Multicast requirements for control over LLN
draft-vanderstok-roll-mcreq-00

Abstract

   This is a working document intended to focus discussion on
   requirements for multicast in Low-power and Lossy Networks in the
   area of M2M communication for control purposes. The Trickle
   algorithm, which uses re-broadcasting to assure that messages arrive
   at all destinations, is proposed as the ROLL multicast protocol. In
   this draft additional requirements on Trickle, such as timeliness and
   ordering, are motivated by building control. Re-broadcasting and
   timeliness can be mutually exclusive properties. To alleviate that
   problem, this draft considers the possibility of reducing
   interference by limiting the number of transmission paths between any
   two devices in the mesh.

Status of this Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on August 27, 2012.

Copyright Notice

Table of Contents

1. Introduction

   The ROLL working group is chartered to design and standardize a
   routing protocol for resource constrained devices in Low-power and
   Lossy Networks (LLN) [I-D.ietf-roll-rpl]. The requirements for ROLL
   are documented in [RFC5548] [RFC5673] [RFC5826] [RFC5867]. For
   building control it is recognized that most communication is local to
   the wireless mesh network, and does not necessarily pass through the
   edge router. The point-to-point RPL routing algorithm is developed
   to efficiently support such applications [I-D.ietf-roll-p2p-rpl].
   The Trickle multicast was developed to support the RPL routing
   algorithm, and later proposed to support general multicast delivery
   in LLN. [RFC6206].

   This draft discusses the multicast requirements for constrained
   devices participating in M2M building control networks. An important
   requirement is the delivery of control commands to a subset (group)
   of neighbouring devices in the LLN within some latency bound.

1.1. 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].
   Addtional privileged words are described below.

   A "device" is a physical processor connected to at least one link
   through a network interface. Each interface has at least one IP
   unicast address. The IP address is optionally bound to a host name,
   which may be a Fully Qualified Domain Name (FQDN).

   One device communicates directly with another device by wirelessly
   transmitting packets to it over a link. The link quality is divided
   in three regions:

   1. good: where a transmitted packet will be correctly received by a
       destination with a probability higher than 99%.
   2. transitional: where the probability of correct reception
       fluctuates.
   3. bad: where almost no transmission is successfully received.

   It is empirically known that good links can become bad occasionally
   due to dynamic effects such as multipath interference.

   A distinction is made between reception and delivery of a message. A
   message is received when it is stored in the reception buffer of the
   receiver after transmission and all error checks have been
   succesfully passed. The message is delivered when the message is

passed from the reception buffer to the client application. We also
say the client application accepts the message.

Broadcasting is used for the link-local sending of one packet to all
reachable 1-hop neigbours. This is equivalent to the term link-local
multicast.

1.2. Motivation

In this draft, we focus and develop discussions on requirements
pertaining to multicasting, in the context of building control
applications.

2. Application characteristics

Multicast is important for building control applications. Two types
of applications are considered:

1. Discovery messages to all members of the mesh (multicast GET)
2. Control messages to a subset of the mesh (multicast PUT)

The first type requires the message to be sent to a subset which may
be randomly distributed over the building area. Some of the
destinations return unicast messages to the source.

The second type requires the message to be sent to a closely spaced
subset. No return messages are generated. This second type is the
subject of this draft, although most of the requirements equally
apply to case 1.

An office building typically consist of multiple floors, divided in
working areas. The working areas can be open or enclosed by walls.
Within a working area sensors measure temperature, presence, humidity
and other parameters. On the basis of these measurements, equipment
within the working area can receive commands to change settings. A
well-known example is presence detection to switch on or dim lights.
The equipment configuration is quite stable, because devices are
installed in the ceiling, and modifying (or servicing) the
installation can be costly.

The equipment is interconnected in a wireless network. The RF
transmissions pass through the walls and generate interference to the
wireless equipment in other working areas.

The lay-out of a network may be different from installation to
installation. However, it is expected that many wireless networks
extend over one floor and include many working areas. Another

working hypothesis is that most of the time sensors will multicast
their values to a group of devices within the working area.
Consequently, multicast messages are often meant for a subset of
neigbouring devices.

A LoWPAN is a mesh of wireless devices that share the same IPv6
address prefix. A typical LowPAN in a building may cover the area of
an entire floor. A commercial installation may cover 1000 m2 per
floor. A length of 50 m can easily mean a hop count >5 for a message
to pass from end to end. For example, devices may be installed in
the ceiling in a grid with a grid pattern distance of 40 cm between
devices.

