Broadcast Communications in Cartesian Networks

Larry Hughes*
Routing Research Group
Department of Electrical and Computer Engineering
Dalhousie University
Halifax, Nova Scotia, B3J 2X4, Canada

June 27, 2003

Abstract

A Cartesian network is a modified, regular 2D-mesh network that supports unicast (point-to-point) communications without the use of routing tables. In this paper, a protocol for the transmission of broadcast packets in a Cartesian network is proposed. The protocol is loop-free and ensures that broadcast packets visit each Cartesian router at most once, by dynamically creating a spanning tree. The protocol requires no additional state information other than that already maintained by the existing Cartesian routers.

This paper describes the Cartesian broadcast protocol, the state to be carried by the broadcast packet, and the different routing algorithms.

Keywords: Broadcast, Cartesian routing, 2D-mesh networks.

1 Introduction

The ability to broadcast information to a group of entities or networks is central to many routing protocols and distributed applications. How broadcast is achieved depends upon a number of factors, including the underlying network's topology, router design, and address structures.

In a broadcast transmission, the broadcast packet must visit all stations in the network. How broadcast is realized depends upon the underlying network architecture. Inherently broadcast networks, such as the bus and ring, ensure that all packets, regardless of the destination address, are made available to each station. However, in most other network topologies, broadcast becomes problematic since multiple paths introduce the possibility of cycles and flooding.

*This research is funded by a Discovery grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada.

Employing time-to-live or hop counts can help minimize the impact of flooding; however, these techniques do not eliminate it. Furthermore, hop counts and other flooding avoidance techniques can require storage or additional processing, thereby increasing the complexity of the router design. Protocols such as Reverse-Path Broadcast (RPB) have eliminated many of these problems by requiring routers to maintain and exchange state that allows them to determine if they are on the shortest path back to the source [3].

In a Cartesian network, routers maintain limited state information (typically only their own address), which is sufficient for unicast communications [5]. This lack of state makes broadcast problematic, although modification to the basic router algorithm does allow broadcast to be achieved using flooding and hop counts. Implementing protocols such as Reverse-Path Broadcast in a Cartesian network requires the router to maintain additional state by packet exchange with its neighbours. This paper proposes a broadcast protocol that is in keeping with the limited state objectives associated with the original Cartesian network unicast protocol.

2 Cartesian Networks

A Cartesian network consists of collectors and arterials [5]. A collector is a 'horizontal' network (running east-west), connecting collector routers, while an arterial is a 'vertical' network (running north-south) that intersects collectors. Collector routers have an east and a west port, connecting the router to its collector, whereas arterial routers have sufficient ports to allow connection to its collector and its arterials. To ensure connectivity, an arterial is not permitted to bypass a collector.

Each router is given an address that indicates its location in Cartesian space (represented, for example, using latitude and longitude); within a collector, all collector routers share a common latitude. A sample Cartesian network is shown in Figure 1.

2.1 Collector Router Unicast Algorithm

A collector router has three ports: east, west, and local. The east and west ports connect to other routers, while the local port acts as a gateway to a machine or network of machines. Packets for unicast (or point-to-point) transmission are supplied to the collector router via the local port:

- Packets for destinations on the same collector are queued on the collector port (i.e., east or west) leading to the destination.
- Packets for destinations on different collectors must be routed to an arterial router. Each local port maintains an Arterial Director Indicator (or ADI) that gives the local port a 'hint' as to the direction of an arterial for unicast transmissions to destinations on different collectors. For reliability, collector routers maintain two ADIs, indicating the state of arterials to the east and west (see Figure 2).

When a packet arrives at a collector router (from either its east or west port), the router inspects the packet's destination latitude; if it is different from that of the router, the packet is forwarded out the port opposite from which it was received. However, if the destination latitude is the same as that of the router, the router inspects the destination longitude; the actions taken by the router depend upon incoming port, the packet's destination longitude, and the station's longitude (see Table 1).

