Cartesian Core Routing

Larry Hughes, Adefisayo Adegoke, Ganesh Subramaniam, Haiyu Song
Department of Electrical and Computer Engineering
Dalhousie University, Halifax, Nova Scotia, Canada B3J 2X4

Presented to the 21st Biennial Symposium on Communications
Department of Electrical and Computer Engineering
Queen's University, Kingston, Ontario, Canada K7L 3N6

2-5 June 2002

Abstract

The dominant backbone protocol implemented in the Internet is the Border Gateway Protocol (BGP). Each router implementing BGP maintains a routing table. As networks increase in size, the memory requirements for the routing tables and the time taken to search the tables increase proportionally.

In this paper, the authors discuss Cartesian Core Routing (CCR), an adaptation of Cartesian routing which transparently replaces BGP as the backbone routing protocol. ASs communicate with the CCR network border device using BGP. The CCR border device uses the Cartesian routing algorithm within the CCR network.

Cartesian routing is a novel packet routing technology that differs from existing provider-based routing in that routers maintain a minimal amount of state information. Routing tables are also unnecessary, reducing routing decisions from O(log(n)) and O(n) time using routing tables to O(1).

The CCR algorithm has three cogent advantages over the existing exterior gateway protocols: faster routing and transmission of data, simplification of router design and dramatic reduction in the number of routing tables used in the CCR network (as they exist only in the border devices).

Keywords: Cartesian routing, Exterior Gateway Protocol and Border Gateway protocol.

1 Introduction

The Internet is experiencing rapid growth: the number of connected host is doubling every year and the traffic is doubling every 6 to 10 months [1]. Even with Classless Inter Domain routing (CIDR) and network aggregation, the growth rate is still astonishing. Current research suggests that the growth will continue at the same rate into the foreseeable future [2].

Existing routing protocols, employing distance vector, link state, and path vector algorithms, require the exchange of routing information in order to construct and maintain the routing table. For example, the Border Gateway Protocol (BGP), the de facto inter-Autonomous System (inter-AS) routing protocol, requires that every BGP speaker establishes and maintains a Destination-Next Hop routing table. Each BGP speaker maintains a routing table by exchanging Network Layer Reachability Information (NLRI) via broadcasts and incremental updates with other BGP speakers. Whenever a packet is received, each BGP speaker must perform a look-up for the destination entry, searching for the 'next hop'. If the 'next hop' is not reachable via a direct connection, a recursive look-up for a route to the 'next hop' is required [3][4].

As networks increase in size, the memory requirements for the routing tables, the time taken to search the tables and the bandwidth for exchanging routing information increase proportionally. This can lead to serious bottlenecks, especially in Internet backbone inter-AS routers.

With the wide deployment of Global Positioning System (GPS) in the near future, the location information could become a primary attribute of each routing node [5].

Therefore, a high-speed geographic packet routing technique known as Cartesian routing is developed in [6]. The Cartesian routing algorithm allows the router to maintain the minimum state information and route packets without a routing table. The routing decision is topology dependent via integrating geographical location with the Cartesian address. Furthermore, Cartesian Core Routing (CCR), an extension of Cartesian routing algorithm is discussed as an inter-AS routing protocol to transparently replace BGP.

The remainder of this paper is organized as follows. In section 2, the fundamental

Cartesian routing methodology is described. Section 3 examines the extension of Cartesian routing, CCR. A comparison of BGP and CCR is presented in Section 4. Current and future research in this field is highlighted in section 5. The last section then concludes the paper.

2 Cartesian Routing and Cartesian network [6]

The fundamental principles of Cartesian routing can be illustrated using a linear routing algorithm in a one-dimensional network topology. In this topology, each router is associated with two ports (east and west), allowing it to connect to, at most, two other routers. Every router is bound to a unique address and maintains no state information other than this address.

Linear routing is achieved in the network by imposing an ordering on the routers which is based upon the unique router addresses; for example, west-to-east in ascending order (unless otherwise specified, west is always considered 'less than' east). When a packet is to be transmitted on the network, the transmitting router's layer 3 determines the packet's initial direction by examining the destination address. If this is less than the router's address, the packet is queued for transmission on the west port, otherwise on the east port. When a packet arrives at a router, its address is compared with the router's address. If the addresses are the same, the packet can be kept; if the packet arrived from the east (west) and is greater than (less than) the router's address, it is discarded; otherwise the packet is forwarded out the opposite port from which it arrived.

