Performance and Cost of Broadcast Routing Algorithms in the Strategic Defense System Terrestrial Network †

Thu Van Vu

Harris Corporation
GASD 102/4844
P.O. Box 94000
Melbourne, FL 32902

Steven P. Risner

U.S. Army Strategic Defense Command
CSSD-SA-BT
P.O. Box 1500
Huntsville, AL 35807

ABSTRACT

Algorithms to perform the network control function for the ground-based elements of the Strategic Defense System (SDS) have been developed and evaluated. This paper compares the analytical and simulation results of two broadcast routing algorithms, namely Flood-and-Forward Routing and Constrained Flood Routing to identify a suitable candidate for routing traffic in a fiber-optic communication network which connects the SDS Command Center and the distributed ground-based sensor and weapon sites which also host various battle management functions. In addition to the usual performance and cost measures such as end-to-end delay, reliability, bandwidth consumption, etc., this paper specifically addresses the packet processing rate required at each relay node. Realistic packet processing delays are also included in the simulation to provide a complete picture of the overall delay statistics.

1. SDS TERRESTRIAL NETWORK

In an early concept of the SDS communication architecture, the ground-based elements are interconnected by a wideband fiber-optic network. As shown in Figure 1, this backbone network connects the seven sensor/weapon sites, the four ground-based radar (GBR) sites, and the Command Center. The three northern fixed sites (North West, North Central, and North East) have both Ground Surveillance and Tracking System (GSTS) and ground-based interceptor (GBI) assets. The two west coast sites (West-1 and West-2) and the two east coast sites (East-1 and East-2) are GBI-only sites. There are also six nodes in the fiber-optic network which serve only as communication network switches in order to provide connectivity to existing fiber-optic cables. The backbone fiber-optic network consists of leased point-to-point DS-3 services and SDS-unique cable (as required to reach the sites from existing DS-3 services). Each link between nodes consists of one fiber (or one DS-3 service) in each direction, and has a capacity of 45 Mb/s in both directions.

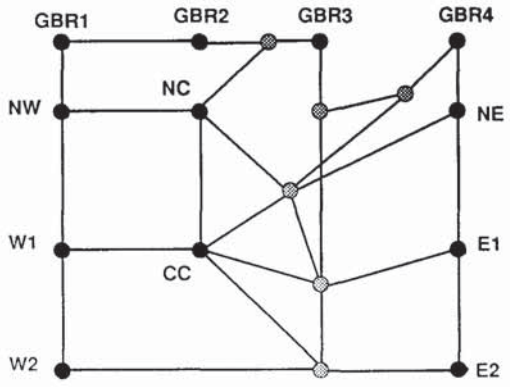

Figure 1. SDS Terrestrial Network

Each fixed site or mobile facility has a number of ground entry points (GEP) to Boost Surveillance and Tracking System (BSTS), Space Surveillance and Tracking System (SSTS), Ground Surveillance and Tracking System or ground-based interceptors. A BSTS/SSTS GEP is a single transmit-receive system that can link to either a BSTS or SSTS satellite. Multiple GSTS GEPs are required at each GSTS site

† This work was funded by the U.S. Army Strategic Defense Command under contract number *DASG60-88-C-0102*.

CH2981-9/91/0000-0622 $1.00 © 1991 IEEE

TABLE 1. Candidate Message Characteristics and Types of Services Required

Message Type	Size (data bits)	Peak Rate (msgs/sec/src)	Distribution	Reliability
C1—Sensor Data	176 - 632	10 - 5,000	Multicast	Medium
C2—Health & Status	256	0.2	Multicast	Medium
C3—Command & Control	128	100	Both	High
C4—Software Upload	Unlimited	100	Directed	High
C5—Communication Internal	TBD	Low	Both	Med. - High
C6—Tracking, Telemetry & Command	256	Low	Directed	Medium
C7—InFlight Target Update	205	30	Directed	High
C7—Target Object Map	3,276 - 24,396	10	Directed	High
C8—Discrimination Data	Short	Low - Med.	Multicast	High

to handle the number of GSTSs expected in flight simultaneously. A GBI GEP is a transmit only system to service as many as 100 inflight interceptors simultaneously.

A major ground site is itself a complex interconnection of subnetworks (LANs and MANs) which distribute the data entering the site through the backbone network or through one of the GEPs, as well as data generated by a local battle manager processor. Candidate message characteristics and the types of services required are summarized in Table 1 to provide a glimpse of the enormous amount of data that must be transferred in the SDS terrestrial network concept considered here.