Table 1: Collection Router Actions for Unicast Packets
(Collector latitude equals packet destination latitude)

Incoming port	Address Comparison	Action
West	Dest.Lon == Stn.Lon	Keep packet
	Dest.Lon < Stn.Lon	Discard packet
	Dest.Lon > Stn.Lon	Forward packet east
East	Dest.Lon == Stn.Lon	Keep packet
	Dest.Lon < Stn.Lon	Forward packet east
	Dest.Lon > Stn.Lon	Discard packet

2.2 Arterial Router Algorithm

An arterial router has at least three ports (east, west, and local); with the exception of the most northerly or southerly routers, it is also expected to have ports leading to the north and south.

If a unicast packet's destination address indicates a destination on the same collector, the local port places the packet on the east or west port, depending upon the destination's longitude. However, if the destination is on a different collector, the packet is placed on the north or south queue (north is assumed to be 'greater than' south).

When a packet arrives at an arterial router from a collector, the actions taken by the router depend upon the destination address: if the packet's destination latitude is the same as that of the router, the router acts like a collector router (see Table 1). However, if the destination latitude is greater that the station's latitude, the packet is forwarded north, otherwise it is forwarded south.

A similar algorithm is applied to packets arriving from arterials: packets from the north must have a destination latitude equal or less than that of the station, otherwise the packet is discarded. Similarly, packets from the south must have a destination latitude equal or greater than that of the station.

2.3 Virtual Arterials

To reduce path lengths and offer increased fault tolerance, an arterial router can be associated with multiple arterials, as shown in Figure 1. In situations where an arterial router has lost all its northbound (or southbound) arterials, it can use a 'virtual arterial' to traverse a collector to an arterial router that continues the northbound (or southbound) path.

In order that an arterial router 'know' the state of the arterials to its east and west, each arterial router maintains a Reachability Indicator (or RI). The RI indicates the arterial path(s) available by taking the eastbound or westbound virtual arterial. A virtual arterial is shown in Figure 3.

2.4 State

The Cartesian unicast routing algorithm eliminates the need for routing tables. This has two significant benefits: first, routing decisions are reduced from between O(log(n)) and O(n) time to O(1); and second, routing maintenance is greatly simplified (restricted to updating ADIs and RIs should the state of an arterial router change).

3 Design of the Broadcast Protocol

A Cartesian network with more than one arterial introduces the possibility of cyclic paths, making broadcast by flooding problematic. Overcoming this limitation is typically achieved using hop-counts or by requiring the router to maintain other state information, a solution that runs counter to the fundamental objectives of Cartesian networks.

In order to design of the broadcast protocol, it is necessary to examine how broadcast could be achieved, first within a collector and then between collectors.

3.1 Intra-collector broadcast

In its simplest form, a broadcast transmission on a collector requires the source router to transmit two broadcast packets: one out its East port, the other out its West port. To distinguish between broadcast and unicast, the destination address must indicate that the packet is a broadcast packet. When a broadcast packet is received, each collector router must take a copy of the packet and forward it out the opposite port from which it was received.

3.2 Inter-collector broadcast

Inter-collector transmissions are packet transmissions that are handled by arterial routers across an arterial (that is, the packet is transmitted North or South). When a broadcast packet reaches an arterial router, the router is responsible for forwarding the packet on the collector, as described in the previous section.

In an inter-collector broadcast, cycles or repeated visits must be avoided, meaning that it is necessary to ensure that the packet:

- traverses no more than one arterial in a given direction.

  If a broadcast packet is transmitted across two or more arterials in the same direction, the packet could be forwarded more than once on the same collector. To avoid this from happening, it is necessary to ensure that a broadcast packet traverse at most one arterial.

- does not ‘back-track’.

  When a packet reaches an arterial router from one direction, the router must not forward the packet in the direction from whence it came; to do so could result in duplicate copies of the packet appearing on the network.

