# Asymmetric Routing

Jani Lakkakorpi
jani.lakkakorpi@nokia.com

## Abstract

**New transmission media have been proposed to provide higher bandwidth to the Internet – satellite links offer unidirectional and asymmetric access. Both unidirectional and asymmetric links should be integrated into the Internet transparently. However, there are some specific problems that must be solved in the case of unidirectional links. The UniDirectional Link Routing working group [1] focuses on supporting unidirectional links on top of bi-directional internetwork. In order to support unidirectional links, there are two approaches: 1) modifications to common Internet routing protocols (RIP, OSPF etc.), or 2) a new layer between the network interface and the routing software to emulate bi-directional links through tunnels.**

## 1 Introduction

The Internet consists of networks interconnected via bi-directional and symmetric links with different bandwidths. Thus, Internet routing and transmission control protocols have been designed to optimize the transmission assuming bi-directionality and symmetry.

The use of satellite links is a relatively new technique [1] that could offer higher bandwidths for the Internet. These links have high downlink bandwidth via satellite, and a lower uplink bandwidth via regular connections. Broadcast satellite links could be used to provide a high speed receive only access to the Internet, with the return traffic sent through a low speed link. Both unidirectional and asymmetric links need to be integrated in the Internet transparently. Unidirectional links are a special case of asymmetric links.

However, there are some problems (such as dynamic routing and transmission control) to be solved in the case of unidirectional links Three cases for unidirectional and/or asymmetric links can be established:

1. Unidirectional links on top of bi-directional underlying network (wired Internet).

2. Bi-directional islands connected via unidirectional links.

3. The general case of asymmetric and possibly unidirectional links.

The focus in the UniDirectional Link Routing working group [1] (and in this paper as well) has been on the Case 1.

Common Internet routing protocols such as RIP (Routing Information Protocol), OSPF (Open Shortest Path First) and DVMRP (Distance Vector Multicast Routing Protocol) are not able to work properly with unidirectional links, because they assume link bi-directionality and symmetry. In the Case 1, there are two proposed approaches. The first one is based on the modification of the common routing protocols in order to support unidirectional links [2] while the second one is adding a layer between the network interface and the routing software to emulate bi-directional links through tunnels [3]. Although the second approach seems to be the short-term solution [4], this paper is mostly focused on the first approach.

Section 2 briefly describes the main points of the Internet Draft 'Supporting Unidirectional Paths in the Internet' [2] while Section 3 takes a look at the problem of unidirectional link and OSPF [5]. Section 4 is about an asymmetric routing experiment in the U.S. Navy [6], and Section 5 contains my own experiments in asymmetric routing with the Network Simulator [7].

## 2 Supporting Unidirectional Paths

There exists a proposal for a low-cost solution to deliver high bandwidth services over wide geographical areas via the use of broadcast satellite networks [8]. Since this solution is based on low cost receivers with zero bandwidth return, the connection over the satellite link is unidirectional. The integration of these satellite networks with the global Internet requires changes in common routing protocols.

### 2.1 A New Architecture

The advantage of a satellite network is to provide high bandwidth services independent of the user's location over a wide geographical area.

A satellite network consists of two types of stations: feeds and receivers. Every receiver has a satellite dish connected to a user station (basic access) or to a router (subnetwork access). The user station has an extra interface, and the router has one or more extra interfaces, connected to the Internet.

After the information is sent from the feed to a satellite, it will be broadcast to all the receivers that belong to the satellite coverage. Installing feeds in strategic positions over the Internet will create shorter paths and higher bandwidth provided by the satellite network.

**Basic Access**

Basic access corresponds to the case when each receiver has a private satellite dish. The user is also connected to the Internet via modem connection. This station has therefore two IP addresses, one for the satellite subnetwork and the other for the regular connection subnetwork.