A comprehensive set of algorithms to perform the network control function for the ground-based elements of the SDS communication network has been developed and evaluated. The algorithms cover end-to-end error control, broadcast routing, multicast routing, point-to-point routing, flow control, and failure recovery. Due to limited space, this paper can only describe and compare the analytical and simulation results of two broadcast routing algorithms.

2. BROADCAST ROUTING

The combination of broadcasting over point-to-point links and a very high generation rate of sensor data presents a unique routing problem in the SDS communication network. Data packets have to be duplicated and forwarded on every (two-way) link in order to reach all destinations. If broadcasting were not frequent, one could afford to be sloppy in routing a broadcast packet—for example, generating more duplicates than necessary. However, the combined sensor traffic from all sources in the SDS communication network may reach levels that exceed even the highest link capacity, making it imperative that the routing algorithm not compound the problem by generating more duplicates than necessary. Thus in addition to the ultimate goal of delivering messages promptly and reliably, the selected broadcast routing algorithm should not incur high bandwidth consumption, high packet processing rate, excessive memory usage, and long computation time.

Algorithms available for routing broadcast packets in a packet switching network include separate destination addressing, multidestination addressing, constrained flooding, minimum spanning tree forwarding, and reverse path forwarding [1,2]. Because of its simplicity and reliability, constrained flood routing is the algorithm of choice in many networks.

2.1 Constrained Flood Routing

In constrained flood routing, an arriving packet that has not been seen before is copied and forwarded on all outgoing links except the one it arrived on. Packets that have been seen are simply discarded. To keep track of already seen packets, a sequence number is assigned to each packet by the source, and a "constraint table" (or bit map) is maintained at each node to log packets. The log for a particular packet has to remain in the constraint table for some time, at least for the duration of the broadcast, before it can be cleared and reused. The size of the constraint table is therefore proportional to this predefined time for packets to live, the maximum traffic generation rate by all sources, and the total number of sources. On top of this hefty memory requirement, constrained flood routing also generates a large number of packet copies, $N(L-1)+1$, where N is the number of nodes in the network, and L is the number of links per node.

As for advantages, constrained flood routing is most noted for its robustness. Packet copies that seem to be such a waste at first glance are much needed to

replace lost copies. The algorithm always finds the shortest routes possible, and its routes adapt instantaneously to changes in the network.

When constrained flood routing is applied to the SDS communication network where there are very large quantities of sensor data, the disadvantages clearly outweigh the advantages. The abundance of packets not only puts serious strains on the link capacity, induces higher levels of network congestion, but also dictates a high packet processing capacity at each relay node. These concerns have prompted the development of a new algorithm for routing broadcast data.

2.2 Flood-and-Forward Routing

Flood-and-forward is a broadcast routing algorithm that was designed for handling the extremely high volume of sensor data generated by sensor platforms in the SDS communication network. The objective of the algorithm is to closely match the delay performance of constrained flood routing without consuming as much bandwidth and memory. This objective is achieved by using constrained flood routing only for a small number of special packets. As these packets crisscross the network to get to their destinations, the routes on which they travel are carefully recorded by every node into a broadcast routing table. When the broadcast routing table is complete, the minimum spanning tree contained therein is used to forward other broadcast packets.

Thus when a node has to broadcast a message, it will decide whether to *flood* or *forward* the message depending on the current status of the broadcast routing table.

Flood Phase

As determined by a preset rate (for example, ten per second), the source node occasionally issues a *Scout Packet* to establish the broadcast routes. The scout packet reaches all nodes by means of a constrained flood. A slight modification of the protocol for constrained flood routing is needed to implement the flood phase of flood-and-forward routing:

- If the scout packet, identified by *Source Id* and *Scout Label,* has not been seen before by the node,
  - — Accept the scout packet.
  - — Log the scout packet in the constraint table.
  - — Copy and forward the scout packet on all links except the link it arrived on.
  - — Set an *Ack Scout Timer* for receiving acknowledgements of the scout packet.
  - — Send back an *Ack Scout Packet* to the neighbor node that forwarded the scout packet.
  - — Enter that neighbor node in the *Received From* column of the broadcast routing table indexed by *Source Id* and *Scout Label.*
- Otherwise, discard the scout packet.

The ack scout timer should be set roughly to a round trip delay between two adjacent nodes, enough time for the scout packet to reach a neighbor and for the ack scout packet to return to the forwarding node.

