

RTGWG
Internet-Draft
Intended status: Standards Track
Expires: April 14, 2013

P. Thubert, Ed.
cisco
IJ. Wijnands
Cisco Systems
October 11, 2012

Applying Available Routing Constructs to bicasting
draft-thubert-rtgwg-arc-bicast-00

Abstract

This draft introduces methods that leverage the concept of ARC to enable bicasting operations.

Requirements Language

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

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 April 14, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Downward Bicasting Operation	5
4. Upward Bicasting Operations	6
4.1. Resolving crossing ARCs	6
4.2. Single Point of Failure	7
5. Applicability	9
5.1. In conjunction with Protocol Independent Multicast	9
6. Manageability	9
7. IANA Considerations	9
8. Security Considerations	9
9. Acknowledgements	9
10. References	9
10.1. Normative References	9
10.2. Informative References	9
Authors' Addresses	10

2. Terminology

The draft uses the terminology defined in [I-D.thubert-rtgwg-arc].

This specificatin also introduces Sided ARCs, that is ARCs with at least an Edge that is known as Left and an Edge that is known as Right. The sense of Left and Right adds up to the existing sense of height that is already defined in [I-D.thubert-rtgwg-arc].

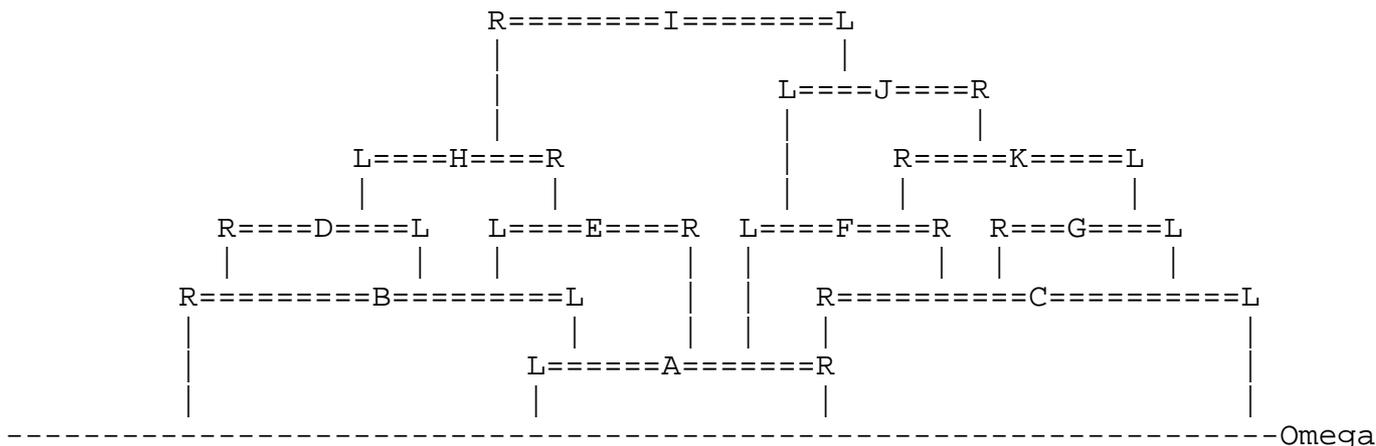


Figure 2: Orienting ARCs

One way of doing this is

- o The basic rule is that an ARC MUST have at least one Left and one Right Edge.
- o The leg of an ARC between the cursor and the Edge inherits the side of the Edge. In a Comb, the whole buttressing ARC inherits the side of the Edge.
- o An Edge ending in Omega can arbitrarily become Left or Right as long as the basic rule is satisfied.
- o An Edge that does not end in Omega inherits the side of an ARC it terminates into, again as long as the basic rule is satisfied.
- o A collision occurs if all the Edges end up on the same side. The shortest path is used to resolve the collision and restore the basic rule: the Edge closer to Omega and all buttressing ARCs on the same side of the cursor keep the side, and the other Edges are toggled. In case of equal cost, an other tie breaker must be used.