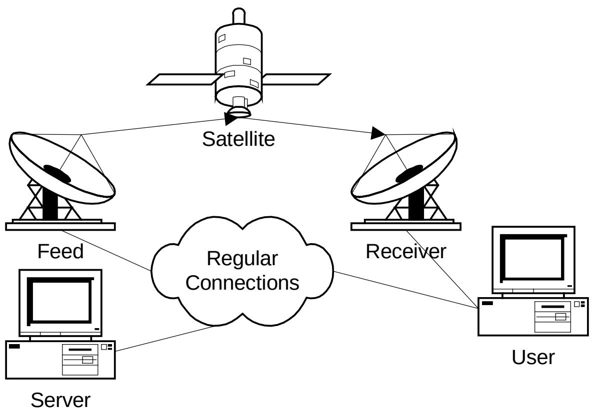


**Figure 1. Basic access**

All requests to a remote server are sent via modem connection, and responses from the server should be routed by a feed on the satellite network.

**Subnetwork Access**

Subnetwork access corresponds to the case when the subnetwork router has a satellite dish. This router also has regular connections to the Internet.

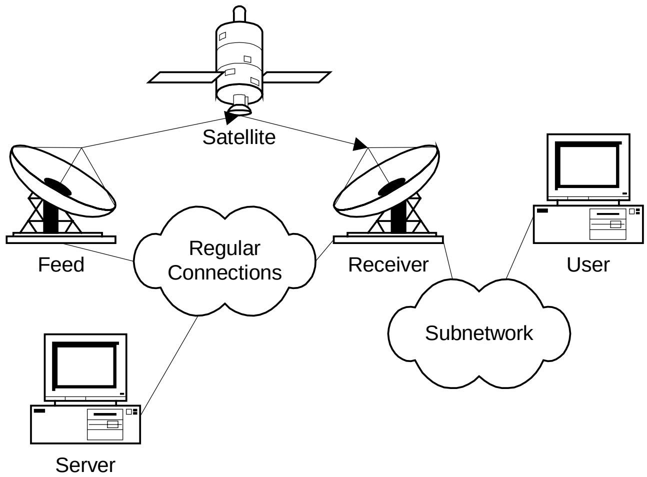


**Figure 2. Subnetwork access**

This configuration requires only one satellite dish for the whole subnetwork. The management is also located in a single place – the subnetwork router.

### 2.2 Solutions

For the basic access and the subnetwork access the authors of [2] propose the following solutions.

**A Dynamic Routing**

In order to handle unidirectional links, some modifications should be applied to the routing protocols. Most of them namely assume that communication between neighbor routers is bi-directional.

**Basic Access**

Since an ARP (Address Resolution Protocol) request sent by a feed to a host belonging to the satellite network is not able to have a response from the receiver, the ARP protocol is not able to function properly.

Routing for that type of user station differs from traditional routing; the station must have two IP addresses – one for the satellite network, and one for the slower uplink connection. Users have to send their packets via the slow interface, and incoming packets should be routed to the default address, which is the satellite network address.

**Subnetwork Access**

Feeds and receivers are now considered as IP routers. But how can a receiver announce routes to feeds since it is not able to communicate directly with them?

The work of UniDirectional Link Routing Working Group is mainly based on the study of the most common routing protocols that will be used by feeds and receivers such as RIP, OSPF, and DVMRP for multicast routing.

Unlike receivers, feeds are able to broadcast routing messages via the satellite network. A feed will expect to receive responses from all of its interfaces. However, a feed can not receive messages from the satellite network.

In order to announce routes, receivers must send their routing messages to the unicast address of each feed via regular connections. Because of the long distances between the feeds and the receivers, their regular Internet interfaces seldom belong to the same subnetwork. Routing protocols, however, ensure security by checking that the sender's address belongs to the same subnetwork as the interface, which received it. Therefore this routing information will not be taken into account because the packet will be rejected.

These are the problems that occur when trying to handle unidirectional links by common routing protocols. Specific problems related to RIP, OSPF, and DVMRP are described in other documents [1, 4]. Next section takes a quick look into handling of unidirectional links with OSPF.

# 3 Handling of Unidirectional Links with OSPF

An Internet Draft 'Handling of unidirectional links with OSPF' [5] describes the modifications which should be applied to OSPF in order to make the communication over unidirectional links feasible.

OSPF is a dynamic routing protocol used in the Internet known as Internal Gateway Protocol. It was designed to work on networks where adjacent gateways communicate using the same link in both directions. Links may have different delays and throughputs in different directions, but they have to be duplex.

