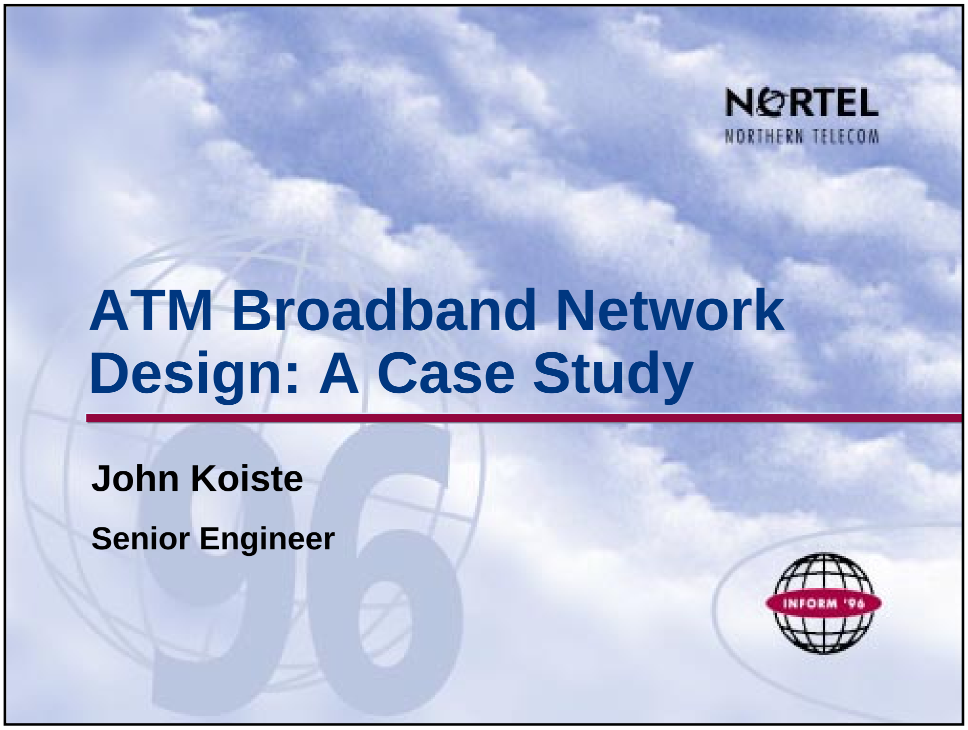

This presentation describes a useful layered model which has been called a Survivable ATM Network Management Architecture in the recent literature [2]. This architecture, which uses a broad meaning for the term network management, is a perceptive and elegant way of viewing an ATM network as it relates to the cell layer, the call layer, the virtual path layer and to the top-most layer: the facility network layer. That upper layer is where network design belongs, and that is the primary focus of the rest of the presentation. The layering concept is presented up front because there tends to be confusion about how different layers interact. The major part of the presentation is devoted to a case study on ATM network topology design. Three different design approaches are used to design 12 different topologies all contending to be the winner under rules set by “the customer”—a realistic (but fictitious) voice, data and video U.S. service provider.

About the presenter:

John Koiste has been associated with the Magellan product line for more than 14 years, first in R&D at BNR, then in product management, marketing, technical consulting, competitive analysis, and network engineering at Nortel.

The last four years at Nortel have been spent working on network designs with Nortel engineers and customers around the world, and designing computer-aided tools to make Magellan network engineering easier, faster and more accurate.

Prior experience with three other companies includes the design of aircraft radar, navigation computers, local area networks, special-purpose small computers, and CAD/CAM. Koiste was a guest lecturer at McGill University in Montreal for several years, teaching second-level digital and analog electronics courses to electrical engineering undergraduates.

Agenda

- **Where network design fits into ATM**
  - overview of a survivable network management architecture
  - meeting QOS at architectural layers
- **A case study of network design**
  - input provided by customer
  - how the best topology is to be selected
  - the three design approaches used
  - resulting 12 topologies, cost, performance
  - selection of best topology
  - generalizations drawn from case studies

2

Network design follows the same general principles regardless of the type of network being considered but there are major differences in the details—for example between synchronous transfer mode (STM) and asynchronous transfer mode (ATM) networks.

This presentation starts off by describing where network design fits into the ATM concept. It then presents an overview of a survivable network management architecture, based on interesting and valuable articles in the current literature. Using that architectural model, network design can be viewed as a facility network planning activity for one of the management layers. The concept of meeting quality of service (QOS) for the layers is reviewed next because it helps to tie the architectural concept together nicely.

This is followed by an ATM network design case study with the headings as shown above.

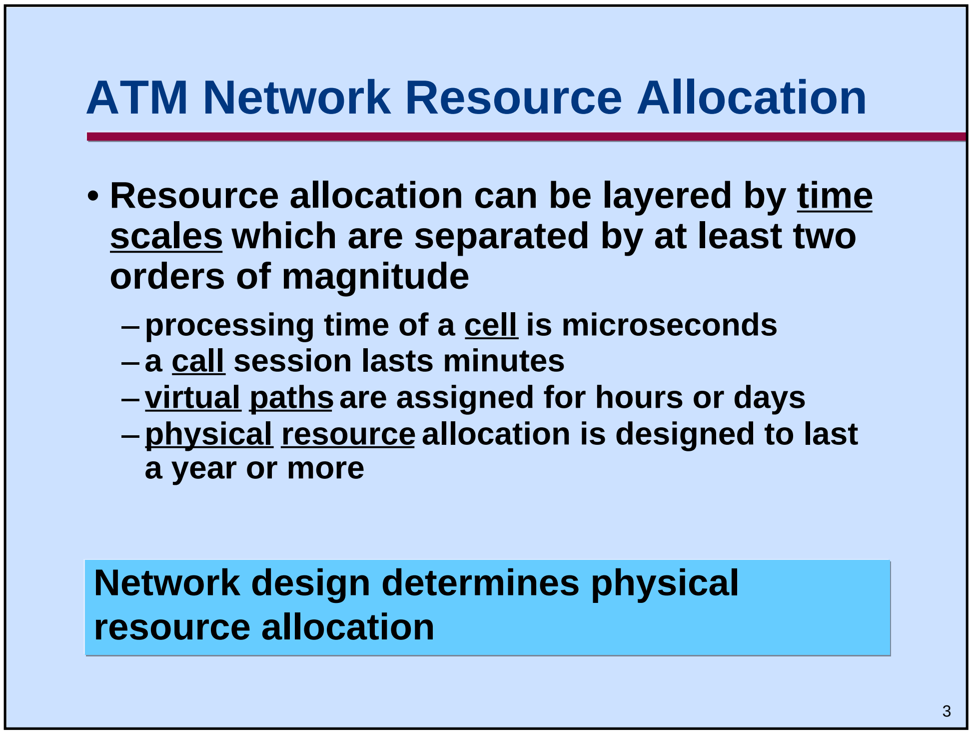

The concept of ATM resource allocation has been covered in the literature over the last few years, and a succinct summary is provided in [2].

