MULTIMOB Group
Internet Draft
Intended status: Informational
Expires: April 22, 2013

S. Figueiredo
Universidade de Aveiro
S. Jeon
Instituto de Telecomunicacoes
R. L. Aguiar
Universidade de Aveiro

October 22, 2012

IP Multicast Use Cases and Analysis over Distributed Mobility Management

draft-sfigueiredo-multimob-use-case-dmm-03.txt

Status of this Memo

Copyright Notice

Abstract

Mobile networks are changing towards distributed mobility management (DMM), tackling inefficiencies of existing mobility protocols regarding network management and packet routing. Identifying IP multicast use cases applicable to DMM is a logical step before exploring solution spaces. This document describes use cases where IP multicast is applied in DMM environments, considering two main deployment options: multicast router or MLD-Proxy deployment at a Mobility Access Router (MAR). Due to the lack of standard terminology, we refer to MAR as the entity embedding mobility-related functions, e.g. providing network access and flow anchoring capabilities. Each deployment option is thoroughly analyzed regarding its advantages and disadvantages, and both multicast listener and source mobility support are considered.

Table of Contents

1. Introduction

Centralized mobility management approach brings several drawbacks such as non-optimal routing or severe overloading on the anchor point. Such problems are expected to be more severe as mobile devices data consumption (and generation) increases.

In order to tackle these problems, the concept of distributed mobility management (DMM) has emerged. It couples per flow and distributed anchoring, bringing the mobility anchor closer to the MN.

IP multicast, as one of the enablers for efficient distribution of multimedia content, is composed of two main functions: multicast routing and membership subscription. When those are coupled with IP mobility, it is very critical to know possible use cases and respective issues, since those techniques were mainly designed for fixed networks.

This document presents possible use cases of IP multicast in a DMM environment, by aligning to DMM Requirements [DMMREQ]. The different use cases result from the different functionalities provided at the MAR - MLD Proxy or MR -, and both mobile listener and sender support are analyzed. The goal was to identify the advantages (e.g. ease of deployment, resource-lightness) and constrains (e.g. lack of resource efficiency, non-optimal routing) of each deployment option, considering its impact for the support of mobile sender or mobile listeners.

2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

This document uses the terminology defined in [RFC5213], [RFC6275], and [RFC3810], and [RFC4601]. Also, new entities are defined relying on the PMIPv6 entities specified in [RFC5213]:

- Mobility Access Router (MAR): A router with the capability of acting both as a mobility anchor and as an access router, in a per flow basis. It can act as either a P-MAR, a N-MAR, or both.

- Previous Mobility Access Router (P-MAR): The MAR where the MN was attached to previously to the IP mobility event, and which is acting as an anchor for one or multiple flows.

- New Mobility Access Router (N-MAR): The MAR to which the MN is currently attached. It provides the network access and thus delivers all the flows destined to the MN's HNPs - including those anchored to previously visited P-MARs.

- Multicast Listener Discovery Proxy (MLD-P): An entity providing MLD based forwarding following the operation defined in [RFC4605]. In the current document, only MLDv2-based signaling is considered, targeting IPv6 networks (REQ3 from [DMMREQ]).

3. Use Cases Description

We identify different use cases that result from the application of existing standards for IP multicast support over mobile environments. We first consider mobile multicast listeners and then mobile multicast senders, and for each case we analyze the impact of deploying distinct functionality at the MAR: either MLD Proxy or Multicast Router (namely PIM-SM capability).

3.1. Assumptions

A1: This draft refers to the requirements in [DMMREQ] as the base DMM framework.

A2: Network access and data anchor functionalities are based in [RFC5213], and are assumed to be provided by a Mobility Access Router (MAR).

A3: It is assumed that when the IP mobility happens, at least one multicast flow is being received (listener) / sent (sender). and a mobility tunnel will be created between the P-MAR and the N-MAR.

A4: Using MLD Proxy, and when a user without mobility session starts a multicast session at a MAR, the upstream interface of MLD Proxy is configured towards an upstream multicast router in the multicast infrastructure. But when a user moves to another MAR (N-MAR), under assumption A3, the upstream interface of MLD Proxy will be configured towards the anchor that enables unicast IP address continuity to be kept (P-MAR, analogous to LMA function).