## 3.1 OSPF Restrictions

Receivers are not able to send any packet via the satellite link. Nevertheless, they have to communicate with the designated router to indicate that they are ready to receive packets, and that they are synchronized with their neighbors.

If we had a network with only feeds, OSPF could be used almost unchanged. In a network that consists of both receivers and feeds, OSPF requires some modifications.

## 3.2 Proposed Solution

The authors of [5] present an example case, where they assume that two gateways, G1 and G2, are connected to symmetric and asymmetric networks. Since G1 uses OSPF, it will never consider G2 as a neighbor because the link is unidirectional, and therefore G1 will send its packets to the regular connections. In order to make G1 consider G2 as a neighbor; OSPF should be modified to take unidirectional links into account.

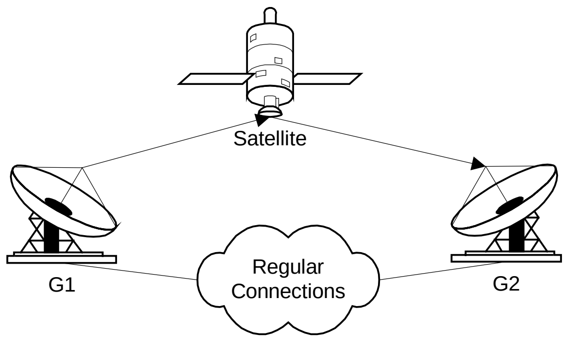


**Figure 3. Symmetric and asymmetric networks**

**Authentication Scheme**

All OSPF protocol packets (such as Hello and Database description) share a common header of 24 bytes. The OSPF packet header includes an 8-bit long authentication type field and 64 bits of data used by the appropriate authentication scheme determined by the type field.

It is suggested [5] that all packets sent to the satellite channel should be authenticated, using either simple password authentication or, preferably, a stronger type of authentication. The authentication code will be specific to the satellite network stations.

Protocol packets sent over the satellite network will be authenticated, and their processing will be different as routers receive them.

### The Hello Protocol

Feeds are set up to the highest router priority on the network in order to make them designated routers of their area.

The Hello protocol usually ensures that communication between directly connected routers is bi-directional. This must change in order to allow the protocol to work asymmetrically between feeds and receivers connected to the satellite network.

When sending Hello packets over the satellite network, feeds will authenticate them as 'satellite packets' by setting a type field and add a specific authentication field. After receiving these Hello packets, receivers examine the authentication code. They will note that this packet was sent by a satellite feed, and add the packet source to a list of 'potential neighbors'.

Receivers periodically send Hello packets to their potential neighbors. These packets are not sent to the multicast address 'to all OSPF routers', but a copy is sent to the unicast address of each potential neighbor. These packets are also authenticated as 'satellite packets'. When receiving these Hello packets, feeds will process them even if they are routed by another interface.

### Network Link Record

The first steps in the formation of the shortest path tree are the links between routers and transit networks. The network links advertisement describes all the routers that are attached to a transit network.

The authors of [5] suggest that the network link record should be extended to supply further information concerning the connected routers. Instead of having just a list of connected routers, we could have a list of routers, which can only send or receive packets.

| Octet 0 | Octet 1 | Octet2 | Octet 3 |
|---|---|---|---|
| Network Mask | | | |
| Number of Attached Routers | | | |
| Attached Router #1 | | | |
| … | | | |
| Number of Senders Only | | | |
| Attached Sender #1 | | | |
| … | | | |
| Number of Receivers Only | | | |
| Attached Receiver #1 | | | |
| … | | | |

**Figure 4. List of connected routers**

- Network Mask is the IP network mask for the network.
- Number of Attached Routers represents the number of routers attached to the transit network, which can send and receive packets.
- Number of Senders Only represents the number of routers attached to the transit network, which can only send packets. This field is followed by the list of router's ID.
- Number of Receivers only: represents the number of routers connected to the transit network, which can only receive packets. This field is followed by the list of router's ID.

Senders and Receivers Only are detected when a router notices that a packet is authenticated as 'satellite packet'.

