A Continental Communication Network for Remote Sensing and GIS Data

Graham J. Hicks[1], David G. Goodenough[2,3], A. S. (Pal) Bhogal[2], and K. O. Niemann[4]

[1]Information Science and Technology Agency, Government of BC, Canada
[2]Pacific Forestry Centre, Natural Resources Canada
[3]Department of Computer Science, University of Victoria
[4]Department of Geography, University of Victoria
4000 Seymour Place, Victoria, BC, Canada, V8X 4S8
(250-387-4803, office; 250-387-8419, fax); graham.j.hicks@gems1.gov.bc.ca

ABSTRACT – Monitoring Canada's forests with remote sensing necessitates the handling and distributed processing of vast quantities of remotely sensed data. A new project (EOSD) has been designed to create national forest products for Canada [1]. Creation of the EOSD products will require the use of Landsat and Radarsat data, topographic GIS files, historical forest cover GIS files, and field plot ground data. The GIS and remote sensing data sources are distributed across Canada and the United States. During the next five years, there will be large increases in the volume of remotely sensed data.

Canada through Canarie, a federally funded agency, created high-speed electronic networks, CA*net2 and recently CA*net 3, for the rapid movement of large volumes of data. As a CA*net 2 activity, we created a capacity for moving large volumes of imagery over ATM and high-speed Ethernet networks. A terabyte of remote sensing data held within a robotic data store has been accessible from Pacific Forestry Centre over this network. In this paper we present the structure of the high-speed network for the EOSD project. This network links data stores at five research centers and the headquarters of Natural Resources Canada. It also connects to some provincial information systems and to numerous advanced networks of the United States. This latter connection will facilitate the transfer of EO-1 and Terra data from NASA. Issues related to distributed processing, meta data, and security are discussed.

INTRODUCTION

A new project has begun to monitor the forests of Canada using remote sensing. This project is called Earth Observation for Sustainable Development of Forests (EOSD) [1]. In addition to indicators of sustainable development and attributes of national forest inventory, EOSD will report on products related to the Kyoto Protocol, such as reforestation, afforestation, deforestation, and above-ground carbon distribution across Canada.

Canada has 10% of the world's forests or 418 million hectares and 20% of the world's fresh water. The federal government is responsible for providing reports for international protocols and the United Nations. Several countries have announced that they will use remote sensing to monitor the compliance of other nations with international agreements. Within Canada, provincial governments manage the forests within their boundaries. For effective national monitoring, the EOSD project will draw on the historical descriptions of forest and land information from the provincial ministries. These agencies are distributed across Canada, but have large holdings of GIS and historical remotely sensed data.

Research and EOSD processing will be distributed across Canada and managed out of 5 forest research centers: Pacific Forestry Centre (Victoria, BC), Northern Forestry Centre (Edmonton, AB), Great Lakes Forestry Centre (Sault Ste. Marie, ON), Laurentian Forestry Centre (Laval, QB), and Atlantic Forestry Centre (Fredericton, NB). The Headquarters for the Canadian Forestry Service (CFS), Natural Resources Canada, in Ottawa will also be connected to the EOSD system. In addition to provincial and federal data sources, Canada has enjoyed close cooperation with US remote sensing and forest agencies. It is anticipated that the EOSD project will use and share data with NASA, NOAA, EDC, USFS, and major US forest monitoring projects. Thus, the GIS and remote sensing data sources for the EOSD project are distributed across Canada and the United States.

In the early stages of CA*net2 network development, a partnership was formed between the Pacific Forestry Centre (PFC) and BC Systems Corporation, now restructured as part of the Information Science and Technology Agency (ISTA). The PFC-ISTA team created a capacity in 1993 for moving large volumes of imagery over ATM and high-speed Ethernet networks. A terabyte of remote sensing data held within a robotic data store has been accessible from Pacific Forestry Centre over this network. In this paper we present the structure of the high-speed network for the EOSD project, which draws upon the new CA*net3 network with transfers up to 2.5 gigabits per second (Gb/s). This network links data stores at five research centers and the headquarters of Natural Resources Canada. It also connects to some provincial information systems and to many of the advanced R&D networks of the United States. This latter connection will facilitate the transfer of EO-1 and Terra data from NASA and interactive video dialogs.