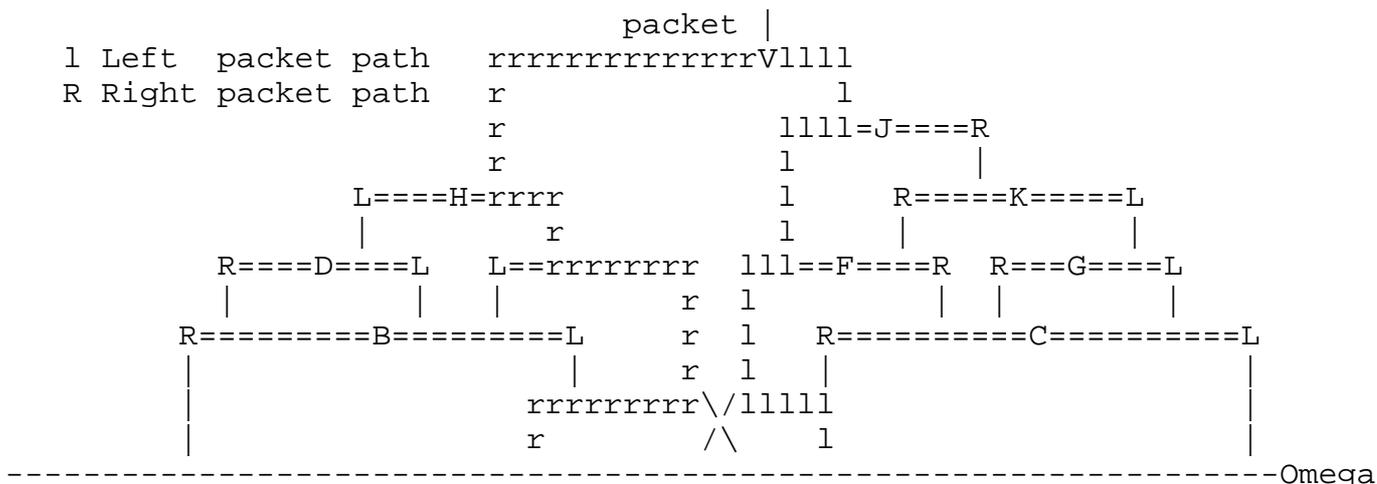


Figure 4: Breakage at a Congruent Link

4. Upward Bicasting Operations

It is also possible with a downward bicasting to place states in the intermediate routers in order to provision an upward bicast path from Omega to a source D. In that case, if the graph is biconnected, it is possible to resolve the pathological cases so as to ensure a real separation of the left and Right paths.

4.1. Resolving crossing ARCs

The first pathological case occurs when both Left and Right packet cross over the same ARC, as illustrated below. Say that the Right reservation comes first and sets up the right path:

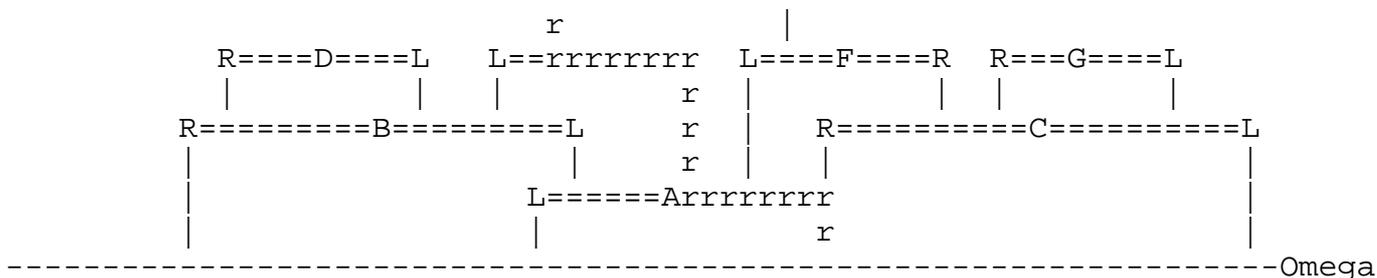


Figure 5: crossing: Right packet

Then comes the left reservation which collisions:

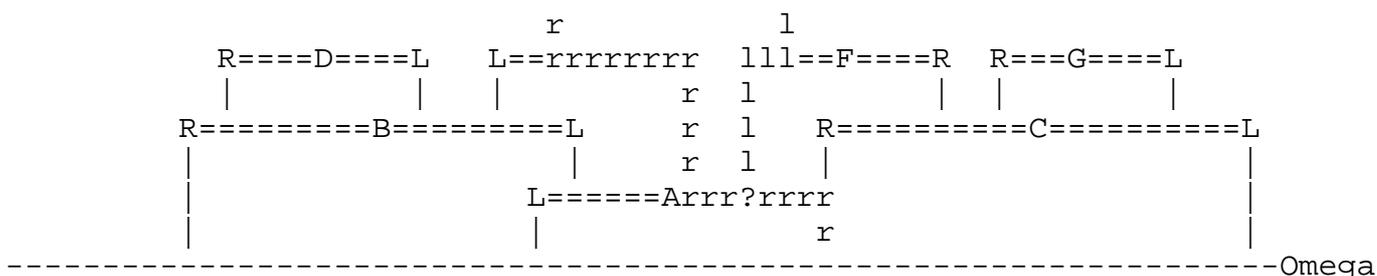


Figure 6: crossing: left packet approaching

The segment between the two incoming point of the common ARC is common to both path and expose the bicasted traffic. The resolution is to leave the second packet through but prune the unwanted states along the collision segment of the ARC afterwards.

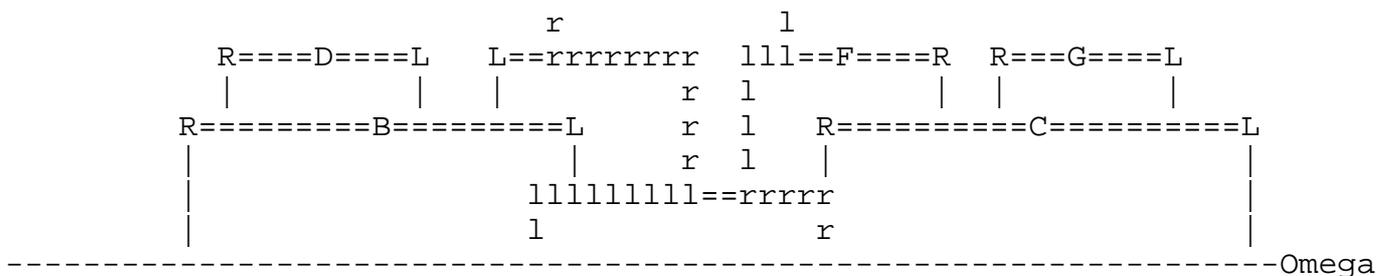


Figure 7: crossing: Resolved state

States along the ARC between the two incoming points are cleaned, up and the paths that were generated by the Left and Right packets are now crossed. This results in two non-congruent upward paths.

4.2. Single Point of Failure

The second pathological case occurs when both Left and Right packet reach a same ARC at the same node, which is this a Single Point Of Failure (SPoF) for both paths.

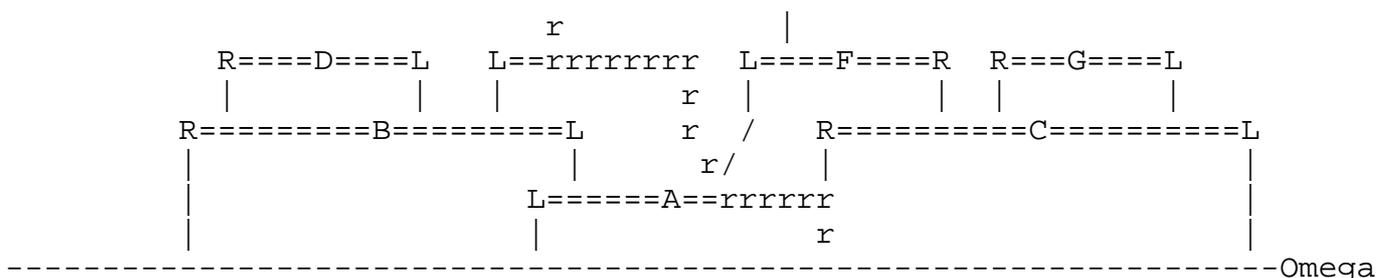


Figure 8: SPoF: Right packet

The resolution is to reject the second packet and send it back along the incoming ARC to exit on the other side. The rejected packet cleans up the states that it has created on its way back and then creates states on the other side of the ARC.