As in the basic OSPF, a graph is built based on the information found in the network link record and the router link record. Unidirectional communications are represented by single vertices and thus the shortest path tree can be calculated with unidirectional links, too.

### Processing Protocol Packets

All protocol packets sent by the feeds via the multicast link are authenticated.

From the list of 'potential neighbors', receivers can find feeds' IP addresses. Receivers send their protocol packets to the unicast address of each feed through the regular connection. These packets are also authenticated as 'satellite packets'.

When receiving packets authenticated as 'satellite packets', feeds will process them even if they are routed by another interface.

# 4 Asymmetric Routing in the U.S. Navy

R. S. Starsman from the U.S. navy has written a paper about his implementation of asymmetric routing. This section describes the main points of that paper.

The technique presented in [6] requires no modification to the client applications or workstations. An asymmetric routed connection can be set up with slight router configuration changes and minor hardware modifications at the router-network interface. An extension of asymmetric routing is also discussed. This extension eliminates the need for any modification to existing communications equipment.

According to [6], a cruiser operating at sea is allocated a SHF (Super High Frequency) 56 kbit/s full-duplex IP data connection to the SIPRNET (Secret Internet Protocol Routed Network). It means that the ship is allocated fixed 56 kbit/s uplink and 56 kbit/s downlink paths regardless of the ratio of up- and downlink usage.

However, tactical data flow is often quite asymmetric. Tactical units are usually information consumers and thus receive considerably more data than they transmit. For example, the ship might request some weather information with a few kilobits worth of IP packets. The returned data might in turn be several megabits worth of weather forecasts with graphics. In a situation like this, the 56 kbit/s uplink may contain some overcapacity.

Instead of allocating 56 kbit/s in both directions, it is proposed in to allocate the bandwidth in asymmetric fashion. In the example above, we could allocate 12 kbit/s of bandwidth for uplink and 100 kbit/s for downlink. Thus it is possible to effectively double the throughput while keeping the aggregate load on the satellite transponder the same.

## 4.1 Description of the Technique

Usually, routers process incoming and outgoing data from a single connection at the same rate. However, a router can be configured to process and route data at two different speeds. This technique has been demonstrated to provide an improvement in data throughput of 10 to 15 times over conventional techniques.

The improvement is limited due to the acknowledgments required in TCP. Increasing the packet window size could further increase the throughput improvement. UDP based connections have no limitations of this kind.

Asymmetric routing can be implemented on any router with two free serial ports. One port is configured as the transmit connection and the other as the receive connection. Static routes are established to force data over the appropriate paths. To achieve this, a few pins on the router serial port must be jumpered to make the router think that it has a full-duplex path. Once this is completed, workstations behind the router have a transparent asymmetric TCP/IP connection and can make unrestricted use of any IP-based application such as a web browser, e-mail, etc.

## 4.2 Implementation

An asymmetric routed connection provides two unidirectional paths with different data speeds between the client and the network. Client applications have transparent asymmetric access to the network.

The tactical router and the gateway router each use two serial ports to make this connection. One pair of ports operates at a lower speed and passes information requests from a tactical client to a server anywhere in the network. This connection might operate on an UHF or EHF channel, or it could be a modem connection via telephone cables. The other pair of ports operates at a much higher speed and returns the requested data. This connection could run over any higher bandwidth system.

Most of the technology that is required to implement this kind of connection already exists and is implemented. However, splitting an IP feed from a router into two unidirectional streams with different data rates had never before been implemented.

## 4.3 Results

Data throughput tests were run at low-speed port rates of 2.4, 9.6, and 56 kbit/s to simulate various tactical data connections ranging via the low-speed serial port. The high-speed link was set to 1.024 Mbit/s for all tests. Data transfer was measured by transferring a 1.5 MByte file from server to client and measuring the transfer rate. Throughput measurements were taken with a full-duplex low-speed connection in place and then with the connection augmented by the unidirectional high-speed path. The results of these tests are shown in Table 1.