NETWORKING BACKGROUND

The current North American broadband R&D networks are now reviewed. The networks are all continental in scale, are internationally connected and built on state-of-the-art

0-7803-6359-0/00/$10.00

transmission technologies. This means necessarily fiber optic cables with, often, multiple wavelengths of transmission. Bit rates supported are typically OC192 (10Gb/s), OC48 (2.5Gb/s) or OC12 (622Mb/s), on each of, currently, up to 32 wavelengths; systems carrying hundreds of wavelengths are being developed.

In Canada the Federal Government has funded the 'Canarie' initiative with increasing budgets since 1993. The mandate has been to support networking R&D and to stimulate high-performance applications development. The third generation of the Canadian backbone network CA*net3 was completed in 1999. The design employs Dense Wavelength Division Multiplexing (DWDM). Out of a set of 8 wavelengths initially deployed by the carrier supplier, Canarie has access to 2 with options on a further 6. The initial configuration uses the two wavelengths at OC48, intersecting each province at interchange points or gigapops. The topology is shown in Figure 1. An Acceptable Use Policy (AUP) and subsidiary networks within each province termed ORANs, or Optical Regional Advanced Networks govern access to CA*net3. User communities such as universities, colleges, and research labs are attached via an ORAN infrastructure.

CA*net3 represents a state-of-the-art deployment of a multi-wavelength, packet-over-Sonet network. To facilitate diversity and also to offer transit and peering with other networks, an early effort was made to interconnect with other international networks. Consequently, CA*net3 and its precursor network CA*net2 have been connected to the University of Illinois' STARTAP interchange point in Chicago, a meet-me exchange for most countries' advanced R&D networks. This exchange point has provided a highly effective meeting point for some 11 international and 7 national US networks. CA*net3 presently connects via an OC3 link.

In the US, the major advanced networks are the VBNS (NSF), Abilene (UCAID/I2), ESnet (DOE), Dren (DOD), NTON/SuperNet (DARPA/NGI) plus two, highly relevant to remote sensing, namely, NASA's NISN and NREN. These networks all conform to a broadband model in that they employ either Asynchronous Transfer Mode (ATM) at OC3, 12 or 48 over a Sonet core, or for the more advanced cases (e.g. NTON) a combination of Packet-Over-Sonet (POS) and ATM over OC192 Sonet on a specific wavelength. These latter cases, like CA*net3, are either already operating multiple wavelengths typically each at a 10Gb/s rate, or are multi-wavelength ready.

With respect to the needs of EOSD, the salient networks within Canada are CA*net3 and the ORANs. Within the US, NASA's NISN and NREN are most relevant. The NISN provides the foundational network for regular production interchange between NASA sites and agencies and NREN is part of the NGI initiative.

Present state-of-the-art networking in the US is typified by the NGI test bed NTON and SuperNet networks. Currently only operating on the west coast, NTON-II is an OC48 and OC192 Sonet based network operating on a single wavelength, but ready to deploy up to 8 or 16 wavelengths shortly. SuperNet is a national NGI affiliation of many advanced networks.

A common limitation has been the inability of typical end-systems (workstations in the $10K to $30K range) to generate challenging network traffic using conventional applications and protocols. In part, this is due to typical processor and input/output bus designs and to historical antecedents that (used to) dictate applications be coded assuming shared 10Mb/s Ethernet as the communications standard.

An aspect that this paper addresses is the relative disparity between current network performance and the customary style of processing used by the remote sensing community. This is described in a later section and is exemplified by Figure 3, an example process flow-chart, showing processing step delays and data volumes. This imbalance issue is currently a perceived barrier to deploying high performance networks into the realm of remote sensing.

NETWORKING TECHNOLOGY.

The dominant architecture for broadband networks in use today is still based on Sonet as the base physical layer with either ATM, or POS at layer 2. Internet Protocol (IP) is then carried over either dedicated Sonet containers or ATM Virtual Circuits (VCs). The inherent scalability and traffic management functionality of ATM has ensured its place as a layer 2 framing and switching technology providing transport from 25 Mb/s to 2.5Gb/s [2]. Most commonly, ATM trunk links are running at OC3 or OC12, with recent introduction of some OC48 capability. Recent incorporation of WDM techniques has allowed ATM to be bypassed and the communications (protocols) stack to be considerably simplified as a consequence.

CA*net3 is one of the first large-scale continental networks employing DWDM at multi-Gb/s rates. DWDM has clear advantages as an intermediate step between ATM-over Sonet and photonic networking proper (wavelength routing); needed to provide sustained Gb/sec rates.