An ack scout packet is generated every time a node accepts a scout packet. Ack scout packets for the same scout represent the branches of the (minimum spanning) tree of routes traversed by the scout packet. The *Source Id* and the *Scout Label* are sufficient for proper acknowledgement of a scout. Due to their brevity, ack scout packets may be piggybacked and redundantly transmitted for added reliability. As ack scout packets are received by each node, segments of the broadcast routes are extracted and recorded in the broadcast routing table as follows:

- If ack scout packets are received before their corresponding ack scout timer expires, enter the list of acknowledging neighbors in the *Send To* column of the broadcast routing table indexed by *Source Id* and *Scout Label.*
- Otherwise, enter *Null.*

Forwarding Phase

A short time after generating a scout packet, the source node may use the routes in the broadcast routing table indexed by its *Source Id* and the *Scout Label.* The activation of the routes must be timed precisely so that when the first non-flood packet is forwarded along those routes, the broadcast routing table is ready at the forwarding nodes as the packet reaches them. In normal situations, the table is ready as soon as ack scout packets are received; i.e., a link round trip delay after the scout packet has been forwarded. Furthermore, if variations in the link delay are small, the time separation between two packets following the same route is expected to remain the same at every node along the route. Thus the source node should wait for a maximum link round trip delay plus some small delay before activating the routes.

Multiple values of *Scout Label* allow the source node to send out another scout packet without having to wait for the previous scout packet to be acknowledged. As soon as a new set of routes is available, it becomes the active set of routes at the source node, but older sets of routes are still maintained by other nodes long enough for packets to reach their destinations. By

recycling a number of preallocated scout labels, one after another, the source node manages the multiple sets of routes in an orderly fashion. The number of scout labels is equal to the product of scout rate and routes' life.

In the forwarding phase, the source node sends out broadcast packets identified by its *Source Id* and a *Scout Label,* thereby signaling to receiving nodes that these packets ought to be forwarded using the routes that have been built by the scout packet with the specified *Scout Label.*

Packets are copied and forwarded by nodes according to the routing information contained in their broadcast routing tables:

- Use *Source Id* and *Scout Label* to look up entries in the *Received From* and *Send To* columns.
- Suppress packets that did not come from the neighbor on the *Received From* list.
- Forward packets to neighbors on the *Send To* list.

3. PERFORMANCE AND COST

Both flood-and-forward and constrained flood routing algorithms have been simulated for the network depicted in Figure 1, using high fidelity models of OSI layers 2 through 4 and representative traffic loads. The total network offered load may reach 20 Mb/s at times. Listed here are some of the performance and cost measures that are used to evaluate the algorithms:

- Packet processing requirements
- Bandwidth consumption
- End-to-end delay

3.1 Packet Processing Requirements

The worst-case packet processing requirements are considered first since these figures dictate the necessary speed of the packet processor.

In constrained flood routing, each packet is received, at worst, L times at each node, where L is the number of links per node. Hence, processing is proportional to L times the offered load. Assuming an offered load of 50K pkts/s and 5 links/node, the packet processing rate must be 250K pkts/s per node.

In contrast, there are three types of packets in flood-and-forward routing. Each non-flood broadcast packet is received once at each node. Processing of these packets is thus proportional to 1 times the offered load. Each scout packet is received at most L times at each node, but the scout rate is fixed (at ten per second per source node, for example). Each accepted scout is acknowledged to a neighbor node, and processing of ack scout packets is also fixed by the scout rate. Hence, the total packet processing requirement for flood-and-forward routing is approximately 51K pkts/s per node.

Flood-and-forward routing requires less processing since the majority of broadcast packets is received only once at each node, whereas in constrained flood routing, each node has the potential of receiving an identical packet over each of its links. This fact is confirmed in Figure 2, where the simulated packet processing rate is plotted. The cumulative distribution function (CDF) curves show that fifty percent of the nodes in the network process approximately 45K packets/sec for flood-and-forward routing, compared with 90K packets/sec for constrained flood routing.

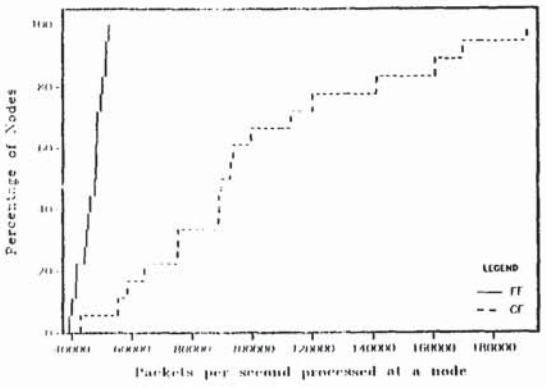

Figure 2. Packet Processing CDF

3.2 Bandwidth Consumption

With the exception of scout packets, packets in flood-and-forward routing are forwarded along branches of a spanning tree. For each non-flood broadcast, only one one-hop packet is generated per destination, thus achieving the optimal use of bandwidth. The bandwidth consumption CDF is portrayed in Figure 3. As shown, fifty percent of the links in the network consume less than 15% bandwidth when using flood-and-forward routing, compared with 30% bandwidth when using constrained flood routing.