Messages may consist of sensor measurements performed or commands
issued in a given working area, which then must be acted upon by
neigbouring devices in the same working area. Given that source and
sink are located in one working area, sink and source of a multicast
message are often between 3 - 6 m from each other. Consequently, it
is required to send a multicast to a subset of the devices in the
LoWPAN.

In case of commands to luminaries, messages must be delivered within
a clear deadline of about 200ms. In [RFC5867] a deadline of 120 ms
is suggested for other building applications.

Although most control messages are exchanged between closely spaced
devices, it is sometimes necessary, say every hour or less
frequently, to send a message to a subset of devices covering the
whole building. In that case the multicast message will need to pass
the edge router of the lowpan and to propagate to other subnets.

3. Multicast requirements

The Multicast requirements are derived from the characteristics of
the applications. A device is said to be correct it it follows the
multicast algorithm. The application characteristics and the network
installation make it possible to add an additional set of network
properties to make the multicast algorithm more efficient.

The basic traditional multicast requirements (PUT and GET)are:

o Validity: If sender S sends message, m, to a group, g, of
   destinations, a path exists between S and a destination D, and S
   and D are correct, D eventually accepts m.
o Integrity: A destination accepts m at most once from sender and
   only if sender sent m to a group including destination.

   o Agreement: If a correct destination of g accepts m, then all
      correct destinations of g accept m.

   The set of intended destination devices is identified by the
   multicast (group) IP address. Every device in the associated
   multicast group is a destination of the multicast. Each destination
   accepts messages with as destination the specified IP multicast
   address. Additional multicast requirements are:

   o Timeliness: There is a known constant C such that if m is sent at
      time t, no correct destination accepts m after t+C.

   For lighting control applications the value of C is taken as 200 ms.
   This requirement considers the PUT case and not the return of a
   response in the GET case.

   o Ordering: When m1 and m2 sent to the same group g, and a receiver
      in g accepts message m1 before m2, every receiver in g accepts m1
      before accepting m2

   Ordering applies to PUT and GET case. Ordering can be partial or
   total. Partial ordering means that for specified message pairs one
   of the pair precedes the other. In case of total ordering, every
   message pair is ordered. Partial ordering is obtained by adding
   message counters in the message such that destinations can order the
   messages of a given sender. Messages from different sources are not
   ordered. Total ordering can be obtained with vector clocks or using
   synchronized clocks. Vector clocks require a large overhead that
   increases linearly with the number of devices in the network. As
   long as no synchronized clocks are available, partial ordering seems
   the most realistic. Total Ordering is interesting for the discovery
   application. When two devices announce themselves simultaneously
   with conflicting properties, all participants can come to the same
   decision by favoring the first arrival. Partial ordering is
   necessary when a multicast message needs multiple packets (for
   example discovery messages) or when multicast messages are sent with
   intervals shorter than the throughput delay.

4. Wireless link characteristics

   It is possible to broadcast from a source to a set of devices
   reachable over good links in one hop. This is not sufficient because
   the set of reachable devices is often a subset of the set of
   destination devices. Consequently, additional measures are needed to
   make sure that the Agreement requirement is met. A standard
   technique, to reach all devices instead of a subset, stipulates that
   every receiver of a broadcast message rebroadcasts this message

(flooding). When the multicast address corresponds with a specified
multicast address in the receiver device, the message is delivered.
Thanks to this technique it is assured that when a path exists
between the source and the destination device, the destination device
will eventually receive the message from the sender.

Given the network density described above, the multicast can generate
a broadcast storm with lots of interfering senders. The technique,
also used in Trickle, is to randomly delay the message rebroadcast.
However, the long delays can seriously jeopardize the timeliness
requirement. This draft proposes three ways suggested by the
application characteristics, to reduce the interference between re-
broadcasting devices:

1. Restrict the scope of the multicast.
2. Restrict number of rebroadcasting devices.
3. Weaken the Timeliness requirement.

In the application characteristics it is mentioned that most control
messages have a set of destinations which are closely spaced to the
source. The interference between multicast sources can be reduced by
limiting the scope of the broadcast message. The ensuing proximity
condition can be formulated for PUT and GET as:

o Proximity condition: A multicast message is accepted by a subset
   of devices closely spaced to the sender.

In practice, this condition means that most multicast messages can be
constrained to 1-2 hops. It is recommended to put the multicast
range under control of the multicast source.