At present advanced networks are still deploying a mixture of TCP/IP carried over Sonet and ATM. The large investments in ATM infrastructure made by carriers over the last 8 years ensures this technology will be an effective base for many broadband networks in the next decade. Subsequently the various proposals and research trials employing C (Coarse) and/or DWDM with wavelength cross-connects, switching and routing will permit 10's of Gb/s per wavelength with much reduced complexity in the communications stack.

Broadband networks currently employ routers and switches capable of packet forwarding rates in the $N*10^6$ PPS (packets per second) range, N, 1 -> 3. These devices are now designed around a structured switching fabric as a backplane, rather than a shared bus, and achieve aggregate switching bandwidths in the multi gigabit per second range. Research emphasis has shifted to developing switch and router designs that aim for aggregate i/o rates in the multi-giga bit to Terabit range with forwarding packet rates in the 10^7 PPS and higher range [3]. CA*net3 currently employs Cisco Systems* GSR 12008 series routers which offer 16 ports each operating at

2.4Gb/s, and packet forwarding rates of around $1.2*10^6$ PPS [4]. Precise interpretation of the PPS and switching rates depends on uniform and standardized agreement on packet sizes, simplex or full/half duplex switching and treatment of overheads. Transit delays through multi-hop routes are a concern of the classical routing and forwarding paradigm of the present Internet, as well as most advanced broadband networks. The use of Multi-Protocol Label Switching (MPLS) is in its early stages as an alternative forwarding paradigm with an expectation that this will allow expedited forwarding in a manner similar to that offered earlier by ATM and MPOA. Additionally, MPLS offers means to design VPNs and to create paths that have assured high performance. There are differing views regarding whether it is better to reduce complexity, i.e. 'flatten' the networks to gain performance possibly at the expense of restricted reachability, or, to seek enriched connectivity and risk lowering performance.

The set of protocols commonly known as 'TCP/IP' dominate most Internets. There are non-IP examples, in use, such as proprietary protocols over Fiber Channel, used frequently in Storage Access Networks (SAN)s and remote back-up facilities. There are also cases where native ATM solely is used typically for high quality video streaming. This paper assumes that the TCP/IP suite (Versions in common usage are TCP (Reno), TCP (Tahoe) and TCP (Vegas)) is used to provide reliable, universal communications between EOSD systems. A well-researched attribute of long distance broadband networks is the 'LFN' (Long Fat Network) phenomena well described in [5]. Essentially the TCP window, flow control, scheme behaves in a complex way under conditions of: a) large transit delays, and b) high bandwidth and congestion or packet/cell loss. Maintaining optimal (large bandwidth) link utilization has been shown to be very difficult to impossible with present realizations of TCP design and a non-zero probability of packet loss.

As a consequence it is difficult to determine, a priori, the expected throughput say between PFC and the University of Ottawa, over CA*net3 under unknown congestion conditions. The CA*net3 Network (West coast to East coast spans some 6400 Km) and with an estimated bottleneck bandwidth limit of say 1 Gb/sec, well known TCP parameter tuning is necessary to ensure throughput is maintained should there be packet or cell loss experienced [6], [7]. With incorrect buffer sizing and transport level timer values, under conditions of high error rate throughput can be reduced to close to zero due to repetitive cycles of slow-start and backoff caused by packet or cell loss.

Transit delays are attributable to three basic phenomena. The first is the clocking delay to transmit a frame, the second is propagation delay, reduced to a factor*c*L, with c = the speed of light, L = the cable run length and the third delay is switching or forwarding delays in network elements. Depending where an actual measurement is to be made, delays in the end systems also may be significant. Propagation delay through single mode fiber is usually taken as 0.7*c or 5 µsecs per km. Typically layer 1 amplifiers, repeaters and multiplexers will have fixed, deterministic delays in the nsec to low µsecs range. ATM switches will switch cells in 3 to 6 µsecs and may have queuing delays that are variable. Current frame switches and routers will forward IP packets typically in 10's to 100's of µsecs and high performance routers currently will forward frames in < 10 µsecs; both devices also subject to queuing delays. The CA*net3 routers offer forwarding delays on the order of 10 to 20 µsecs when operating at 50% of rated capacity [4].

For the CA*net3 network, (Figure 1), delays are due to: transmitting, + propagation, + DWDM devices + routers. For cases where links are running at 2.5Gb/s the delay cross-country is dominated by the propagation delay. Approximately, this would be 32 millisecs for a 6400 Km cross Canada path. Assuming a lightly loaded situation, West to East transit delays through the 10 routing hops should be of the order of 40 millisecs.