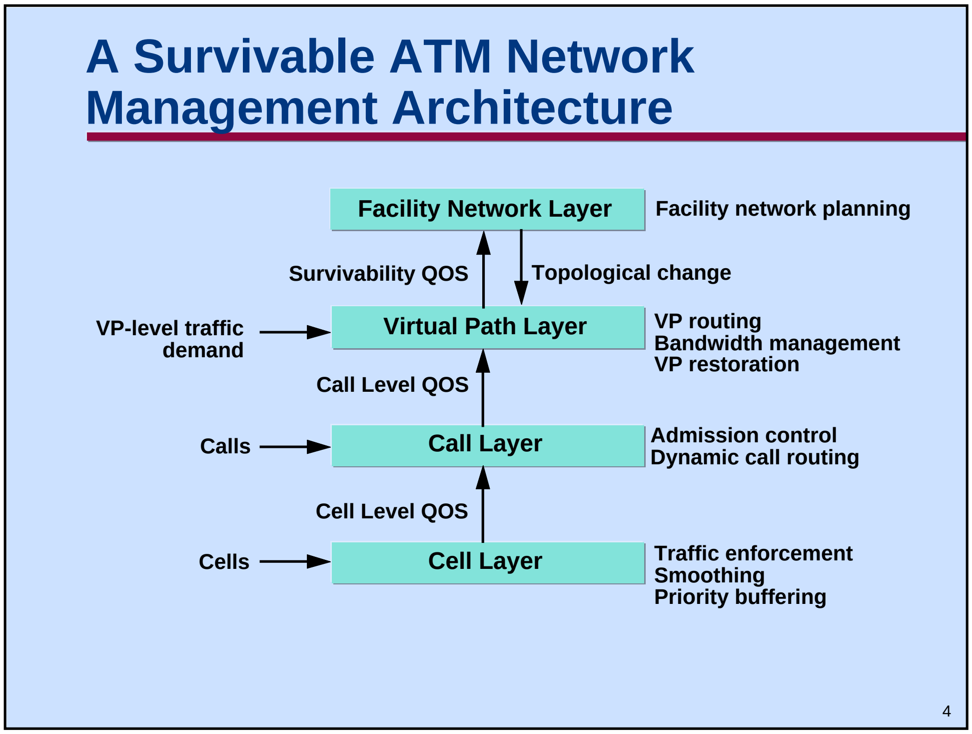

Two of the reference papers, [1] and [2], explain the network management layering concept, and the diagram above from [2] is a useful summary. The purpose of the layering is to simplify the network management process, which, in its broad sense, includes everything on the chart above.

The highest layer is the facility network layer, and it is the focus of this presentation because it includes network design. It is important to be constantly aware of the other (non-physical) layers also, however, because it is all too easy to inadvertently use capabilities in any layer to try to solve problems that are outside its scope. Because of the large difference in operating time scales between layers, the transmission resource configuration at a higher layer can be regarded as quasi-static.

Resource assignment in a layer is performed so that the quality of service (QOS) of the next lower layer is guaranteed. This is accomplished by assigning sufficient bandwidth at the layer. For example, a program which might generically be called the VP-layer manager builds a VP subnetwork over the physical network resources that have been provided by the network designer. Given the VP-level demand that was used to design the physical layer, and which is adequate for the VP manager to build the required subnetwork, the call-level QOS can be met.

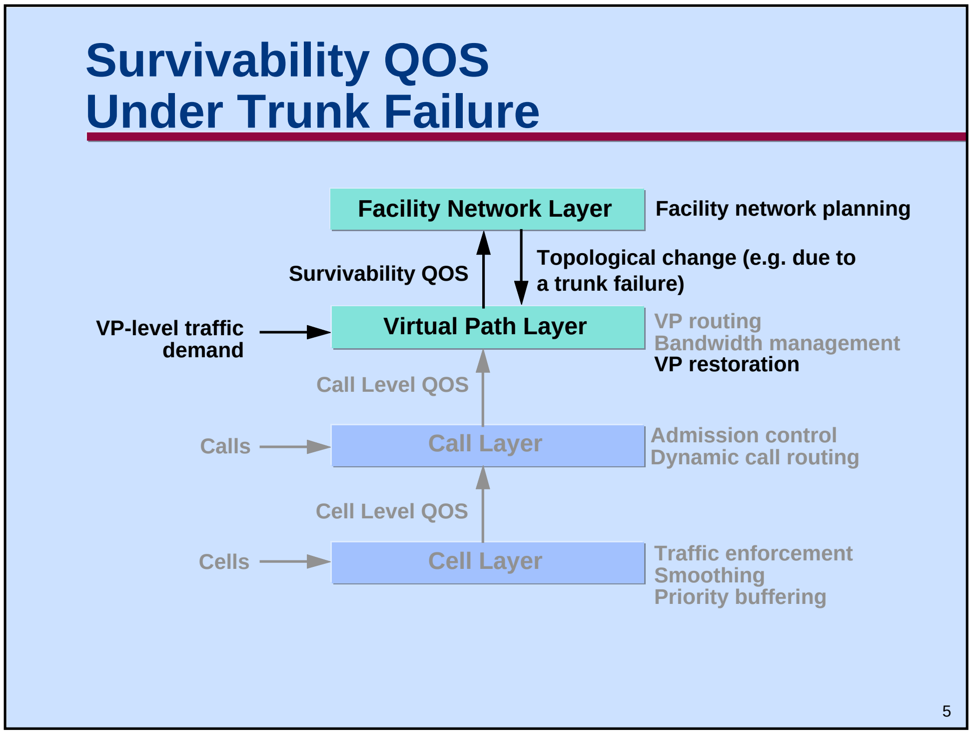

The survivability functions of the network, which are among the most difficult aspects of network design, VP routing design, bandwidth management capability, and fast VP restoration design are all in the upper two layers. The concept of survivability quality of service (SQOS) is introduced in [2].