A Cartesian network consists of a set of collectors and one or more arterials, as shown in figure 1. Each collector is a chain of collector routers running east west sharing a common latitude. Collector routers have two ports (east and west) to exchange packets "horizontally". Each collector router also has a local port, which allows it to connect to a set of local hosts. Arterials exchange packets between collectors. Each arterial router, except the most northerly and the most southerly, has, at least, four ports (north, south, east and west). Arterials need not share a common longitude.

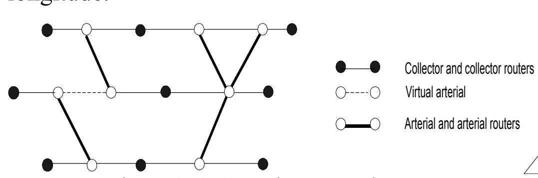

Figure 1: A Cartesian network

In a Cartesian network, the imposed topological structure relieves each router from maintaining routing tables. Each router is bound to a unique address (for example, a latitude and longitude). Both collector and arterial routers implement the linear routing algorithm: collector routers examine latitudes while arterial routers examine both latitudes and longitudes.

The state information maintained is minimal: each router maintains an Arterial Direction Indicator (ADI) that indicates which of its latitudinal ports leads to an arterial and whether the arterial connects to the north, south, or both. A router's ADI is updated when an arterial router sends an Arterial This Way (ATW) message out its collector ports or when a router detects a change in the state of its links.

An arterial router differs slightly from a collector router in that it can have multiple links leading to other arterials, for both fault tolerance (should an arterial link fail) and potential shortcutting. A virtual arterial is a collector that doubles as an arterial, allowing a continuous path from north-to-south, as shown in figure 1.

3 Inter-Autonomous System routing: Cartesian Core Routing (CCR)

Existing routing protocols, like BGP, are both time-consuming and resource-intensive. CCR, an adaptation of Cartesian routing, is proposed to transparently replace the inter-AS routing protocol.

3.1 Topology

In the CCR implementation, a Cartesian network is the core (or backbone) network and is responsible for exchanging inter-AS traffic. Each AS is connected to the core network via Internetwork Router (IR) – Encompassing Router (ER) links as shown in figure 2. An ER is a device connected to a collector in the CCR network and handles all internetwork traffic for its AS.

Figure 2: Cartesian Core routing Topology

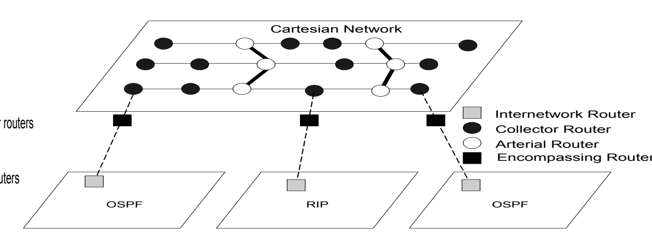

An IR, on the other hand, exchanges packets between the ER and its local AS. Each AS continues to use its own Interior Gateway Protocol (IGP), such as RIP [7] or OSPF [8], while the core network employs the Cartesian unicast routing protocol.

3.2 Address structure

The Cartesian address includes the geographical information (for example, longitude and latitude) which is enough to route packets. However, in order to connect IP-based ASs employing existing IGPs, the IP address needs to be translated to a CCR address. The CCR header consists of the destination ER's Cartesian address. Therefore, the Cartesian packet consists of two parts: IP packet received from the AS and the CCR header as shown in figure 3.

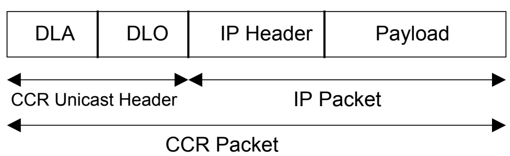

DLA - Destination Latitude

DLO - Destination Longitude

Figure 3: CCR Unicast Packet

Since the equatorial circumference of the Earth is 40,075km, and the polar circumference is 39,940km, a 48-bit Cartesian address (24 bit for longitude and 24 bit for latitude) could identify an area approximately 2.4m × 2.4m.