CA*net3 network topology and routing considerations are based on static wavelength assignments, Border Gateway Protocol (BGP) 4 as the Exterior Gateway Protocol (EGP) and Open Shortest Path First (OSPF) as the interior routing protocol. The network is built as a 'long pipe' with the ORANs attached via spurs usually using GigE. This appears as a series of links terminated on routers. Traffic inserted on the West Coast destined for say a University on the east coast will be forwarded through one routing hop per ORAN point of present or gigapop. A sample trace route listing currently shows 11 hops end to end. BGP is used as a policy tool enabling route filtering and denial of access to unauthorized source stations, and to establish peering sessions with adjacent R&D networks. Current Canarie policy is that the network is available to any educational institution as well as federal and provincial research agencies and labs. Recent adjustments to this AUP have included connections to schools also. All sites that connect must also maintain a separate link for commodity Internet access. Commercial use is denied unless legitimized by being part of a public sector or educational or research project.

Local access is often the Achilles heel of high performance networking since, frequently, it takes a relatively large cost to deliver fiber to an arbitrarily placed end system. For cases such as an University campus where fiber has been installed universally, connecting such campuses to backbones is relatively easy since often carriers have pre-deployed fiber to such institutions. Smaller or less well-situated organizations may incur considerable expense to join such a network. For CA*net3 local access the critical element is the development of an ORAN in each province. At the time of writing each province is involved at various stages in this process. The likely result of these efforts will be ORAN structures that will connect at Gigabit rates and local ORAN throughput performance is expected to approach sustained 10's to 100's of Mb/sec. For the near term it is anticipated that throughput performance limits will be set by end system software, and I/O channels rather than network infrastructure.

The expected scope and coverage of EOSD, see Figure 2, is largely confined to provincial governments, Universities, Federal Laboratories and related contractors. The effect of distributing resultant products to the general public will be described in a future paper. Previous projects and demonstration activities in Canadian remote sensing image processing and distribution at PFC have meant that several relevant sites have fiber installed directly to end systems. For the purposes here we assume institutions that become part of EOSD are already Giga-bit ready within their establishments.

End system effective performance has been the most common limitation in achieving very high throughput for many years. Performance of routers and switches show doubling in PPS rates every 6 to 12 months, Ethernet frame i/o rates from a typical Unix workstation, however, have not kept pace with network performance growth. Table 1 shows expected transfer rates for specific network bandwidths using ftp, overhead included. This is a theoretical limit and not likely to be achievable when end system behavior is included. For example, an 8.9 MByte file transfer (FTP) between, OC3 connected, UltraSparc 10 and Sparc 5 over a local, unloaded ATM Elan exhibited an effective throughput of 22Mb/s. Tests over long distances (e.g. Victoria BC to NUS, Singapore) completed in 1999 were unable to sustain throughputs above 3Mb/sec. due to the 'LFN' problem. Conversely, as a measure of what is feasible, Internet 2 in the USA recently achieved a new 'Land Speed Record' by moving 8.4 GB of data across the US, Redmond to Washington, DC, in 81 secs, or equivalently at an 830Mb/s rate.

From the foregoing, network capacity will be substantially more than that required to service any one image transfer for EOSD. Downloading a Landsat 7 source image over CA*net3 should be possible in 20 to 40 secs assuming a sustained effective throughput of 100Mb/s can be attained.

NATIONAL MONITORING COMMUNICATION REQUIREMENTS

Remote sensing offers information at spatial and temporal scales to monitor Canada's 418 million ha of forests that cannot be realized by using any other means. The Landsat series of satellites have created a rich legacy of remotely sensed data over Canada since 1972; an historical record that holds the key to assessing landcover changes.

Using Landsat TM (Landsat and Radarsat are principal sensors for EOSD), a single coverage of Canada requires 300 GB of data (unprocessed) corresponding to 750 scenes. A buffer of 30% to maximize the chances of a cloud-free coverage implies that 1000 TM scenes are required to cover Canada adequately. Processing such a data set for forest attributes would generate at least 10 intermediate scene products for every scene of TM data, amounting to over 10,000 scenes upon completion of data processing. Assuming that the data are located at several data warehouses across Canada and the USA, the only viable option for data interchange is through electronic means utilizing a fast network.