If the VP manager cannot maintain an SQOS at the desired level due to growth in VP-level traffic demand, the facility network layer must trigger a network planning process.

If an unplanned topological change occurs due to one trunk failure, the SQOS should be maintainable by the VP manager if the network was designed for that. With two overlapping trunk failures—a rare but possible event—the VP manager may have to ask for assistance from the facility layer again. Because of the slowness of the facility network layer to respond, VP layer may have no option but to force a lower QOS to the call layer.

Contrasts Between the Upper Two Layers

- Facility network layer
  - largely product independent
  - ultimate provider of network bandwidth
  - partly or largely under manual control
  - can only react slowly to trunk failures or requests for more resources
- Virtual path layer
  - potentially significant product dependency
  - must use given bandwidth effectively
  - totally under automatic control
  - must react quickly to trunk failures and to increases in call-level requirements

6

Note the contrasts between the upper two layers. The facility layer is a strategic layer when it comes to network resources while the VP layer is more tactical.

Role of Network Design in Facilities Layer

- **The activity in which network bandwidth resource placement is assigned to meet**
  - **projected traffic demand and QOS**
  - **network survivability requirements**
  - **lowest possible bandwidth resource cost**

7

This chart summarizes the role of network design in the facilities layer.

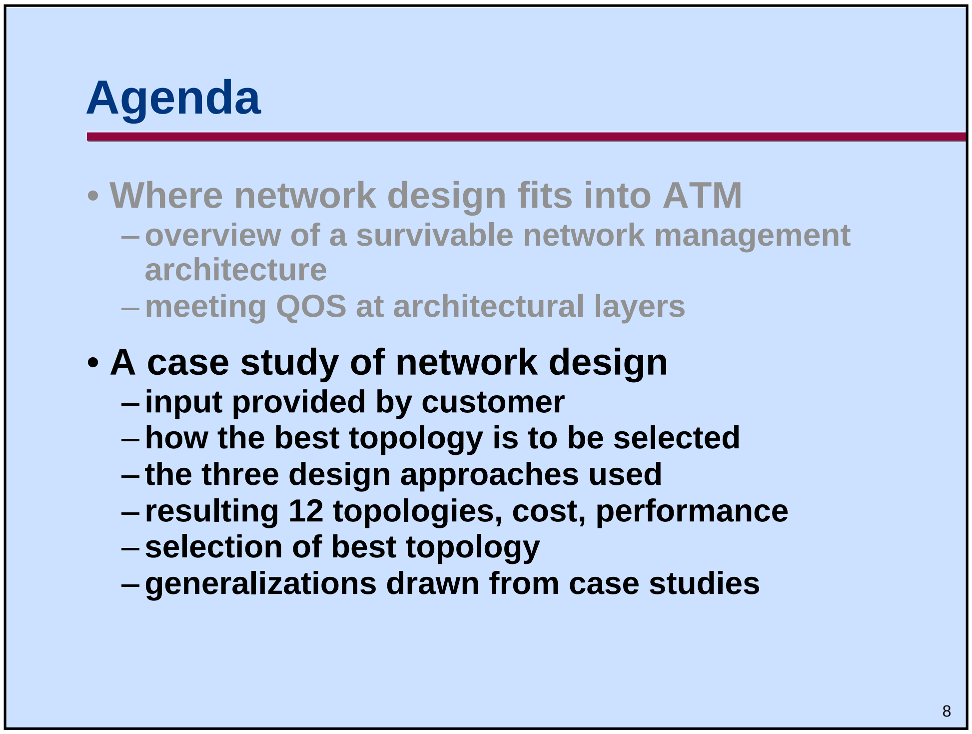

This section describes a case study which is not real but was designed to be realistic based on Nortel network design experience.

Case Study Background

- The company is a U.S. public telecommunications service provider of nation-wide voice, data and video services
- The company has 16 major locations being used in this case study, and a large number of smaller locations not being considered here
- The company has sent out a request for quotation of a 16-site ATM network design
- The bidders are asked to design their most cost-effective network using T3, OC-3 and OC-12 inter-nodal connections (trunks)

9

This is just a preamble to set the stage.

Network Design Information

- The network is to carry the company's customers' voice, data and video traffic between the 16 locations
- Traffic levels for busy-hour voice, data and video have been estimated and are provided for use in the design
- Costs to be used for T3, OC-3 and OC-12 are provided (next page)
- Each location is to have two or more connections to the network for reliability
- OC-12 trunks can be assumed fully protected by the SONET level
- Trunk utilizations are to be kept below 90% for any trunk when all trunks are functional
- The network must be able to carry 65% of busy period traffic in the event of a single trunk failure

Design a network topology with acceptable service and delay performance and lowest possible internodal trunk cost

10

The information provided is typical of what one might expect in a real request for quotation from a service provider.

The goal of the design process is the usual one: Meet a performance requirement at lowest cost. What are the performance requirements?

- The network must be able to carry the offered traffic
- Trunk utilizations must not exceed 90%
- The reliability requirement must be met and that probably means that utilizations of 90% are not possible with any likely topology. Also the protection of OC-12 by the SONET level for this case study means that at the ATM level any OC-12 trunk can be assumed to be totally reliable, which significantly reduces cost and network complexity.
- As in this case study, if the customer has not specified QOS, it is important for the network designer to know enough about services in general, as well as network design principles and methodology, to design for acceptable QOS.

Cost of Facilities to be Used in Design

Trunk	Distance	Base Rate	Cost/Mile*
T3	0—50 mi.	$ 10627.47	$ 91.59
	51—100 mi.	$ 11399.97	$ 76.14
	101—500 mi.	$ 12999.97	$ 60.14
	501+ mi.	$ 16864.97	$ 52.41
OC-3	0—50 mi.	$ 21554.95	$ 183.17
	51—100 mi.	$ 23099.94	$ 152.28
	101—500 mi.	$ 26299.94	$ 120.28
	501+ mi.	$ 34029.94	$ 104.83
OC-12	0—50 mi.	$ 54295.84	$ 586.15
	51—100 mi.	$ 59239.81	$ 487.28
	101—500 mi.	$ 69479.81	$ 384.88
	501+ mi.	$ 94215.81	$ 335.45

* For entire distance

11

By providing realistic costs for facilities, the customer makes possible the comparison of various submissions. The costs shown are gauged to be typical of U.S. tariffs at this time, but are not any service provider's actual tariffs. Note the four linear cost segments, which allow for a convenient separation of switching and long distance costs.