A5: Mobility occurs within a single PIM-SM multicast routing domain.

3.2. Mobile Multicast listener

3.2.1. MLD Proxy deployment at MAR

3.2.1.1. Operation

In this use case, MLD Proxy is considered as being deployed at all MARs, and operating in an analogous way to the Base Multicast support solution for PMIPv6 [RFC4605]. As such, after mobility, the upstream interface MUST be configured towards the mobility tunnel endpoint, i.e., P-MAR.

When a MN initially attaches to the P-MAR (as shown in Figure 1), it receives a HNP address which will be associated with communications started at that MAR. As the P-MAR detects the new logical link, it transmits a General MLD Query message to which the MN will not yet reply, as it is not yet running any multicast session. The P-MAR then adds the downstream logical link to the MLD Proxy instance [RFC4605]. In this case, i.e. when users subscribe to multicast content only after associating with the MAR, the MLD Proxy will set its uplink to the multicast infrastructure. When the MN intends to start receiving the multicast session, it will send an unsolicited MLD Report, triggered by its application. On receiving the latter message, the MLD Proxy tries to join the multicast channel(s) by sending an aggregated MLD Report through the MLD Proxy upstream interface. Note that the same MLD Proxy instance will be assigned to all MNs which initiated their multicast subscriptions in the current MAR (i.e. the MNs having no multicast mobility session). When the joining procedure ends, multicast data is transmitted through the same interfaces, until reaching the MN.

```
                  +----------------+
                  |   Multicast    |
                  | Infrastructure |
                  +----------------+
                          *
                          * (S,G)
                          *
                     +----------+              +----------+
                     |  P-MAR   |--------------|  N-MAR   |
                     |          |**************|          |
                     | (MLD-P)  |--------------| (MLD-P)  |
                     +----------+              +----------+
                          *                           *
                          *                           *
                       +------+                  +------+
                       |  MN  |      ----->      |  MN  |
                       +------+                  +------+
```

Figure 1 Multicasting architecture using distributed mobility management

When the MN moves from P-MAR, the N-MAR is required to establish a tunnel for IP session continuity of the flows being sent towards and from the MN's HNP assigned by the P-MAR. This implies that N-MAR has an appropriate method to know the P-MAR. Multiple ideas are supposed to be made at the solution stage of DMM WG, therefore it is out of scope of this document. Following the operation of the MLD Proxy [RFC4605], after the bidirectional tunnel establishment, the following process takes place. First, the N-MAR sends a General MLD Query, and verifies whether the MN is admissible for multicast sessions. Then, the MLD Proxy at the N-MAR adds the downstream interface corresponding to the MN, and configures the upstream interface towards the MN's P-MAR, i.e. the tunnel endpoint.

3.2.1.2. Detailed analysis

This scheme is simple and directly applicable to a network-based multicast DMM approach, as no extensions for multicast-related operation are required. Besides, the MLDv2 Proxy was devised with a less resource-demanding nature compared to MRs, aimed for scenarios where multicast routing is not beneficial or required. The same applies in DMM scenarios. However, this approach leads to a couple of relevant issues. Traffic duplication is the result of the tunnel convergence problem, occurring e.g. in [RFC6224]. As shown in Figure 2, MN1 and MN3, which moved from MAR1 and MAR3, respectively, are currently located at the MAR2. Through their respective tunnels, they