The source router must therefore select an arterial and transmit the packet towards that arterial, with an indication that the packet is to be sent along the arterial in question. Furthermore, since a broadcast packet actually travels across two arterials (one northbound and one southbound), it is necessary to indicate to each whether or not the packet is already traversing the network in the ‘other’ direction.

3.3 Broadcast state

The requirements described in the previous section for the broadcast protocol mean that state must be maintained for a successful broadcast transmission. The state is most easily maintained in the packet itself, indicating that the packet is a broadcast packet and whether the packet is to traverse the northbound arterial, the southbound arterial, or both.

A broadcast packet requires indications that it is:

- a broadcast packet. A single bit can be used to distinguish the broadcast and unicast addresses; this is referred to as the ‘Broadcast bit’.
- being forwarded on a northbound arterial. A single bit can be used to indicate whether or not the packet is to be forwarded on a northbound arterial; this is referred to as the ‘Northbound bit’.
- being forwarded on a southbound arterial. A single bit, equivalent to the ‘Northbound bit’, used for the southbound arterial; referred to as the ‘Southbound bit’.

The Northbound and Southbound bits are referred to collectively as the ‘direction bits’.

Routers receiving a packet with the broadcast bit set are expected to take a copy of the packet. Forwarding of the packet depends upon the type of router, its connections, and the state of the direction bits; how the direction bits are manipulated and interpreted is described in subsequent sections.

Cartesian unicast packets carry a fixed-length destination address; several benefits can be achieved by if the broadcast address is the same size:

- Both broadcast and multicast addresses can be supported. For example, by adopting the Ethernet address structure, the broadcast address can be represented by setting the address bits, while multicast addresses can be distinguished from broadcast by clearing at least one of the address bits [1].
- Router design is simplified with a single address structure for a packet's destination address.

3.4 Routing Protocols

Broadcast routing protocols are required for both collector and arterial routers.

3.4.1 Initial packet transmission

Broadcast packets that are sent by the source router have the Broadcast bit set; the values of the direction bits depends upon whether the router is a collector router or an arterial router:

- If the source router is a collector router, it is expected to transmit two packets along its collector. The first packet is sent in the direction of a known arterial (using the ADI) with the direction bits clear.

  The second packet is sent in the opposite direction with the direction bits set. In those situations where the collector router lies between two arterial routers connected both north and south, it can send the first packet in either direction.
- If the source router is an arterial router, the packet is sent in both directions on its collector with the direction bits set. The packet is also transmitted on the northbound arterial and the southbound arterial; the direction bits are set.

3.4.2 Collector router packet forwarding

When a broadcast packet is received by a collector router, the direction bits are neither examined nor processed. The router takes a copy of the packet for its local port, while forwarding the packet out the 'opposite' port from which it was received.

3.4.3 Arterial router packet forwarding

An arterial router can receive a broadcast packet from the collector to which it is connected or from one of its arterials. How the packet is processed depends upon the value of the direction bits and the port from which the packet is received, as shown in Table 2.

Table 2: Arterial Router Actions

Southbound bit	Northbound bit	Incoming Port	Action
Set	Set	Collector (East or West)	Forward on opposite collector port (West or East).
Set	Set	Arterial (North or South)	Forward on both collector ports (East and West). Forward on opposite arterial port (South or North). The router's actions if an opposite arterial is not available are described in the next section.
Clear	Clear	Collector (East or West)	Set direction bits. Forward on opposite collector port (West or East). Forward on two arterials (one North and one South). The router's actions if one or both arterials are missing are described in the next section.
Set	Clear	Collector	See section 3.5.
Clear	Set	Collector	See section 3.5.

As discussed previously, an arterial router may be associated with more than one northbound (or southbound) arterial to minimize path lengths. When a broadcast packet arrives at an arterial router with multiple arterials, the router must transmit the packet on one of the arterials only. If the packet is sent on more than one arterial in the same direction, flooding will occur on all collectors to the north (or south) of the arterial router. (The choice of arterial is dependent upon a variety of factors, most notably the maximization of network resources; this is presently under investigation.)