Locations of the 16 Node Sites

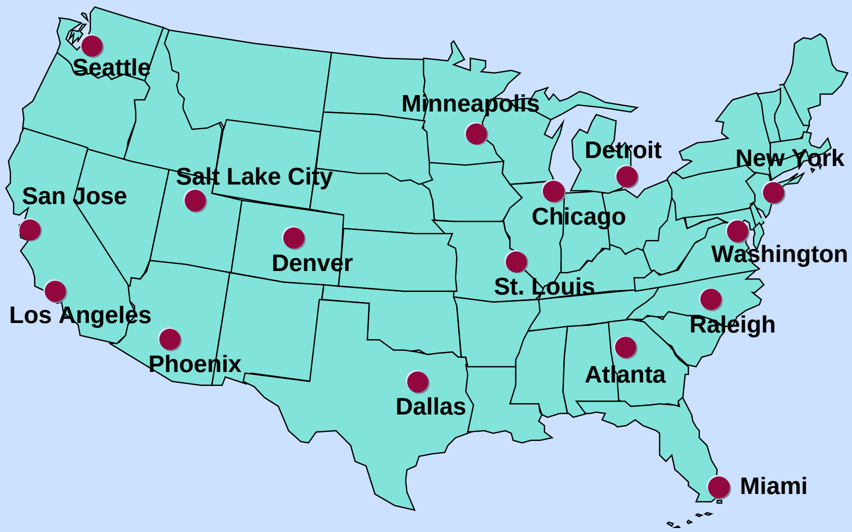

12

This just identifies the 16 sites.

Part of Offered Traffic Table Provided by Customer

- Multiply voice channels by 12
- Use 2% blocking factor

	3	4	5	6	7	8	9
8				---------------------- VBR ----------------------			
9							
10	City 1	City 2	Voice Chs	PC kbps	CC kbps	Video kbps	Total VBR kbps
11	Atlanta GA	Chicago IL	10	1000	1000	10000	12000
12	Atlanta GA	Dallas TX	40	1000	1000	10000	12000
13	Atlanta GA	Denver CO	10	1000	1000	10000	12000
14	Atlanta GA	Detroit MI	10	1000	1000	10000	12000
15	Atlanta GA	Los Angeles CA	25	1000	1000	10000	12000
16	Atlanta GA	Miami FL	100	20000	1000	10000	31000
17	Atlanta GA	Minneapolis MN	10	1000	1000	10000	12000
18	Atlanta GA	New York NY	30	25000	1000	10000	36000
19	Atlanta GA	Phoenix AZ	10	1000	1000	10000	12000
20	Atlanta GA	Raleigh NC	15	1000	1000	10000	12000
21	Atlanta GA	Salt Lake City UT	10	1000	1000	10000	12000
22	Atlanta GA	San Jose CA	10	1000	1000	10000	12000
23	Atlanta GA	Seattle WA	10	1000	1000	10000	12000
24	Atlanta GA	St. Louis MO	10	1000	1000	10000	12000
25	Atlanta GA	Washington DC	10	1000	1000	10000	12000
26	Chicago IL	Atlanta GA	10	1000	1000	20000	22000
27	Chicago IL	Dallas TX	60	1000	1000	20000	22000
28	Chicago IL	Denver CO	10	1000	2000	20000	23000
29	Chicago IL	Detroit MI	100	1000	2000	20000	23000

13

The offered traffic table is provided in a typical format: a spreadsheet. Another common format is a text file. In that case, translation to spreadsheet format is done using a utility program to change delimiters such as semicolons to tabs—a process which results in the format shown.

The first VBR column (PC kbit/s) refers to traffic which originates in personal computers. The second (CC) refers to traffic from communications controllers (e.g. 3745). The third is VBR video traffic, for example MPEG-2 coded video using AAL-5. All are in kbit/s and all are busy-period average values. Totals are shown in the last column.

Additional information is given for voice, which makes it possible for the network design program to calculate call-level traffic intensity in Erlangs, and the resulting CBR traffic level.

Note that the table is n(n-1) or 16(16-1) or 240 rows long, since there are 16 nodal sites. There is no requirement in this case to route traffic from any city to any other city more than one way; i.e., all services can use the same route through the topology, if the designer so chooses. (If the designer does not choose that approach, there has to be a good explanation of why not).