receive multicast packets of the same channel through different anchoring MARs. This causes duplicated traffic to converge to the MAR2. Concretely, the magnitude of replicated traffic is expected to be much higher than that of PMIPv6, considering the consequences of having a single mobility entity coupling anchoring and network access provision: we assume that the number of MARs in future DMM domains will be much larger than that of LMAs at core level within a PMIPv6 domain.

```
       +----------------+
       | Multicast Tree | *
       +----------------+     *
               *        *         *
               *          *           *
               *            *             *
         (S,G) *       (S,G)  *               *  (S,G)
               *                *                 *
          +----------+  (-->)  +----------+ (<--)   +----------+
          |   MAR1   |---------|   MAR2   |---------|   MAR3   |
          |          |*********|          |*********|          |
          | (MLD-P)  |---------| (MLD-P)  |---------|  (MLD-P) |
          +----------+  Tun.1  +----------+  Tun.2  +----------+
                                 *   *   *
                                *    *    *
                               *     *     *
           +---+  move      +---+  +---+ +---+   move   +---+
           |MN1|  --->      |MN1|  |MN2| |MN3|   <---   |MN3|
           +---+            +---+  +---+ +---+          +---+

             (<--/-->) : direction of the multicast packet flow
```

Figure 2 Data replication

Another issue is non-optimal routing (Figure 3). If we consider a significantly large domain, multicast packets may traverse a long distance, depending on the setup direction of the upstream interface of MLD Proxy instances. The issue is more noticeable if we assume all MARs are connected to the multicast infrastructure.

When MLD Proxy is used in mobile multicast, low performance handover will result from the need of going through the following process: 1) MLD Proxy sets up the new MLD interface, 2) MLD Proxy sends the General MLD Query, 3) MN sends the MLD Report. Only then MN will receive the content, which for a significant number of IP multicast-based applications, will represent a noticeable service disruption.

```
            +-----------------+
            |    Multicast    |
            | Infrastructure  |
            +-----------------+
                     *
                     * (S,G)
                     *
                +----------+                 +----------+
                |  P-MAR   |------     ------|  N-MAR   |
                |          |****** ... ******|          |
                |(MLD-P)   |------     ------| (MLD-P)  |
                +----------+                 +----------+
                     *                             *
                     *                             *
                  +------+                      +------+
                  |  MN  |       ----->         |  MN  |
                  +------+                      +------+
```

Figure 3 Non-optimal routing problem

3.2.2. Multicast Router deployment at MAR

3.2.2.1. Operation

In this use case, PIM-SM is deployed at all MARs. After the MN attaches to a MAR, its PIM instance will act as the Designated Router (DR) for the MN (and all other attached nodes). The procedure for multicast subscription will be the same as in a fixed network, i.e. the MN sends the Explicit MLD Report, and the MR at the MAR will process that message and join any multicast channel if needed.

When the MN moves, the mobility tunnel will be setup between the two MARs. As soon as the MN sends a MLD Report to the N-MAR, N-MAR will join the multicast group / channel (if needed), the traffic will start flowing from the tunnel and a new downstream state will be configured.

3.2.2.2. Detailed analysis

An advantage of using MRs within MARs, is that multicast routing is not affected due to RPF check. which leads to the selection of a single upstream interface per MAR. This selection is independent of

the number of existing mobility tunnels. On the other hand, the usage of MLD Proxy might lead to the tunneling of the same multicast group to a common MAR, and might mean severe replication, similarly to [RFC6224].

As referred, although the mobility tunnel is activated after the MN mobility, N-MAR will only subscribe the required multicast group / channel after receiving the explicit MLD Report This can translate into several lost packets, as the MN isn't aware of the mobility process, and the General MLD Queries sent by MRs are periodic with a varying frequency in the order of several seconds. Summarizing, smooth handover cannot be assured without extending existing mechanisms. When a MN moves to its N-MAR, there is the possibility that the multicast channel(s) is (are) already being distributed there. In such case, the new MAR will directly forward the multicast traffic to the MN without using the tunnel for subscription function. However, different MARs might be receiving common channels at distinct times, which from the point of view from the receiver, will result in frames replay (if receiving the same frames again) or miss (in case the new MAR transmission is more advanced in time).

A possible non-efficiency is the unnecessary tunnel creation. Following assumption A3, the tunnel is created due to the existence of at least an ongoing multicast flow. The tunnel is created independently of N-MAR's awareness to multicast subscriptions, because the MLD Query occurs after its creation. Although, if the subscription(s) of interest is (are) already being subscribed at the N-MAR's MR, the tunnel will not be used at all for the multicast subscription transmission. Thus, large scale creation of unnecessary tunnels represents non-negligible wasted processing overhead.