Consider a scenario where it takes a team of analysts at various centers, working collectively, one week to analyze a TM scene. At 155 Mbits/s speed, it would take just over 20 seconds to transport a scene across the network. Compare this to 1-2 days to exchange data using tape or CD media, not counting the time to upload and download the data from media, as well delays posed by different time zones in Canada. Assuming that 2 days were used in exchanging 4 TM scenes (which would fit nicely on a 2 GB data cartridge), this would mean that at least 250 data cartridges with unprocessed data, would need to be exchanged amongst various centers and groups. Assuming that there are 125 batches of tapes (2 tapes per batch) that would be exchanged, this would mean a delay of 125 to 250 working days.

Now consider a dedicated professional or a consulting company working on data analysis at a rate of $1000 per day. The following simple model helps to explain the enormous costs associated with data not being available on demand: Cost = Rate * (T + tu + td + D) where Rate = Hourly charges by dedicated professional, T = difference in time zone in hours, tu = time to upload data to medium on at site supplying the data (e.g. one hour), td = time to download data from medium at site using the data (one hour is typical), D = delay in receiving data due to transportation of media (1.5 to 2 days or up to 15 working hours in Canada alone; higher if cross-border).

If we consider a delay of 2 days (15 hours), upload and download times of an hour each, and a time zone difference of 2 hours, the above equation yields a cost to the client of $237,500 for delays posed by awaiting 125 batches of TM data. These costs are only for a single coverage of Canada. However, at least two coverages (leaf-on and leaf-off) are required for a good separation of softwoods and hardwoods. Using the data from Figure 3, the manpower cost associated with using a 1 Gbits/s network would be just over $80.00 for the scenario presented here.

EOSD volume processing will commence in 2002. By that time we must have the communications network in place and optimized for rapid sharing of data and individual video conferencing and consultation on demand. The reporting dates for EOSD are: 2003 for Montreal Protocol; 2005 for progress to Kyoto; 2008 for Kyoto Protocol reference; 2010 for national forest inventory; 2012 for Kyoto Protocol.

PROPOSED NETWORK SOLUTION

Our proposed EOSD network solution is shown in Figure 3. This depicts an overlay network using IP addressing above CA*net3; i.e. each wavelength carries Packet-Over-Sonet at 2.5 Gb/s as the physical transport. We show a CA*net3 general schematic in Figure 2. A common ORAN format uses distributed GigE. ORAN transit capacity is likely to be in the multi-million bits/sec range. However, there are many unknowns, such as local access feeder link capacities (or, for example, where an institution relies on extant carrier ATM offerings such as a VP service that has explicit traffic contract limitations). Not all ORAN-attached organizations will be Gig-E capable.

Local ORAN services are the domain of each ORAN administration and may, or may not support the suite of CA*net3 services. ORAN infrastructure within BC will likely be based on GigE accesses to a DWDM backbone traversing key institutions within the major urban areas of the province. Current planning makes use of dark fiber construction in large urban centers and shared access to wavelengths on long distance DWDM carrier facilities.

Common services that are necessary and must be supported are the border gateway routing protocols, and the interior routing protocols within each ORAN. Current CA*net3 design uses BGP4 for the EGP and OSPF for the IGP. Authorized access is mediated by the use of route tagging in BGP such that specific institutions are tagged according to their position in the CA*net community. EOSD is made up from Federal laboratories and Universities, all of which enjoy fully peered connectivity within the CA*net3 community as well as access to US R&D networks as shown in Figure 2.

The primary warehouse servers for EOSD will reside within PFC. Infrastructure being refined for this project is based on OC3 ATM using LAN Emulation and Fast Ethernet. Where necessary, GigE will be used for layer 3 switches and end systems directly connected to the ORAN, since it is expected that the present ATM-based structure will be discontinued. A detailed design of the EOSD network will be undertaken in the early stages of the project.

CONCLUSIONS AND FUTURE WORK

This paper presents a rationale for using high-speed communication networks in support of the management of very large quantities of remotely sensed imagery. Such networks have been used at the Pacific Forestry Centre for remote sensing data sharing and for demonstrations of video teleconferencing simultaneous with image analysis and GIS sessions. The paper also presents a network topology for the EOSD project, a project to monitor all of Canada's forests. Networks with speeds of 2.5 Gbits/s are recommended for distributed remote sensing of large image volumes. The Canadian network, CA*net3, is connected to the US STARTAP, opening the door to rapid data sharing between the EOSD project and US agencies, such as NASA, USFS, EDC, NOAA, DOD, etc.