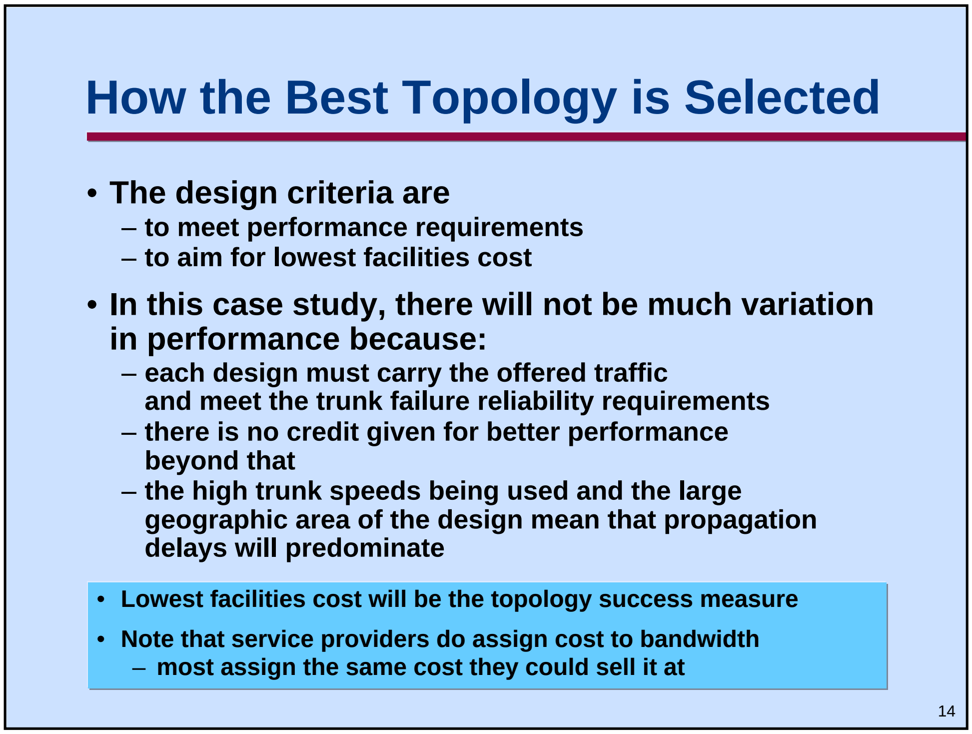

In a competitive bid or design exercise, it is obviously important to understand the rules of choosing the best design. Rules vary, but the ones above are typical. Note that service providers do assign costs to facilities even if they own them, since the facilities can be usually leased out for profit. Here the ATM network has to pay the same for the facilities as if they were leased out because the leasing profits are eliminated. ATM then has to deliver its profit on ATM services.

What is the Best Design Approach?

Three design approaches often used at Nortel are applied and compared here:

(1) Pruning approach: Start with a full mesh topology and remove connections
- a methodology is defined for sequence of connection deletion, and when to stop

(2) Additive approach: Start with zero connections and add connections
- a methodology is defined to determine sequence of connection addition, and on when to stop

(3) A four-node hub approach: Select a four-node hub, fully meshed, and add other connections
- a simple methodology is used to select the hubs and add connections

15

The three topology design approaches above are often used at Nortel. Each has its advantages and proponents. They come at the solution from three different directions, and this case study was judged to be an interesting one in which to compare them as part of the process.

Nortel Magellan Network Design

- **Nortel has skilled and experienced network designers in many parts of the world**
- **Nortel uses state-of-the-art programs for quick and accurate designs**
  - **some programs are internally developed, such as NetCalc2 used in this case study**
  - **others are externally procured, e.g. automatic topology generators**
- **The people and the programs are equipped to design networks with many hundreds nodes using various techniques**
- **Fast design iterations are possible if input requirements change or if the customer expands the network**

16

NetCalc2, used for this case study, is an internally created program used for network design by Nortel in many parts of the world. It is a computer-aided design program which assigns most of the clerical and computation processes to the computer but requires knowledge and direction of a network designer to design the network. NetCalc2 does not design a topology by itself, but provides fast feedback on how good any topology entered into it is, given the other inputs such as offered traffic.

The current crop of automatic topology generators would not speed up the design process for a network of the size and type being considered here, nor would they likely result in a significantly better design. But they are the way to go in the future; it's only a matter of time. To further improve design productivity, there are also other design areas to automate to a greater extent than what NetCalc2 does.

The speed of design for this case study is mentioned later.

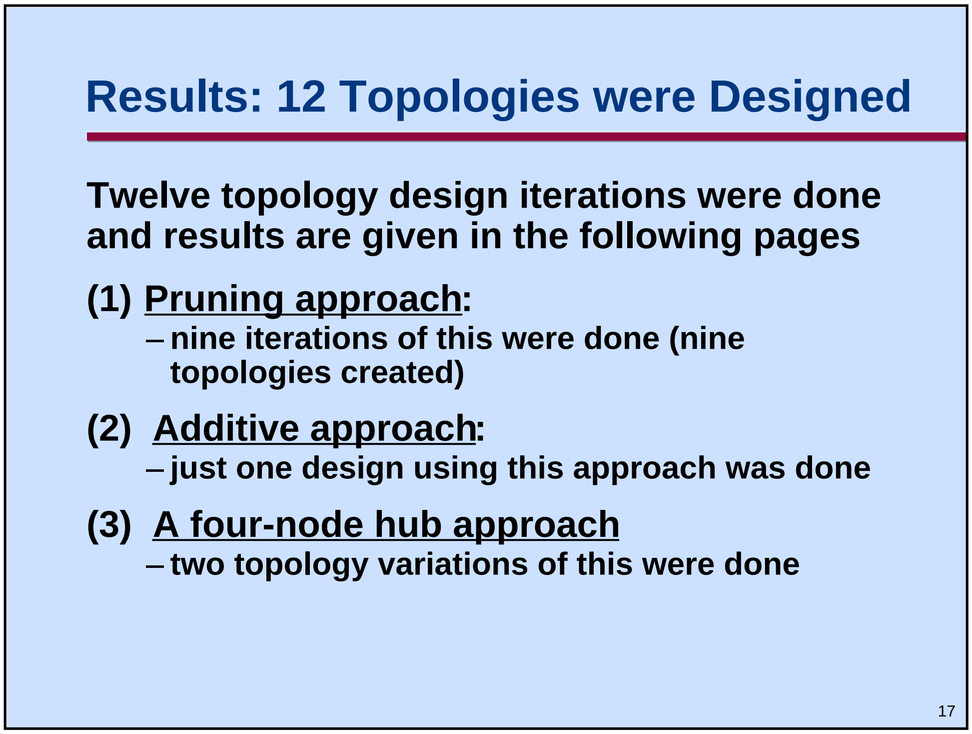

Twelve topologies were created for this case study, with the aim of finding the best one. The pruning method is fairly time consuming but once the up-front work is done, the iterations don't take very long.

To put some perspective on how long iterations took for this case study, all 12 designs were done in one day. They were also repeated twice because it was decided to change the input traffic profile and facilities costs after the first and second rounds were done. The whole design process took two days for one person using a desk-top machine comparable in speed to a 66-MHz 486 PC.

Sometimes creating a case study takes longer than doing the real thing. If Nortel had been designing to a bid, the topology design would have been done in one day or less. But if the input parameters are incomplete, the process can be slowed considerably while seeking clarification. Larger networks (e.g. 50 nodes) take more time because the amount of work for the computer and the person seems to increase as the square of the number of nodes. But dividing the network into hierarchical subnetworks reduces the work required. There are also other tricks of the trade which can significantly reduce the design effort for something like a 50-node design.