3.3. Multicast sender support

3.3.1. MLD Proxy deployment at MAR

3.3.1.1. Operation

Sender mobility support is known to lead to service disruption problems impacting all multicast tree, in particular if SPT is active. In [SENDER], the utilization of MLD Proxy is proposed, being the upstream interface always configured towards the fixed anchoring entity - the LMA. To allow the sender to transmit multicast content to the multicast tree in a DMM framework, the MLD Proxy should configure its upstream interface towards a multicast router [PM-HOME]. Depending on the network topology, it may also be configured towards a MLD Proxy placed on a neighbor MAR. On the multicast source's mobility (Figure 4), an identical operation to the listener mobility

case is expected from the MLD Proxy behavior. In this case, the source uploads multicast traffic through one of MLD Proxy's downstream interfaces, and the traffic is forwarded through the uplink interface towards the P-MAR.

```
                      +------+         +----------------+
                      |  RP  |---------|   Multicast    |
                      +------+         | Infrastructure |
                          *            +----------------+
                          * (S,G)               |
                          *                     |
                     +----------+          +----------+
                     |  P-MAR   |----------|  N-MAR   |
                     |          |**********|          |
                     | (MLD-P)  |----------| (MLD-P)  |
                     +----------+          +----------+
                          *                       *
                          *                       *
                       +------+              +------+
                       |   S  |   ---->      |  S   |
                       +------+              +------+
```

Figure 4 Multicast sender mobility

3.3.1.2. Detailed analysis

The utilization of MLD Proxy carries the previously referred advantages, as ease of deployment and operation lightness.

When a source moves to N-MAR from P-MAR, multicast data will be sent through the mobility tunnel between N-MAR and P-MAR (Figure 5). If a listener (L1) attaches to the same MAR (N-MAR), it will receive the multicast data through multicast infrastructure, following the configuration of MLD Proxy. Hence, the multicast data is routed non-optimally between the source and the listener, going from N-MAR to P-MAR, to the multicast routing tree, and then back to N-MAR again before reaching the listener.

```
               +------+         +----------------+
               |  RP  |*********|   Multicast    |
               +------+         | Infrastructure |
                   *            +----------------+
                   * (S,G)             *
                   *                   *
              +----------+       +----------+
              |   P-MAR  |-------|   N-MAR  |
              |          |*******|          |
              | (MLD-P)  |-------| (MLD-P)  |
              +----------+       +----------+
                   *                *     *
                   *                *     *
                +------+  move +------+ +-----+
                |  S   |  ---> |   S  | |  L1 |
                +------+       +------+ +-----+
```

Figure 5 Triangular routing after source mobility

A similar problem occurs in the opposite process, i.e. if a multicast source starts transmitting multicast content at a MAR, and a listener moves to the same MAR while receiving the source's content (Figure 6).

```
               +------+         +----------------+
               |  RP  |*********|   Multicast    |
               +------+         | Infrastructure |
                   *            +----------------+
                   * (S,G)             *
                   *                   *
              +----------+       +----------+
              |   N-MAR  |-------|   P-MAR  |
              |          |*******|          |
              | (MLD-P)  |-------| (MLD-P)  |
              +----------+       +----------+
                *       *              *
                *       *              *
            +------+  +----+  move  +----+
            |  S   |  | L1 |  <---  | L1 |
            +------+  +----+        +----+
```

Figure 6 Triangular routing after listener mobility (MLD-Proxy case)

When the source and the listener are within a same MAR (MAR2), if both the source and listener try to send the session and receive it, respectively, the traffic will be optimally sent to the listener without going through native multicast infrastructure. As the traffic reaches the MLD Proxy via the downstream interface to which the source is attached, it will be sent through the downstream interface to which the listener sent the MLD Report. However, if the source and the listener move to different MARs, the traffic will traverse the following non-optimal path, even though they share a common anchor:

