Network Working Group
Internet-Draft
Intended status: Experimental
Expires: August 18, 2016

J. Chroboczek
IRIF, University of Paris-Diderot
February 15, 2016

Diversity Routing for the Babel Routing Protocol
draft-chroboczek-babel-diversity-routing-01

Abstract

This document defines an extension to the Babel routing protocol that allows routing updates to carry radio frequency information, and therefore makes it possible to use radio diversity information for route selection.

Table of Contents

1. Introduction and background 2
2. Operation of the protocol 3
  2.1. Changes to data structures 3
  2.2. Receiving updates 4
  2.3. Sending updates 4
  2.4. Metric computation and route selection 5
  2.5. Protocol encoding 5
3. IANA Considerations 6
4. References 6
Appendix A. The Z3 algorithm 7
Author's Address 7

1. Introduction and background

The Babel routing protocol [RFC6126] does not mandate a specific algorithm for computing metrics; Appendix A of that document suggests using an additive integer metric. While this works well in many topologies, it fails to take into account the possibility of interference between radio links, which is important in multi-frequency wireless mesh networks.

Consider for example the following topology, where the solid lines use one radio frequency and the dashed lines another, and suppose that the solid frequency has very slightly lower packet loss than the dashed one:

```
   B
  / \
 /   \
A     D
 \   .
  \ .
   C
```

When sending data from A to D, Babel will reliably choose the solid route through B. Howerver, this route self-interferes: when B is sending a packet to D, it cannot simultaneously be receiving a packet from A, which halves the effective throughput. No such issue arises with the route through C, which should therefore be preferred.

Interference needs to be taken into account even when it happens between non-adjacent links. Consider the following topology:

```
   B +++ C
  /       \
 /         \
A           F
 \         .
  \       .
   D +++ E
```

When routing data from A to F, the route through B and C has two interfering links: if two packets are sent by A and C at roughly the same time, a collision will occur, and both packets will need to be resent. Again, no such issue arises with the route through D and E.

2. Operation of the protocol

The diversity protocol extension allows a Babel router to attach information about radio frequency to the routes that it maintains -- we call this the route's "diversity information".

We assume that all links can be categorised into one of the following categories:

- non-interfering links, e.g. wired links;

- links that have a well defined frequency, and only interfere with other links at the same frequency; these are described by a single channel number, an integer between 1 and 254;

- interfering links, links that interfere with all other links except non-interfering links.

This model does not describe reality accurately, since distinct but close radio frequencies do in fact interfere, but it works well enough in practical networks, where a small number of discrete radio frequencies are used.

2.1. Changes to data structures

A Babel router maintains a route table ([RFC6126] Section 3.2.5). A router implementing diversity routing has one additional field in every route table entry:

- the diversity data, a (possibly zero-length) sequence of channel numbers, each of which is an integer between 0 and 255.

The diversity data is interpreted as the set of channels of the links that would be followed by a packet sent along this route, omitting non-interfering links. The values 0 and 255 are special: they indicate, respectively, a non-interfering link (a link that doesn't interfere with any other links) and an interfering link (a link that is assumed to interfere with all other links except non-interfering ones).

2.2. Receiving updates

When a node receives an Update TLV, it creates or updates a routing table entry according to [RFC6126], Section 3.5.4. A node that performs diversity routing extends the procedure given in that section with the following actions.

Let D be the diversity information attached to the received Update TLV, or the one-element sequence 255 if there is no such information. Then the diversity information in the routing table entry is set to D', where:

o if the interface over which the update was received is non-interfering, then either D'=0.D or D'=D (the choice is left to the implementation);

o if the interface over which the update was received is interfering, then D'=255.D;

o if the interface over which the update was received is tuned to channel C, then D'=C.D.

Note that zero-length diversity information is different from lack of diversity information: the latter is treated as 255 (interfering, since no information is available) in order to ensure reasonable behaviour when interoperating with the original Babel protocol.

Note further that there are two ways of encoding a non-interfering link: the interference information can be omitted (D'=D) or a 0 can be prepended to the interference information. The latter, less parsimonious encoding MAY be preferred by implementations that wish to ignore diversity information after a given number of hops. Both encodings MUST be correctly parsed by a receiving node.