Pruning Design Approach

- Enter the offered traffic
- Create a full mesh topology (the 120-trunk topology)
- Create other topologies with fewer trunks by repeating the following:
  - delete connections from the current topology as follows:
    - remove connections starting with lowest traffic (in busier direction), longest connection first to resolve ties
    - stop every 10 or 15 connections to record all topology performance data and cost
    - never leave less than two connections to any node

Three of the nine topologies designed using the pruning method are shown here

18

The pruning approach is particularly instructive for new designers (or, as it turned out, to this presenter, who has much more often used the additive or multi-hub approach).

Note that zero traffic connections are eliminated easily—but this case study had none. So, next on the list are the lowest traffic connections. If there is a tie in traffic level for those (and there were many here) the longest connections are eliminated first because they are the most expensive per the costs used here (and that is very typical in real designs).

Where do you stop? Perhaps after every 10 or 15 connections to take a cost and performance reading. Then you keep going until costs stabilize or start to increase, or until you run out of trunks to prune due to the reliability requirement. Beware of local minimum cost situations.

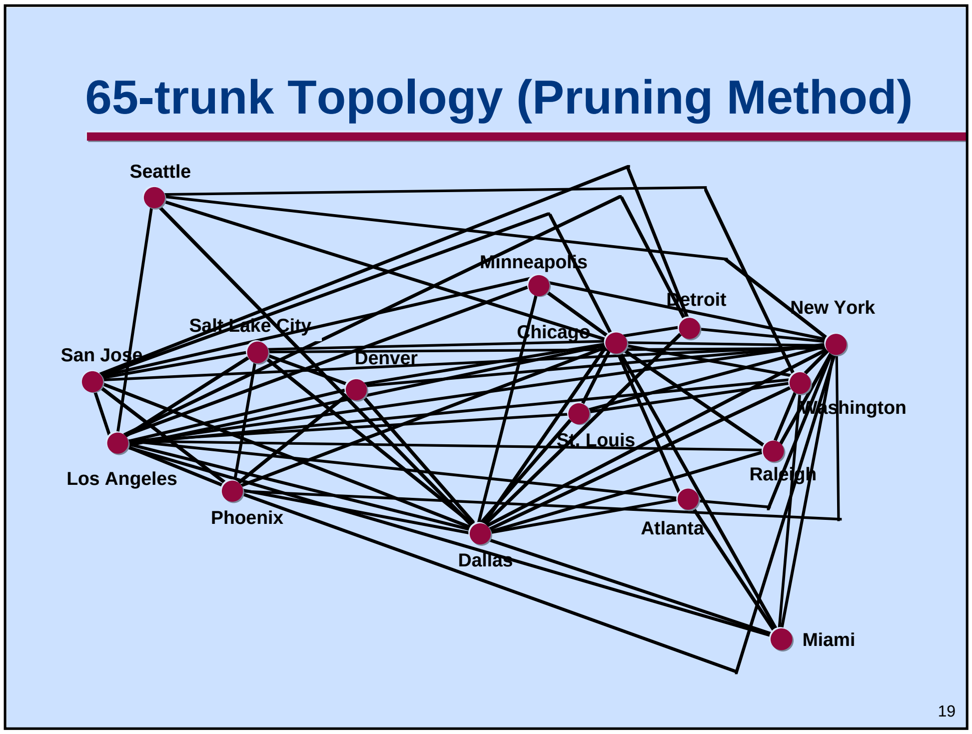

This is not an optimum topology in an age where high-bandwidth fiber is much cheaper than T3 on a per-kbit/s-basis. For the design above, there is actually a lot of fiber, but not enough yet. (Due to the low resolution of the picture, there was no attempt to identify fiber on it).

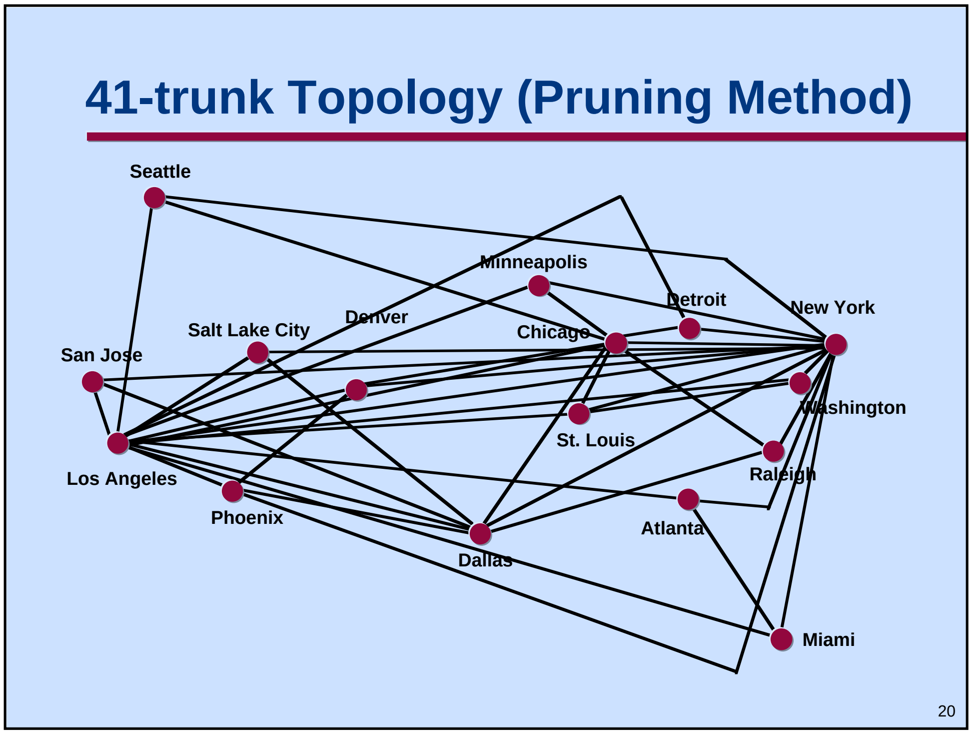

Before the case study tariffs were readjusted (based on more realistic ratios between T3, OC-3 and OC-12 rates, which in round numbers or net bandwidth numbers are 43 Mbit/s, 150 Mbit/s and 600 Mbit/s, respectively), this topology was the best one. Note the three-trunk reliability for the smaller sites, such as Seattle. If one trunk from Seattle fails, only 66% of the traffic can be carried compared to the same utilizations when all three are functional. But the customer allows (in this case) a network-wide traffic reduction of 35% (to 65% of nominal). So having three trunks to a node provides a simple method of protection without much risk if all failure scenarios are not examined (a time-consuming process).