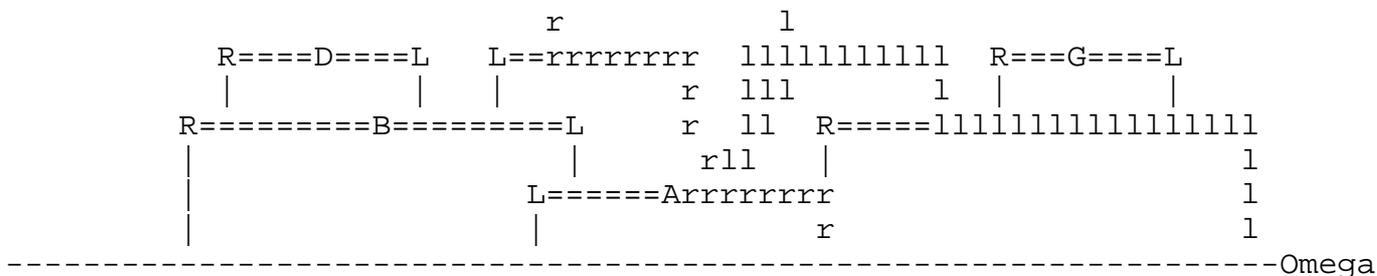


Figure 9: SPoF: Left Packet

At this point the downward packet will exit the incoming ARC in the wrong side for its own indication.

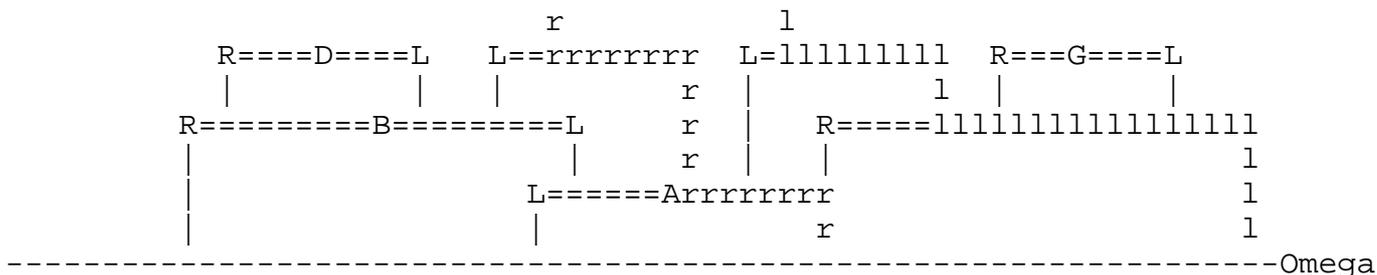


Figure 10: SPoF: Resolved state

This is in fact what happens also in the case of a monoconnected zone, or if a breakage cause the downward packet to bounce.

5. Applicability

5.1. In conjunction with Protocol Independent Multicast

(To be refined in 01) Upwards bicasting can be used for Protocol Independent Multicast PIM [RFC4601] and Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths mLDP [RFC6388]. A bicasted downwards Join message would establish two non congruent return paths, each path joining the receiver and Omega that is the set of existing receivers.

6. Manageability

This specification describes a generic model. Protocols and management will come later

7. IANA Considerations

This specification does not require IANA action.

8. Security Considerations

This specification is not found to introduce new security threat.

9. Acknowledgements

The authors wishes to thank Dirk Anteunis, Stewart Bryant, Patrice Bellagamba, George Swallow, Eric Osborne, Clarence Filsfils and Eric Levy-Abegnoli for their participation and continuous support to the work presented here.

10. References

10.1. Normative References

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

10.2. Informative References

[I-D.thubert-rtgwg-arc]
Thubert, P. and P. Bellagamba, "Available Routing Constructs", draft-thubert-rtgwg-arc-00 (work in

progress), October 2012.

- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, November 2011.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

IJsbrand Wijnands
Cisco Systems
De kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com