3.3 End-to-End Delay

Packets in flood-and-forward routing are forwarded on routes selected by a constrained flood. But with fewer

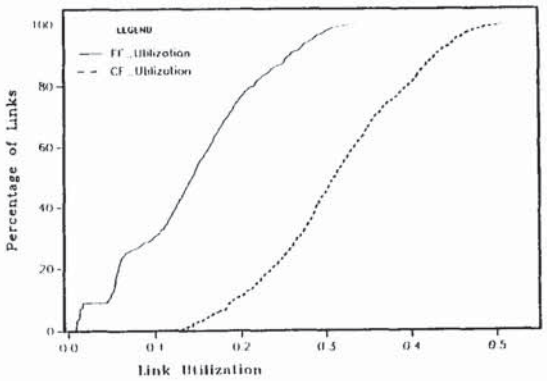

Figure 3. Bandwidth Consumption CDF

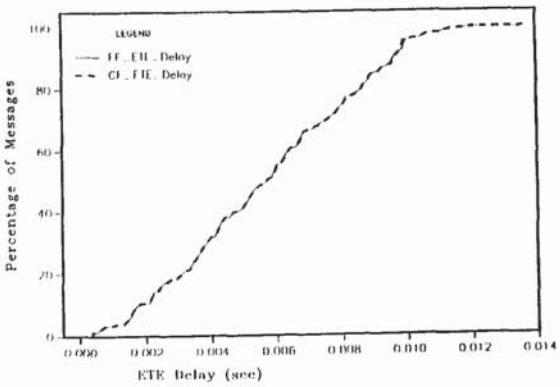

Figure 4. End-to-End Delay CDF

packet duplicates, thus less congestion and queue delay, end-to-end (ETE) delay may be even better than that of constrained flood routing. Figure 4 shows CDF curves for ETE delays from the broadcasting source to all destinations. ETE delay consists of transmission, queue, and propagation delays. The plot shows that flood-and-forward routing exhibits the same excellent delay performance as constrained flood routing. The shorter average queue length experienced by packets in flood-and-forward routing (Figure 5) is a direct consequence of the optimal use of bandwidth.

4. CONCLUSIONS

To summarize, the disadvantages of using constrained flood routing include the high bandwidth consumption, high packet processing and memory requirements. Flood-and-forward routing seeks out the same (least-delay) routes as constrained flood routing, but has optimal bandwidth consumption. The results are that for the SDS terrestrial network under consideration, flood-and-forward requires one half the average packet processing capacity and one fifth the worst case packet processing capacity of constrained flood. It has been rightly argued that flood-and-forward routing is not as reliable as constraint flood routing. However, this is not a real concern since fiber-optic networks have extremely small bit error rates. To tip the scale in favor of flood-and-forward routing, the algorithm also provides some degree of load balancing with a quick response time and can be configured to provide multicast routing and point-to-point routing.

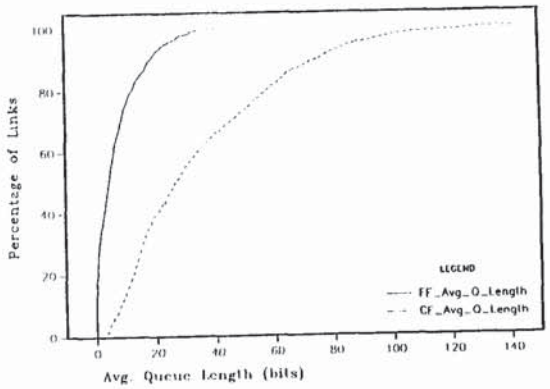

Figure 5. Average Queue Length CDF

5. REFERENCES

[1] Tanenbaum, A.S., *Computer Networks*, Prentice-Hall, Englewood Cliffs, New Jersey, 1981.

[2] Dalal, Y.K. and Metcalfe, R.M., "Reverse Path Forwarding of Broadcast Packets," *Communications of the ACM*, vol. 21, pp. 1040-1048, December 1978.