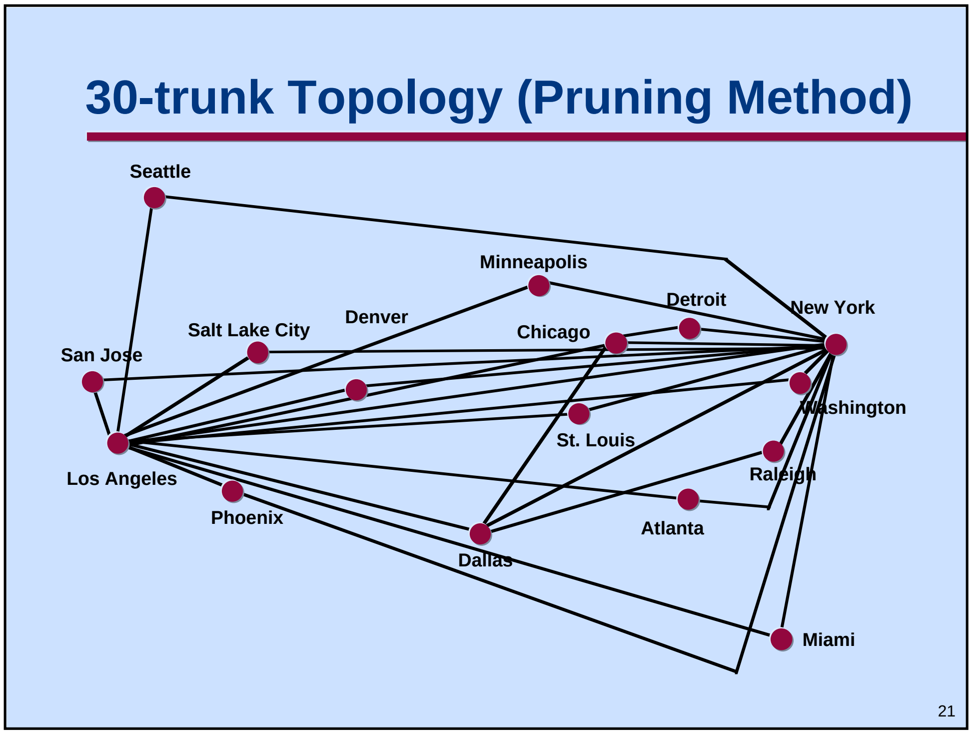

Here, failure scenarios must be done quite rigorously, because many nodes only have the minimum allowable two trunks. If one fails, only half the traffic carrying capacity remains, so those two connections must be tested for adequate bandwidth.

Note that OC-12 trunks, not identified in the picture, are deemed by the given data for this study to be perfectly reliable at the ATM level since they are protected at the SONET level. That makes a significant difference in cost for the topologies with fewer trunks (in their favor).

(Nortel’s Concorde ATM switch ratchets the reliability level one step higher: It uniquely offers such protective capability for any ATM network. Refer to Quyen Bo’s Inform ‘96 presentation “Magellan Concorde—Infrastructure Solutions”, particularly to the slide entitled High Performance: ATM and SONET/SDH Interworking).

Cost of Topologies Created with Pruning Method

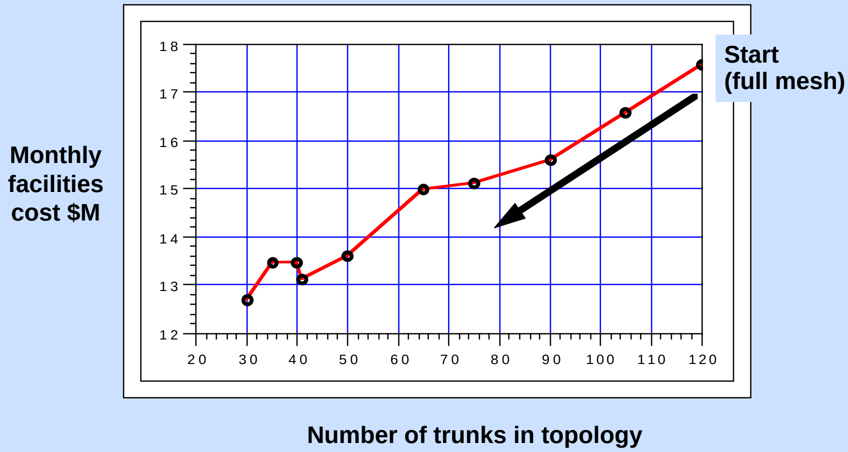

22

This is a plot of the facilities costs in $US per month for the nine topologies created using the pruning approach. The curve is not smooth, particularly at the low end, because when a trunk is removed, often one or more of the remaining trunks must undergo an increase in bandwidth to accommodate greater tandem traffic. For example, another trunk in the network might have to go from one OC-12 to two OC-12s, since there isn't anything in between. That is why there is a jump in cost in going from 41 to 40 trunks.