The next steps in the implementation of the communication network for the EOSD project are the detailed design including the identification of distributed processing sites. The EOSD team will be brought together in 2000 for a workshop on communications, acceptable use policies, distributed processing, and intelligent data management.

ACKNOWLEDGMENTS

We are grateful for the support of the Canadian Space Agency, Natural Resources Canada, NSERC (DGG), and ISTA, Government of BC. We thank Howard Plato of PFC for early contributions to the communications design.

REFERENCES

[1] D. G. Goodenough, A. S. Bhogal, A. Dyk, R. Fournier, R. J. Hall, J. Iisaka, D. Leckie, J. E. Luther, S. Magnussen, O. Niemann, and W. M. Strome, "Earth Observation for Sustainable Development of Forests - A National Project," presented at International Geoscience and Remote Sensing Symposium 1999 (IGARSS'99), Hamburg, Germany, 1999, vol: 4, pp. 2101-2104.

[2] Anonymous, "ATM-Forum BTD-PHY-2.4Gb-01-04," 1999.

[3] C. Partridge, P. P. Carvey, E. Burgess, I. Castineyra, T. Clarke, L. Graham, M. Hathaway, P. Herman, A. King, S. Kohalmi, T. Ma, J. Mcallen, T. Mendez, W. C. Milliken, R. Pettyjohn, J. Rokosz, J. Seeger, M. Sollins, S. Storch, B. Tober, G. D. Troxel, D. Waitzman, and S. Winterble, " A 50-Gb/s IP Router," *IEEE Transactions on Networking*, vol. 6, pp. 237, 1998.

[4] N. McKeown, "Fast Switched Backplane for a Gigabit Switched Router (A Cisco Systems white paper)," Department of Electrical Engineering, Stanford University, Palo Alto, CA, USA 1999.

[5] V. Jacobsen, R. Braden, and D. Borman, "RFC1323 TCP Extensions for High Performance," 1992.

[6] T. Lakshman and U. Madhow, "The Performance of TCP/IP for networks with high bandwidth-delay products and random loss," : Unknown, 1999.

[7] H. M. Chaskar, T. V. Lakshman, and U. Madhow, "TCP over Wireless with Link Level Error Control: Analysis and Design Methodology," *IEEE Transactions on Networking*, vol. 7, pp. 605, 1999.

Table 1. Network performance times for remote sensing image transfers.

Time in seconds	Impossible quality	Moderate	Good	Excellent
Network	Internet	E100	GE	C3 0C48
One image Bytes	At 50 Kb/s	100 Mb/s FD	1 Gb/s FD	2.5 Gb/s
4.00E+08 raw data	74200	37.10	3.71	1.48
1.00E+09 fused data	185500	92.75	9.28	3.71