**Table 1: Starsman's results**

| Low-speed Rate (kbit/s) | Throughput Without Augmentation (kbit/s) | Throughput With Augmentation (kbit/s) | Incr. (%) |
|---|---|---|---|
| 2.4 | 2 | 31.2 | 1560 |
| 9.6 | 8 | 124.8 | 1418 |
| 56 | 52.2 | 664.8 | 1278 |

With this asymmetric arrangement, the end-user's 2.4 kbit/s connection looked like a connection more than 15 times faster.

Although this alone gives some evidence of the benefits of asymmetric routing, there are some improvements that could be made; the routing configuration requires two serial ports per router. It would be useful to reduce the number of router resources required to build an asymmetric router connection. This can be done by building an asymmetric data buffer that handles asymmetric data flow and converts it into a symmetric flow expected by the router serial port [6].

# 5 Own Results

I tried a similar experiment as described in [6] with the Network Simulator [7]. First, I measured file transfer delays for different file sizes with a duplex link, and then I replaced this duplex link with two simplex links. The results, which are similar to the results of [6], are in Table 2. The simulation model is depicted in Figure 5.

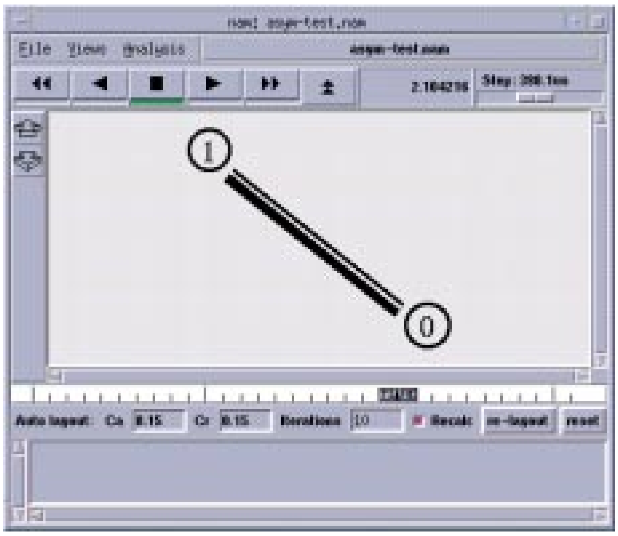


**Figure 5. Simulation topology**

**Table 2: Own results**

| File Size (kB) | Transfer Delay With 50 kbit/s Duplex Link [s] | Transfer Delay with 10 kbit/s Uplink and 90 kbit/s Downlink [s] |
|---|---|---|
| 10 | 1.7606 | 1.067 |
| 20 | 3.4822 | 2.023 |
| 30 | 5.1974 | 2.976 |

# 6 Conclusions

Improving user connectivity to the Internet at reasonable cost seems now possible - both for basic access and subnetwork access.

However, handling unidirectional links by standard routing protocols is not an easy task and currently not supported. It requires changes in the current protocol specifications. These changes should be transparent for routers not connected to satellite networks.

Keeping this in mind, tunnelling will most probably be the short-term solution for the problem [4].

# 7 References

[1] The UniDirectional Link Routing home page: http://www-sop.inria.fr/rodeo/udlr/index.html

[2] E. Duros, W. Dabbous: Supporting Unidirectional Paths in the Internet, Internet Draft, November 1997.

[3] M. Karir, Y. Zhang: An Experimental Study of Asymmetric Routing, WOSBIS'98, Dallas, Texas, October 30, 1998.

[4] W. Dabbous, E. Duros, T. Ernst: Dynamic Routing in Networks with Unidirectional Links, WOSBIS'97, Budapest, Hungary, September 1997.

[5] E. Duros, W. Dabbous: Handling of unidirectional links with OSPF, Internet Draft, March 1996.

[6] R. S. Starsman: Asymmetric Routing, Milcom 97, Proceedings Volume: 2 , 1997 , Pages: 1021 -1025 vol.2

[7] UCB/LBNL/VINT Network Simulator – ns (version 2), http://www-mash.cs.berkeley.edu/ns/ns.html

[8] V. Arora, N. Suphasindhu, J. Baras, D. Dillon: Asymmetric Internet Access over Satellite-Terrestrial Networks, http://www.isr.umd.edu/CSHCN/projects/aia/index.htm