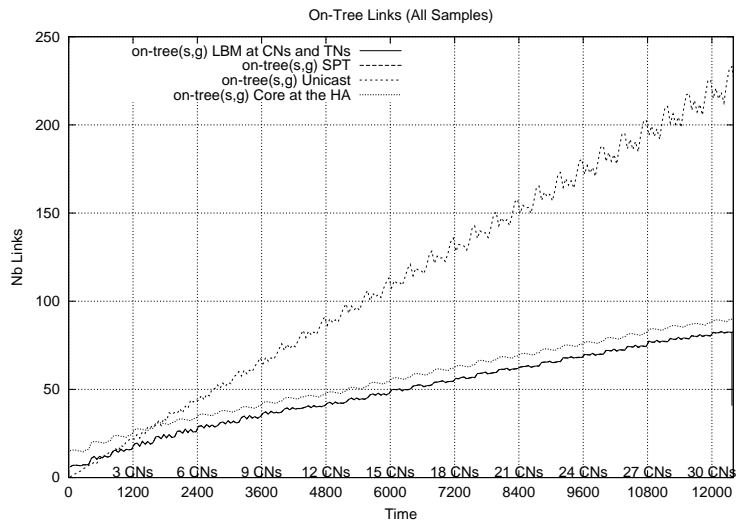


Fig. 6. XCAST vs Unicast: On-Tree Links

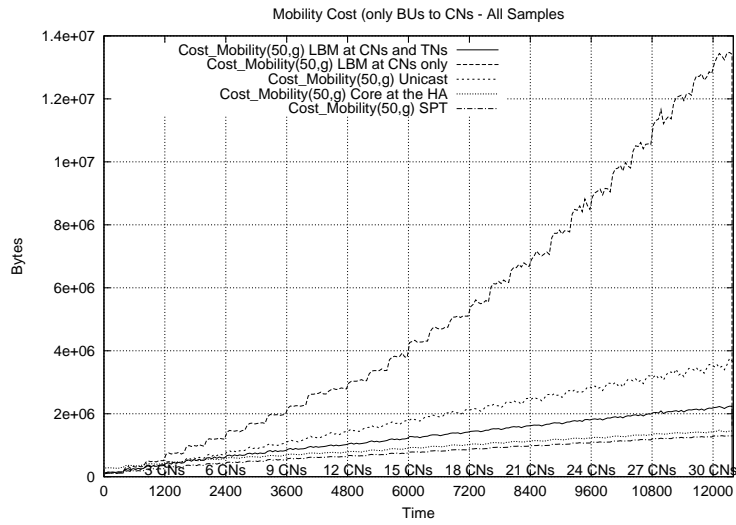


Fig. 7. XCAST vs Unicast: Mobility Cost

5 Conclusion

Routing optimization between correspondent nodes and mobile nodes is a necessary feature to reduce the network load and the transmission delays. However, routing optimization using Mobile IPv6 is made at the expense of periodic BUs that must be sent individually and periodically to every CNs. When these messages are sent to an increasing number of CNs, a periodic signaling burst is propagated in the network. Links close to the node that emits these BUs are most likely to suffer from this, particularly over the air where a significant amount of the available bandwidth is consumed. It is therefore proposed to use Xcast as a means to minimize Mobile IPv6 signaling. The comparison of the unicast delivery of BUs against Xcast shows that Xcast performs rather well provided that a number of routers are Xcast-enabled. On the other hand, Xcast is clearly inefficient when only CNs are able to process the header. We note that the *Xcast Header* requires more processing of the packet at intermediate routers, thus it is not recommended to process the header at each router. The tradeoff is certainly, as highlighted by the performance results, to process this header only at routers well located in the network, probably at transit points like in the backbone. In order to ensure the delivery of BUs to all CNs, all CNs listed in the *Xcast Header* must also be Xcast-enabled. Of course, Xcast is inevitably restricted to a limited number of CNs since the more CNs in the packet, the larger the packet length. To overcome this, CNs may be split into separate groups. As a side note, we observe that our proposal faces a number of security issues that need to be considered in future work. Also, a possible area of application could be network mobility (NEMO).

References

1. Bhagwat, P., Perkins, C., Tripathi, S.: Network Layer Mobility: an Architecture and Survey. *IEEE Personal Communications* **3** (1996) 54–64
2. Ioannidis, J., Duchamp, D., Maguire Jr., G.Q.: IP-based Protocols for Mobile Internetworking. In: Proc. ACM SIGCOMM, Department of Computer Science, Columbia University (1991) 233–45
3. Johnson, D.B., Perkins, C., Arkko, J.: Mobility Support in IPv6. Request For Comments 3775, IETF (2004)
4. Deering, S.: Multicast Routing in Datagram Internetwork. PhD thesis, Stanford University, US (1991)
5. Deering, S., Cheriton, D.: Multicast Routing in Datagram Internetworks and Extended LANs. *ACM Transactions on Computer Systems* (1990) 85–111
6. Boivie, R., Feldman, N., Metz, C.: Small Group Multicast: A New Solution for Multicasting on the Internet. *IEEE Internet Computing* **4** (2000) 75–79
7. Braun, T.: Multicast for Small Conferences. Technical Report IAM-00-008, University of Berne, Switzerland (2000) <http://www.iam.unibe.ch/rvs/publications>.
8. Boivie, R., Feldman, N., Imai, Y., Livens, W., Ooms, D., Paridaens, O.: Explicit Multicast Xcast Basic Specification. Internet Draft draft-ooms-xcast-basic-spec-02.txt (2001) Work in progress.

9. Imai, Y.: XCAST Related Researches, Web Page (As of October 2004)
<http://wiki.xcast.jp/cgi-bin/xcast-wiki.pl>.
10. Ernst, T.: Network Mobility Support in IPv6. PhD thesis, Université Joseph Fourier (2001) <http://www.inria.fr/rrrt/tu-0714.html>.
11. Ernst, T.: MobiWan: NS-2 extensions to study mobility in Wide-Area IPv6 Networks. Technical report, MOTOROLA Labs and INRIA Rhône-Alpes (2001) <http://www.inrialpes.fr/planete/pub/mobiwan/>.