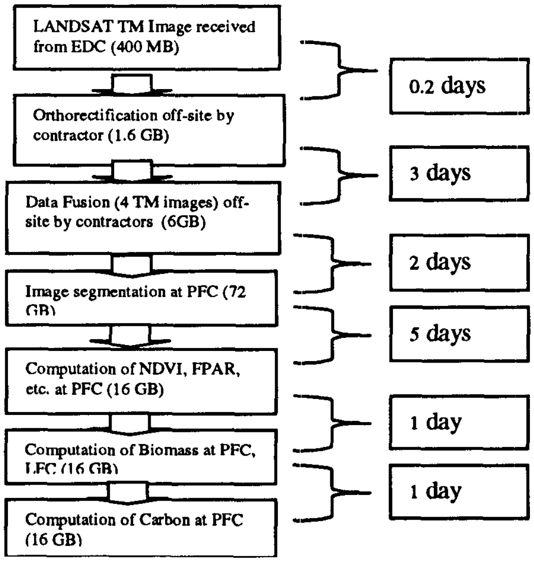

Figure 2: A Typical Application of TM DATA for Computing Above Ground Carbon

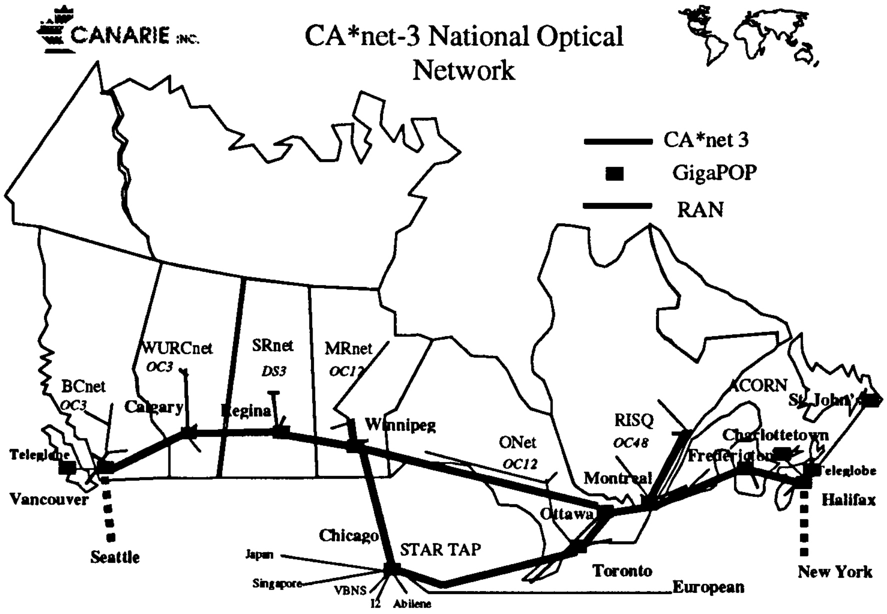

Figure 1

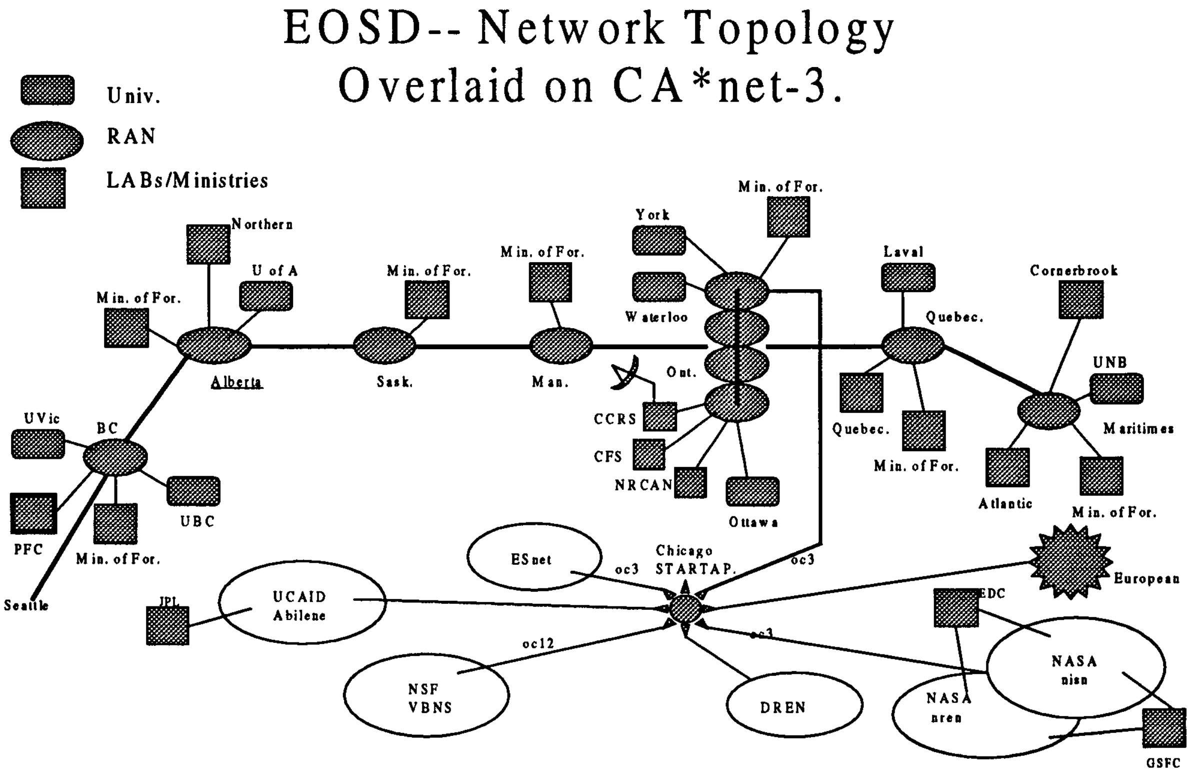

Figure 3