3.5 Virtual Arterials

There are two problems that must be considered if virtual arterials are to support broadcast communications: how packets are to be transmitted when sent by a source router on the virtual arterial, and how packets are to be forwarded when they arrive at an arterial router that requires the virtual arterial.

3.5.1 Source Router

When acting as a source router, a collector router that lies on a virtual arterial must still send a broadcast packet out each of its collectors. The values of the packets' direction bits are determined by the network's topology as reflected in the router's ADI:

- An ADI that indicates a northbound and southbound arterial in one direction on a virtual arterial implies that there is at least one arterial router in the specified direction. Without an indication that the collector router is on a virtual arterial, the packet is transmitted with both its direction bits clear in the direction specified by the ADI (the other packet is transmitted out the opposite port with both direction bits set).
- However, if the ADI indicates a northbound arterial in one direction and a southbound arterial in the other, the collector router still sends two broadcast packets, with the modifications shown in Table 3.

Table 3: Collector Router Modifications

ADI Direction	Action
Northbound only	Clear the Northbound bit. Set the Southbound bit. Forward in direction of northbound arterial router.
Southbound only	Clear the Southbound bit. Set the Northbound bit. Forward in direction of southbound arterial router.

When an arterial router is a source router on a virtual arterial, it transmits three packets:

- The first is sent on its sole (real) arterial with both direction bits set.
- The second packet is sent across the virtual arterial to the arterial router that has an arterial leading in the direction opposite to that of the transmitting arterial router; this is determined from the router's Reachability Indicator (RI).

  The packet's direction bit associated with the transmitting arterial router's arterial is set, while the other direction bit is clear.
- The third packet is sent on the collector in the direction opposite to that of the virtual arterial with both its direction bits set.

3.5.2 Arterial Router

There are three cases that an arterial router must handle when it is part of a virtual arterial:

1. The arterial router receives a broadcast packet from a collector with one of the direction bits clear and the other direction bit set. The router's actions are described in Table 4.

Table 4: Arterial Router Actions

Southbound bit	Northbound bit	Incoming Port	Action
Set	Clear	Collector (East or West)	Set Northbound bit. Forward packet on opposite collector (West or East). Forward packet on northbound arterial.
Clear	Set	Collector (East or West)	Set Southbound bit. Forward packet on opposite collector (West or East). Forward packet on southbound arterial.

The actions shown in Table 4 assume that the arterial router has a connection in the required direction (i.e., north or south); if no such connection exists, the router acts as a collector router and forwards the packet out the opposite port without modifying the direction bits.

2. The arterial router receives a broadcast packet with both the direction bits being clear from a collector router but the router is connected to a single arterial (either northbound or southbound). The actions required by the router are as follows:

   (a) Forward the packet on the arterial to which the router is connected. Both the direction bits are set.

   (b) Forward the packet out the opposite collector from which it was received (this is assumed to be a virtual arterial). One of the two direction bits is set, depending upon the direction the packet was forwarded (determined in the previous step).

3. The broadcast packet is received from an arterial (both of the direction bits are expected to be set), but the router does not have a corresponding northbound (or southbound) arterial. However, the arterial router is connected to a virtual arterial.

   When this occurs, the arterial router performs the following actions:

   (a) Forward the packet on the non-virtual arterial collector (both of the direction bits remain set).

   (b) Forward the packet on the virtual arterial with the Northbound (or Southbound) bit set, corresponding to the direction from which the broadcast packet was received; the value of the 'other' direction bit is cleared.

4 Algorithmic Efficiency

The efficiency of the Cartesian broadcast routing protocol can be expressed in terms of packet transmissions. In an n node network with optimal broadcast efficiency, exactly $n-1$ packet transmissions will be made, as each node will receive the broadcast packet exactly once. This can be illustrated by considering the path a broadcast packet follows; this is done in two phases:

Collector. Intuitively, when a broadcast packet is introduced into a collector with c routers, exactly $c-1$ packet transmissions will be made.