3.3 Routing mechanism

The routing inside an AS depends upon the local IGP. Any packet with a destination address that does not match an entry in the routing table of the interior router is sent to its ER via its local IR.

The routing in the Cartesian network is done using the Cartesian address while routing in the AS is done using IP addresses. For inter-domain communication between ASs via the Cartesian network, IP-CCR address translation is required. The ER, which is outside the Cartesian network, does this by encapsulating the IP packet with its corresponding destination ER Cartesian address. This Cartesian address is looked up in an IP-CCR address translation table. This translation table is established via Cartesian broadcasts [9] during initialization and if a change in the link-state is detected.

Within the Cartesian network, routing tables and NLRI-like global propagation are not necessary since each router determines the next hop based entirely on the result of the comparison of the local router's Cartesian address and the packet's destination Cartesian address.

When the packet reaches the ER of the destination AS, the ER removes the CCR header from the packet and forwards it to the IR inside the AS. Finally, the local IGP is employed to route this packet to destination IP host.

4 Comparison of BGP and CCR protocols

4.1 Similarities

- BGP and CCR are both designed to exchange routing information between border routers of their respective ASs.
- BGP and CCR both update their routing table only when a change of link-state is detected. Also, only the link-state changes will be advertised.

4.2 Differences

- The routing table in CCR is simpler than that implemented by BGP. The routing table of BGP routers contains a list of known routers, the addresses they can reach, and a cost metric associated with each path. On the other hand, the routing table of the ER includes a list of destination addresses and the Cartesian addresses corresponding to them. This suggests that less memory is needed in ERs.
- BGP and CCR are different in the number of look-ups in the routing table. As opposed to BGP's look-up for Next Hop in every BGP speaker on the path, CCR protocol does the look-up only once. This is done at the ER, when the packet enters the CCR network.
- BGP and CCR are different in their best path selection algorithms. BGP routers select the best path by comparing the cost of different paths in its routing table, while in CCR there is no path selection required (the path is determined by destination Cartesian address).

5 Current and future research

5.1 Optical Routing within the CCR network

The implementation of Cartesian routing considered in this paper is purely electrical. That is, data transmission and path determination are done electronically. In [10], the author proposes the use of optical signals in a Cartesian network called CORE (Cartesian Optical Routing Environment). In this scheme, the electrical packets from an IP network are converted into Cartesian packets by encapsulating the IP datagram with the Cartesian Address. This is then converted into optical signals and transmitted to the destination. The novelty of this scheme is the fact that o-e and e-o conversions are done only when there is a change in either latitude or longitude.

Research is ongoing to implement this scheme in the CCR network so as to take advantage of the increased speed and more efficient utilization of bandwidth provided by optical transmission.

5.2 Achieving Fault Tolerance in a CCR network

A redundancy scheme for the CCR network, similar to the one used in BGP to provide fault tolerance is being examined. Fault tolerance in the CCR network is achieved by using redundant pairs of ERs and IRs for each AS connected to the CCR network as shown in figure 4. For simplicity, only two links are illustrated in this paper. One ER acts as the active router while the other acts as the standby router (In this case we can assume ER1 to be active while ER2 acts as the standby router for AS1). The two ERs are connected to different Cartesian routers in the Cartesian network and hence have different Cartesian addresses.

The active ER periodically sends a keep alive message to its standby ER. If the standby ER ceases to receive this update, it assumes that the active ER either has a fault or one or some of its links are down. The standby ER takes over as the new online ER by broadcasting its Cartesian address to all the other ERs. If there is a change in link-state within an AS, the online IR for that AS sends the link-state update to its connected online ER using the existing BGP broadcast mechanism. The ER updates its routing table and broadcasts the new information to the other ERs via the CCR network.