Source -> MAR1 -> MAR2 -> Multicast Tree -> MAR2 -> MAR3

This problem is depicted in Figure 7.

```
                             +----------------+
                             | Multicast Tree |
                             +----------------+
                                    *  *
                                    *  *
                                    *  *
                                    *  *
                                    *  *
           +----------+  (-->)  +----------+  (-->)  +----------+
           |   MAR1   |---------|   MAR2   |---------|   MAR3   |
           |          |*********|          |*********|          |
           | (MLD-P)  |---------| (MLD-P)  |---------|  (MLD-P) |
           +----------+  Tun.1  +----------+  Tun.2  +----------+
               *                  *       *              *
               *                 *         *             *
               *                *           *            *
             +---+    move    +---+       +---+  move  +---+
             | S |    <---    | S |       | L |  -->   | L |
             +---+            +---+       +---+        +---+

                (<--/-->) : direction of the multicast packet flow
```

Figure 7 Multicast traffic non-optimal routing due to both mobile sender and listener

REQ1 from [DMMREQ] refers that "DMM MUST enable a distributed deployment of mobility management of IP sessions so that the traffic can be routed in an optimal manner without traversing centrally deployed mobility anchors". When a MN subscribes to a new multicast session with existing multicast mobility session, the Aggregated MLD Report containing all the MN's multicast subscriptions will be sent from the current MLD Proxy through the same uplink interface, i.e. towards a single multicast mobility anchor. This results in some of previously identified issues (e.g. non-optimality in the path that both the subscription and multicast traffic traverse). It can be stated that the MLD Proxy nature doesn't comply with the aforementioned requirement, leading to the subscription of any multicast flow using the same multicast mobility data path.

This problem is depicted in Figure 8, where both multicast flow 1 and flow 2 reach MAR2 from MAR1, being flow 2's optimal routing path affected by the mobility status of the MN, and in particular by the order in which the multicast flows were subscribed. While this issue is not exclusively related to mobile multicast sources, it is better depicted and its' impact in the routing is more obvious when considering one.

```
                          +----------------+
                          | Multicast Tree |
                        * +----------------+#
                      *                      #
                    *                          #
                  *                              #
                *                                  #
        +----------+  (-->)                 +----------+
        |   MAR1   |---------            -------|   MAR2   |
        |          |#*#*#*#*#............#*#*#*#|          |
        | (MLD-P)  |---------            -------|  (MLD-P) |
        +----------+  Tunnel (used after mob.)  +----------+
            *                                      *    #
            *                        downstream) #    # (upstream)
            *                                      *    #
            *                                      #    #
          +---+              move              +---+  +---+
          | L |          ---------->           | L |  | S |
          +---+                                +---+  +---+

          * : Multicast flow 1
          # : Multicast flow 2 (subscribed after some time in MAR2)
```

Figure 8 Non-optimal routing due to single MLD Proxy uplink

3.3.2. Multicast Router deployment at MAR

3.3.2.1. Operation

When a source starts transmitting multicast traffic, the content will be encapsulated in PIM-Register messages, and sent towards the RP (e.g .statically configured or discovered through a Bootstrapping Router (BSR)). In DMM, the RP can be a core MR or a MAR (including the initial MAR). The RP's SPT and each of the DR's SPTs may then be created. When the source moves, N-MAR's MR will create the state for the new multicast group, and the traffic will be forwarded using the tunnel to the P-MAR, until reaching the RP (unless it is the PIM instance at the P-MAR itself). It is then sent down the RPT. Again, the creation of the SPTs will typically be triggered following PIM-SM regular operation.

3.3.2.2. Detailed analysis

In case the RP's Source Path Tree is built before the mobility process, it will be destroyed due to mobility, and the tree construction process will be reinitiated at the new MAR. Also, in case the Shortest Path Tree between the listener's DRs and the sender's DR is being used, mobility will reset the PIM process to the RPT stage. This means that each sender mobility event results in increased signaling overhead and delay as consequence of the multicast routing convergence (i.e. Phase 2 and Phase 3 from PIM-SM operation). Moreover, non-optimal routing occurs when the RPT is used. When a sender moves to a MAR where listeners are subscribing to the channel it is sending, the multicast traffic will always reach the N-MAR by going through the RP (as shown in Figure 9).