A packet introduced to an arterial consisting of a arterial routers, will require exactly $a - 1$ forwards. This holds true for both uninterrupted arterials and virtual arterials. Packet transmissions across a virtual arterial are ignored since the collector routers in the virtual arterial are considered part of the collector's broadcast.

Arterial. The collector and arterial cases are easily combined. Since every non-source collector must be reached by a forward on an arterial, we can say that every c router collector requires c forwards for a complete broadcast, effectively removing the arterial from the equation.

This leaves the special case of the source collector: it requires only $c - 1$ forwards (because no arterial had to be traversed to initially deliver the message). Therefore, considering all n routers, there must be exactly $n - 1$ forwards to reach all routers on the network.

5 Router Implementations

The addition of the broadcast routing protocol to Cartesian networks has resulted in modifications to both the collector and arterial routers.

5.1 Collector Routers

The initial transmission requires the router to determine the value of the direction bits by inspecting the value of its ADI. Two copies of the packet are transmitted, to the east and the west.

When a broadcast packet is received by a router, it is simultaneously forwarded out the 'opposite' port from which it was received and copied for any local processing.

5.2 Arterial Routers

The modifications to the arterial routing protocol can be divided into two parts: packets from a collector and packets from an arterial.

5.2.1 Packets from a Collector

When a broadcast packet arrives from a collector, the arterial router attempts to forward the packet out its north and south arterials, if it has them and depending upon the values of the packet's direction bits. The values of the direction bits may also be changed. The packet is always forwarded out the opposite collector.

The following code fragment from our OPNET simulation shows the implementation:

```
if (pktptr -> NB == 0 && N_port_exists)
{
  pktptr -> NB = 1;
  fwd_arterial(NORTH, pktptr);
}
if (pktptr -> SB == 0 && S_port_exists)
{
  pktptr -> SB = 1;
  fwd_arterial(SOUTH, pktptr);
}
fwd_collector(opposite, pktptr);
```

Although the collector can act as a virtual arterial, it is not necessary for the arterial router to make a distinction between packets sent from collector routers or packets forwarded across the collector by arterial routers.

5.2.2 Packets from an Arterial

When a broadcast packet arrives from an arterial, the arterial router is expected to forward the packet on its collectors and an arterial. If an arterial exists that leads in the direction opposite to that of the incoming packet, the router simply forwards the packet in that direction. However, if no such arterial exists, the router must use its reachability indicator to determine the collector (acting as a virtual arterial) over which the packet is to be forwarded. When the packet is forwarded over a virtual arterial, the value of one of the packet's direction bits will be cleared.

The following code fragment illustrates the steps associated with forwarding a broadcast packet that is received from an arterial connected to the north (i.e., the packet must be forwarded to the south):

```
if (inport == NORTH)
{
  if (S_port_exists)
  {
    fwd_collector(WEST, pktptr);
    fwd_collector(EAST, pktptr);
    fwd_arterial(SOUTH, pktptr);
  }
  else
  if (reachability . east == SOUTH)
  {
    fwd_collector(WEST, pktptr);
    packet . SB = 0;
    fwd_collector(EAST, pktptr);
  }
  else
  {
    fwd_collector(EAST, pktptr);
    packet . SB = 0;
    fwd_collector(WEST, pktptr);
  }
}
```

In the above code fragment, if a south port exists, the packet is forwarded to the south as well as both east and west. If there is no southbound port, it is necessary to check for a virtual arterial leading to the south by inspecting the reachability. In both cases, the packet is first sent on one collector with both direction bits set, while on the other collector, the Southbound bit is clear.

A similar code fragment is required for broadcast packets received from the south.