2.3. Sending updates

A Babel node sends updates in various circumstances, described in [RFC6126], Section 3.7. A node performing diversity routing attaches diversity data to every update that it send. This diversity data is computed as follows:

- o if the update is for a locally redistributed route, then the value is implementation-dependent (zero-length diversity information is a good choice in most cases);

- o if the update is for a route in the Babel route table, then the diversity information is taken from the route table.

2.4. Metric computation and route selection

How the diversity data is used for metric computation and/or route selection is left to the implementation, as long as it obeys the rules given in Sections 3.5.2 and 3.6 of [RFC6126]. In particular, the strict monotonicity requirement implies that a non-interfering hop must be taken into account in the resulting metric -- it cannot be simply ignored.

An algorithm that has been found to work well in practice is given in Appendix A.

2.5. Protocol encoding

We define one new sub-TLV which is attached to Update TLVs and contains a sequence of channel numbers.

2.5.1. Encoding of channel numbers

A channel number is encoded as a one-octet integer. The values MUST be interpreted as follows:

- o 0: used to represent a non-interfering link;

- o 1-254: radio channel numbers, link-technology specific

- o 255: used to represent an interfering link.

In order to ensure consistent metrics computation, implementations supporting IEEE 802.11 SHOULD use values 1 through 14 and 46 through 165 for encoding IEEE 802.11b and IEEE 802.11a channel numbers respectively, and implementations using other link technologies SHOULD choose values that do not collide with IEEE 802.11a/b channel numbers. However, since choosing inconsistent values does not prevent interoperability but merely leads to suboptimal routing, this is not mandated by this specification.

2.5.2. The Diversity sub-TLV

Diversity data is carried in a Diversity sub-TLV [RFC7557] that is carried by Update TLVs. The sub-TLV contains a sequence of octets that directly encode the diversity data from the route table.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Type = 2   |    Length     |   Channel 1   |   Channel 2   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Channel 3   |  ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

Fields :

Type   Set to 2 to indicate a Diversity Information sub-TLV.

Length   The length of the body, exclusive of the Type and Length fields.

Channel n An integer between 0 and 255, as described in the previous section.

3. IANA Considerations

IANA is instructed to add the following entry to the "Babel Sub-TLV Types" registry:

Type	Name	Reference
2	Diversity	(this document)

4. References

[RFC6126] Chroboczek, J., "The Babel Routing Protocol", RFC 6126, April 2011, <http://www.rfc-editor.org/info/rfc6126>.

[RFC7557] Chroboczek, J., "Extension Mechanism for the Babel Routing Protocol", RFC 7557, May 2015, <http://www.rfc-editor.org/info/rfc7557>.

Appendix A. The Z3 algorithm

In this section, we describe the Z3 algorithm, a particular instance of diversity routing that has seen some modest deployment and that appears to work reasonably well in practice while being extremely easy to implement.

The Z3 algorithm works by announcing a slightly smaller metric than the metric it uses for route selection when announcing over a non-interfering link. In effect, a Z3 router maintains two metrics for each route: the noninterfering metric, which is announced on links that can be proven to not interfere with the route being announced, and the interfering metric, which is used for route selection and announced over all other links.

More precisely, upon receiving an update with metric M over a link with cost C, the interfering metric is set to C+M, as suggested in Appendix A of [RFC6126]. The non-interfering metric is set to alpha*C+M, where 0<alpha<1 is called the diversity factor (with rounding biased upwards in order to ensure strict monotonicity).

Let D be the diversity data of route R, and L be a link. We say that R interferes with L when one of the following is true:

- L is a non-interfering link (e.g. an Ethernet); or

- L is a radio interface tuned to channel C, and neither C nor 255 is an element of D.

When we announce R over L, we announce the interfering metric if R interferes with L, and the non-interfering metric otherwise.

The metric that Z3 yields is non-isotonic; hence, Z3 Babel does not necessarily converge to a set of minimum-metric routes. In fact, the set of minimum-metric routes might not even be a tree in the general case. We believe that Z3 Babel converges to a Nash equilibrium, but this appears to be a difficult property to prove.

Author's Address

Juliusz Chroboczek
IRIF, University of Paris-Diderot
Case 7014
75205 Paris Cedex 13
France

Email: jch@pps.univ-paris-diderot.fr