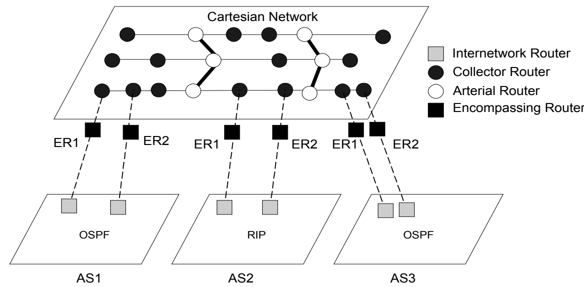

Figure 4: Redundancy in the CCR network

5.3 Other related research

Ongoing related work in our research group includes microcontroller-based and Optical Cartesian router design and implementation, fault tolerance in the CCR network, multicast and broadcast in the Cartesian network and optimal Cartesian network design.

6 Conclusions

This paper has described the Cartesian routing algorithm, which compares the local router's location address with the destination location address of the packet to make the routing decision. With wide deployment of GPS, geographical location information is readily available. This further simplifies the design of CCR networks. By employing a structured topology, routing tables are not necessary in a Cartesian network. CCR, an adaptation of Cartesian routing, was discussed as a transparent replacement of BGP in the backbone network. The purpose of CCR is to inter-connect non-Cartesian ASs through a single CCR network. A translation table is maintained in each ER, the entry point of the core network, to encapsulate IP addresses into CCR addresses. The packet is then routed as a Cartesian packet in the core network. The CCR header is removed when it arrives at the destined AS and the IP address is used by the interior gateway protocol for routing to the destination host within the AS.

In CCR, the ASs communicate with the CCR network border device using BGP while the CCR border device uses the Cartesian routing algorithm within the CCR network. This ensures transparency of the CCR network to the users within the various ASs and a smooth transition in the deployment of CCR from the existing BGP backbone network.

The CCR algorithm has three cogent advantages over the existing exterior gateway protocols: faster routing and transmission of data, simplification of router design and dramatic reduction in the number of routing tables used in the CCR network (as they exist only in the border devices).

Acknowledgements

The research described in this paper is supported by an operating grant from the Natural Sciences and Engineering Research Council of Canada.

The authors would like to thank their colleagues Yanting Lin and Shehzad Pradhan for their invaluable contribution to this paper.

References

1. C. Huitema, “Routing in the internet”, Prentice Hall, 2000.
2. http://www.telstra.net/ops/bgptable.html Accessed on 28th March 2002.
3. Y. Rekhter and T. Li , “A Border Gateway Protocol 4 (BGP-4)”, RFC1771, 1995.
4. P. Traina, “BGP-4 Protocol Analysis”, RFC 1774, 1995.
5. T. Imielinski and J. Navas, *GPS-Based Addressing and routing,* RFC 2009, November 1996.
6. L. Hughes, O. Banyasad and E. Hughes, “Cartesian routing”, Computer Networks, vol. 34, pp.455-466, 2000.
7. G. Malkin, “RIP Version 2”, RFC2453, 1998.
8. J. Moy, “OSPF Version 2”, RFC2328, 1998.
9. L. Hughes, "Broadcast Communications in Cartesian Networks" Unpublished, 2002.
10. E. Broadfoot, "All Optical Networks", Masters Thesis, Dalhousie University, 2001

The authors

Dr. Larry Hughes is an Associate Professor in the Department of Electrical and Computer Engineering at Dalhousie University. His research interests include computer communications, embedded systems, and renewable energy. He is the author of over 50 papers and several textbooks. Dr. Hughes obtained his Ph.D. in Computer Science from the Computing Laboratory at the University of Newcastle upon Tyne in 1988.

Mr. Adefisayo Adegoke received a B.Sc. in Electrical and Electronics Engineering from the University of Ibadan, Nigeria in 1999. His research interests include computer networks, the next generation of the Internet, routing protocol design & implementation, network optimization & design and control systems. He is currently working towards his MASc in Electrical and Computer Engineering at Dalhousie University.

Mr. R Subramaniam Ganesh received a B.E in Electronics Engineering from University of Pune, India in 1998. His research interest include, Cartesian Routing, network routing protocols for the Internet and network security. He is currently working towards his MASc. in Electrical and Computer engineering at Dalhousie University.

Mr. Haiyu Song received a B.Eng. in Electrical Engineering from North China Electric Power University in P.R. China in 1995. His research interests include computer networks, IP Multicast techniques and Cartesian Routing. He is currently working towards his MASc in Electrical and Computer Engineering at Dalhousie University.