6 Concluding Remarks

The paper has shown how broadcast communications can be supported in a Cartesian network. The broadcast packet must maintain state to indicate that the packet is a broadcast packet, while two additional bits signify whether the packet has been sent north, south, or both. Two routing protocols are needed, one for collector routers and the other for arterial routers; these protocols are separate from the existing unicast routing protocol for Cartesian networks. No changes are required to the unicast protocol, the network, or the packet structure.

Furthermore, other than router modifications to identify broadcast packets, no other changes are required to the network. Source collector routers use the existing unicast Arterial Direction Indicators to start packet transmission. No broadcast-specific control information need be exchanged between routers.

Another benefit of the broadcast protocol is that it also permits multicast (or group) transmission at no additional cost [6]. In a broadcast transmission, the transmitted packet must visit all stations in the network, while in a multicast transmission, the packet must visit all stations that belong to the multicast group. How broadcast and multicast are realized depends upon the underlying network architecture: in a broadcast network such as the Ethernet, all multicast packets are made available to all stations on the network. This is one approach we are considering for the support of multicast communications within a Cartesian network.

When used in conjunction with the existing unicast Cartesian routing protocol, the broadcast protocol allows the development of distributed algorithms that require both unicast and broadcast, such as the Border Gateway Protocol (BGP) [7, 8]. By supporting broadcast, routing information about the underlying Autonomous Systems can be distributed amongst the border routers, which, in turn, can forward packets across a backbone network without the use of routing tables [4]. The simplicity of the two Cartesian routing protocols allows for the development of all-optical networks, avoiding o-e and e-o conversions at each router [2].

The broadcast routing protocol has been simulated successfully using OPNET to implement both collector and arterial routers. We are in the process of developing the necessary hardware to support the protocol in our Cartesian network testbed.

Acknowledgments

The author would like to thank Omid Banyasad, Evan Hughes, and the members of the Cartesian Routing Research Group for their valuable suggestions in the development of this protocol.

References

[1] 3Com. Etherlink III Parallel Tasking ISA, EISA, Micro Channel, and PCMCIA Adapter Drivers Technical Reference, August 1994. Part No. 09-0398-002B.

[2] E. Barefoot. A Proposal for All-Optical Packet Routing. Master's thesis, Faculty of Engineering, Dalhousie University, April 2002.

[3] S. Deering et al. Protocol Independent Multicast Version 2 Dense Mode Specification. http://www.ietf.org/proceedings/99nov/I-D/draft-ietf-pim-v2-dm-03.txt, June 1999. Accessed July 2002.

[4] L. Hughes, A. Adegoke, G. Subramaniam, and H. Song. Cartesian Core Routing. In *Queen's University 21st Biennial Symposium on Communications*, June 2002.

[5] L. Hughes, O. Banyasad, and E. Hughes. Cartesian routing. *Computer Networks*, 34:455–466, 2000.

[6] L. Hughes and H. Song. Multicast Communication in Cartesian Networks. In *First Annual Computer Networks and Services Research Conference*. University of Moncton, May 2003.

[7] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC1771, 1995.

[8] P. Traina. BGP-4 Protocol Analysis. RFC 1774, 1995.