Given the stability of the network configuration, the configuration
of good links is also stable over long periods (say several days).
When all good links are available, the number of possible paths
between a source and each of its destinations is probably larger than
required given the sporadic failure of a good link. Under the
assumption that the qualities of the good links of a given device are
unrelated, the failure of good link has no consequence for
alternative good links. Given the installation characteristics
above, the number of paths between source and destinations is much
larger than required to assure that a source device remains connected
to all its destinations when a good link fails. The number of paths
can be reduced by specifying a subset of devices, called relay
devices, to rebroadcast messages. A path can pass from a source via
relay devices to the multicast destinations. A relay device can also
be a destination device. In [RFC5867] it is mentioned that 1 out of
2 devices is a relay device. Given the network densities foreseen
for lighting, a much lower relay density is possible. The reduction

of the relay devices reduces the risk of interference in the dense
networks described above. An appropriate condition to assure the
presence of a path between source and destination can be formulated
as:

o Multiple relay links: any device has good links to at least q
   relay devices

The value of q is determined by the quality of the links in a given
installation.

However, the probability that a path was temporarily unavailable
cannot be excluded. The timeliness requirement is too strong for
wireless sensor networks, where packets get lost for multiple reasons
like hidden terminal, multipath fading, and others. The timeliness
requirement can be reformulated for the PUT case as:

o Majority Timeliness: There is a known constant C, and a subset s
   of correct destinations in group g, such that if m is sent to g at
   time t, with low probability p, all destinations in s accept m
   after t+C.

The agreement requirement specifies that the destinations in s accept
the message eventually. Both probability p and set s are specified
as function of the installation and linked with the value of q.

Using rebroadcast with a low frequency (as in Trickle) assures that
missed messages are eventually repeated. For a lighting application
this means that in general all lights switch on/off within 200 ms and
quite infrequently, (say once a month) one out of all lights
swictches on/off a bit later (say a few seconds).

5. Recommendation

From the text above emerges a number of recommendations to make it
possible to put propagation characteristics of the multicast
algorithm under application control.
1. Take into account timeliness and partial ordering requirements in
    multicast algorithm.
2. Exploit the small range of most multicasts and put multicast
    range under application control.
3. Introduce a subset of devices as relay devices to reduce the
    number of rebroadcasting devices.
4. Use majority timeliness requirement to balance the number of
    relay devices with respect to the probability that a device
    misses its multicast reception deadline.

5. Multicast messages transit through an edge router.

6. IANA Considerations

   This document makes no request of IANA.

   Note to RFC Editor: this section may be removed on publication as an
   RFC.

7. Security Considerations

   TBD

8. Acknowledgments

   This I-D has benefited from conversations with and comments from
   Anders Brandt, Kerry Lynn, Zach Shelby, Emmanuel Frimout, Michael
   Verschoor, Jamie Mc Cormack, Dee Denteneer, Jerald Martocci, Matthieu
   Vial, and Nicolas Riou.

9. References

9.1. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119, March 1997.

   [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel,
              "Routing Requirements for Urban Low-Power and Lossy
              Networks", RFC 5548, May 2009.

   [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney,
              "Industrial Routing Requirements in Low-Power and Lossy
              Networks", RFC 5673, October 2009.

   [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation
              Routing Requirements in Low-Power and Lossy Networks",
              RFC 5826, April 2010.

   [RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen,
              "Building Automation Routing Requirements in Low-Power and
              Lossy Networks", RFC 5867, June 2010.

   [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko,

"The Trickle Algorithm", RFC 6206, March 2011.

9.2. Informative References

   [I-D.ietf-roll-rpl]
            Brandt, A., Vasseur, J., Hui, J., Pister, K., Thubert, P.,
            Levis, P., Struik, R., Kelsey, R., Clausen, T., and T.
            Winter, "RPL: IPv6 Routing Protocol for Low power and
            Lossy Networks", draft-ietf-roll-rpl-19 (work in
            progress), March 2011.

   [I-D.ietf-roll-p2p-rpl]
            Goyal, M., Baccelli, E., Philipp, M., Brandt, A., and J.
            Martocci, "Reactive Discovery of Point-to-Point Routes in
            Low Power and Lossy Networks", draft-ietf-roll-p2p-rpl-07
            (work in progress), January 2012.

Author's Address

   Peter van der Stok (editor)
   Philips Research
   High Tech Campus 34-1
   Eindhoven,   5656 AA
   The Netherlands

   Email: peter.van.der.stok@philips.com