Using PIM-SM in DMM scenarios there is a trade-off between the routing non-optimality of RPT and the non-efficient consequences of frequent SPT establishment. It is important to note that this impact is magnified the more listener's DRs are receiving the multicast channel(s).

```
              +------+         +----------------+
              |  RP  |*********|   Multicast    |
              +------+         | Infrastructure |
                  *            +----------------+
                  * (S,G)             *
                  *                   *
             +----------+       +----------+
             |   P-MAR  |-------|   N-MAR  |
             |          |*******|          |
             |   (MR)   |-------|   (MR)   |
             +----------+       +----------+
                  *                *     *
                  *                *     *
               +------+  move +------+ +-----+
               |  S   |  ---> |   S  | |  L1 |
               +------+       +------+ +-----+
```

Figure 9 Triangular routing after source mobility (MR case)

3.4. Summary of analysis

Table I summarizes the previous analysis, globally depicting the differences between MLD Proxy and MR over DMM. The relevant sections for each feature are appended below, and the meaning of each feature is extended afterwards.

Feature \ Function	MLD Proxy	Multicast Router
Lightweight	Yes	No
Optimal routing	No (3.2.1.2)/ (3.3.1.2)	Yes (using SPT) (3.3.2.2)
Efficient distribution	No (3.2.1.2)/ (3.3.1.2)	Yes (3.2.2.2),(3.3.2.2)
Distributed anchoring	No (3.3.1.2)	Yes (3.2.2.2)
Seamless mobility (listener)	No (3.2.1.2)	No (3.2.2.2)
HO signaling overhead	Low (3.2.1.2), (3.3.1.2)	High (when using SPT) (3.3.2.2)

Tunnel is always used	Doesn't apply	No (3.2.2.2)

Table I. Comparison between MLD Proxy and MR deployment

The meaning of each of the entries is as follows:

Lightweight: this entry reflects whether the deployed multicast feature has a resources-wise lightweight operation.

Optimal routing: This entry reflects whether optimal routing is assured.

Efficient distribution: This entry reflects vulnerability to multicast traffic replication.

Distributed anchoring: This entry assesses whether for a single MN, different multicast streams can be anchored at different mobility anchors or not.

Seamless mobility (listener): This entry reflects whether IP mobility is seamless from the point of view of the mobile listener's application.

HO signaling overhead: This entry assesses the amount of signaling introduced by the IP mobility of a MN represents. This signaling may be relative to the mobility protocol or the multicast routing operation.

Tunnel is always used: This entry assesses whether the mobility tunnel utilization is assured or not.

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

TBD

6. References

6.1. Normative References

[RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC6275] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

[RFC3810] R. Vida, and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," IETF RFC 3810, June 2004.

[RFC5213] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6", IETF RFC 5213, August 2008.

[RFC4605] B. Fenner, H. He, B. Haberman, and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) Based Multicast Forwarding ("IGMP/MLD Proxying")", IETF RFC 4605, August 2006.

[RFC4601] B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

6.2. Informative References

[RFC6224] T. Schmidt, M. Waehlisch, S. Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domain", RFC 6224, April 2011.

[DMMREQ] H. Chan, "Requirements of distributed mobility management", draft-ietf-dmm-requirements-02 (work in progress), September 2012.

[SENDER] T C. Schmidt et al, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains", draft-ietf-multimob-pmipv6-source-01 (work in progress), July 2012.

[PM-HOME] S. Jeon, N. Kang, and Y. Kim, "Mobility Management based on Proxy Mobile IPv6 for Multicasting Services in Home Networks," IEEE Transactions on Consumer Electronics (TCE), vol. 55, no. 3, pp. 1227-1232, August 2009.

Authors' Addresses

Sergio Figueiredo
Universidade de Aveiro
3810-193 Aveiro, Portugal

E-mail: sfigueiredo@av.it.pt

Seil Jeon
Instituto de Telecomunicacoes
Campus Universitario de Santiago
3810-193 Aveiro, Portugal

E-mail: seiljeon@av.it.pt

Rui L. Aguiar
Universidade de Aveiro
3810-193 Aveiro, Portugal

E-mail: ruilaa@ua.pt