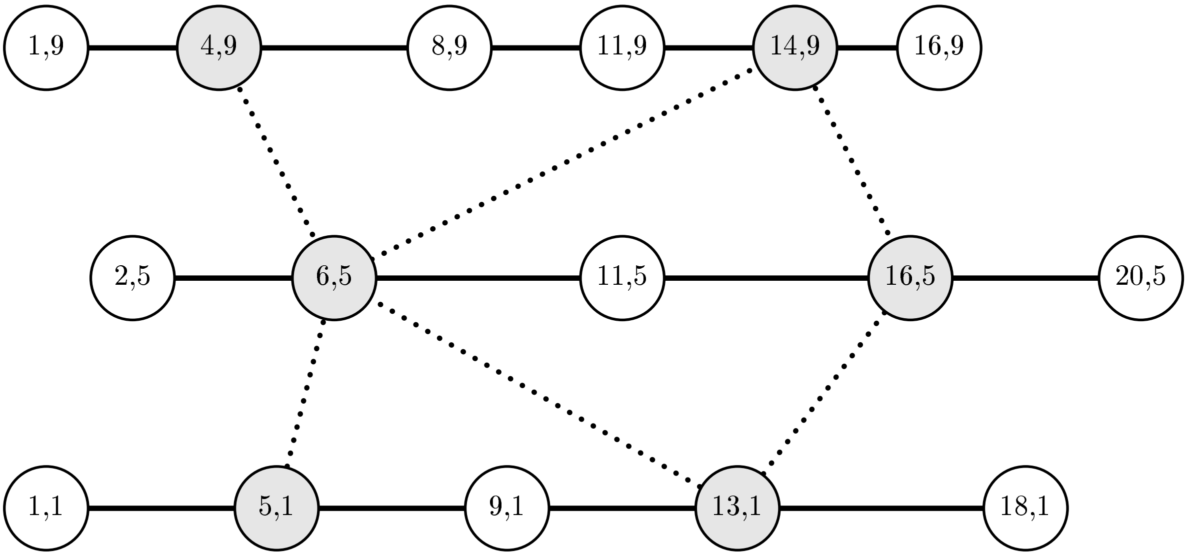

Figure 1: A Cartesian Network (Collector: solid line; Arterial: dotted line; Collector Router: white circle; Arterial Router: light gray circle)

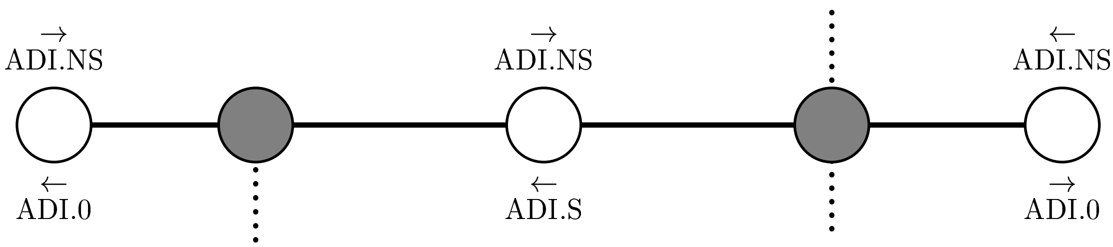

Figure 2: Arterial Direction Indicators

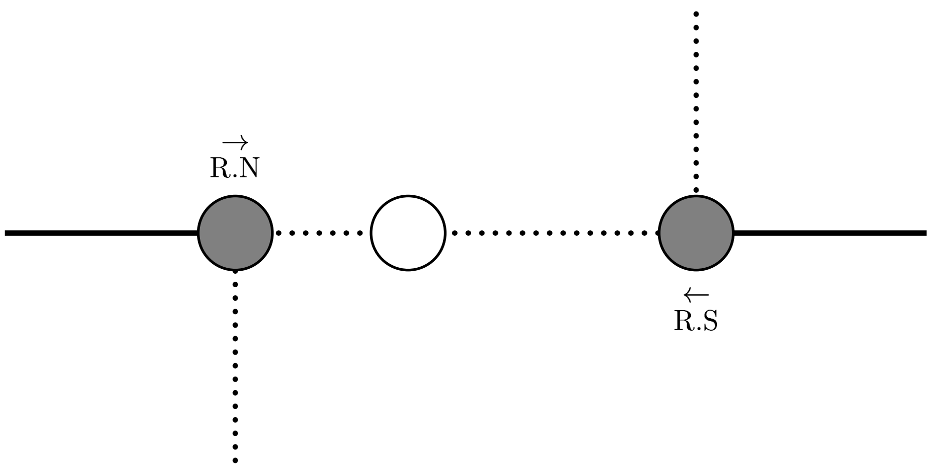

Figure 3: A virtual arterial connecting two arterial routers (Reachability is indicated by 'R' followed by the direction; the arrow points towards the arterial router)