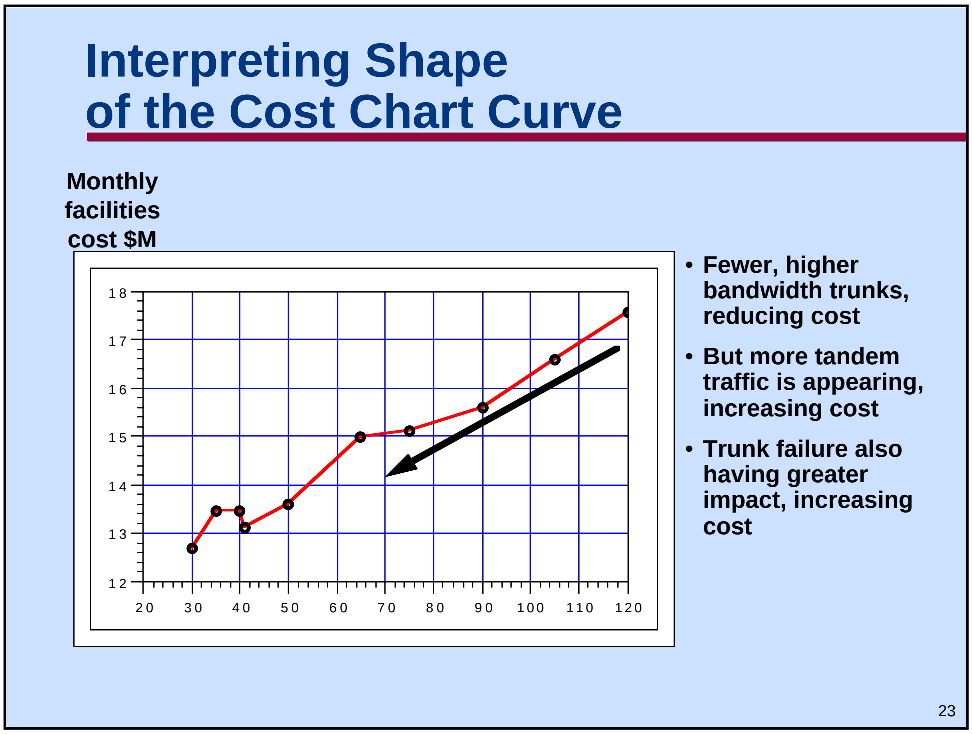

The three effects in the note on the right of the slide are important. Fewer trunks are better in terms of lower cost per kbit/s. But fewer trunks means more tandem traffic. And a more massive impact due to failure, which must be covered (but not for OC-12 in this particular case study). So often, depending on reliability requirements, the traffic profile, and tariffs, the curve bottoms out at some point and then starts to go back up as the number of trunks is further reduced.

Additive Design Approach

- Enter the offered traffic
- Identify high traffic paths
- Interconnect nodes as follows:
  - put in trunks for the highest-traffic paths
  - for any node with lower traffic levels,
    - consider putting in a trunk to the network node which has the greatest traffic from it (or to it)
    - otherwise, connect it to the nearest two network nodes
- Run the autorouter or load in routes from some other source
- Run through traffic calculations
- Do reliability checks by failing busiest trunks

24

The additive approach comes at the problem from the opposite direction compared to the pruning approach. Ideally, one should converge on exactly the same final result. But that doesn't happen as frequently as one might hope.

Where does one stop with this approach? Given time, when the cost has definitely started to increase again. Then one selects the topology with the lowest cost. With high bandwidth circuits, such as OC-12 for this example, there was a sense, based on past experience, that after 26 trunks had been added, the point of diminishing returns had been reached. So, partly due to time constraints, that 26-trunk effort ended up as the sole contender for this approach against the nine topologies done using the pruning approach.

26-trunk Topology (Additive Method)

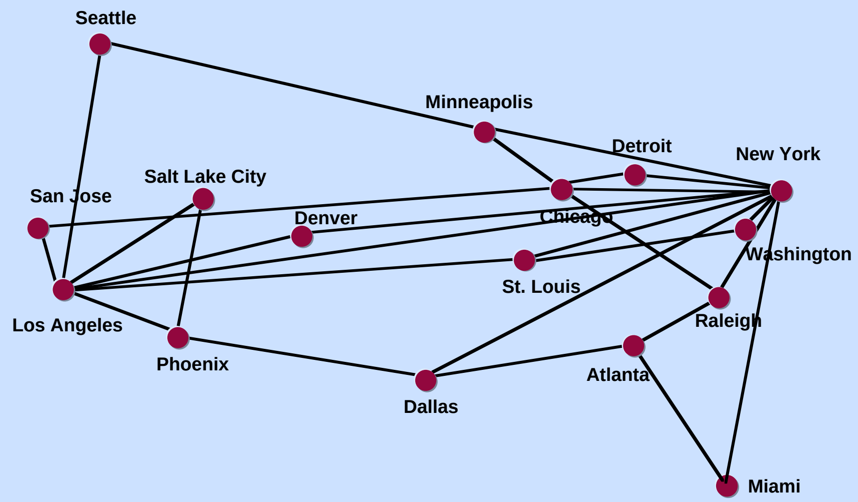

25

This shows the 26-trunk topology designed using the additive approach.

Cost of 26-trunk Topology Created with Additive Method

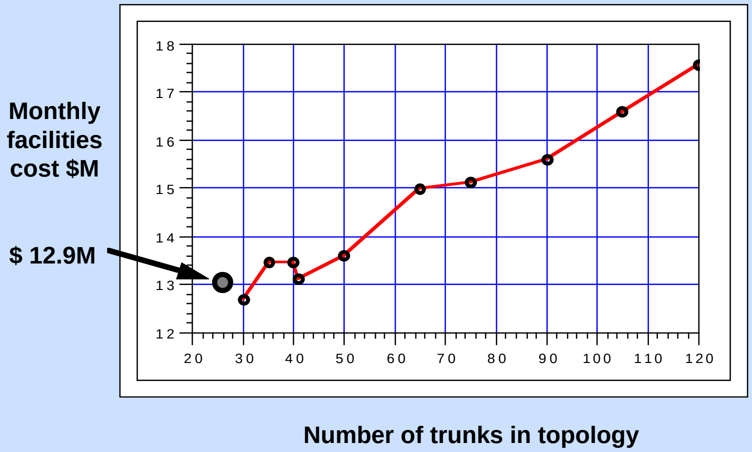

26

The one entrant from the additive direction did quite well as shown, ending up second cheapest of the 10 topologies done to this point.

A Four-node Hub Design Approach

- Enter the offered traffic
- Identify high traffic paths
- Interconnect nodes as follows:
  - identify four cities with the highest traffic
  - fully connect those four cities (hub cities) to each other
  - connect all other cities to nearest two hubs
- Run the autorouter or load in routes from some other source
- Run through traffic calculations
- Do reliability checks by failing busiest trunks

27

Several variations of the multi-hub design approach have worked well in network designs at Nortel. One can start with two hubs, and home all other nodes to those two nodes. Or one could use three hubs, or four or five, and connect other nodes to the nearest two hub nodes. As the number of hubs increases, tandem traffic increases on hub-to-hub connections. Often four hubs are the most cost-effective. With the hub approach, the hubs themselves are usually fully meshed.

This approach is a good cross-check on more complex approaches, such as the other two being used here—and sometimes it does even better.

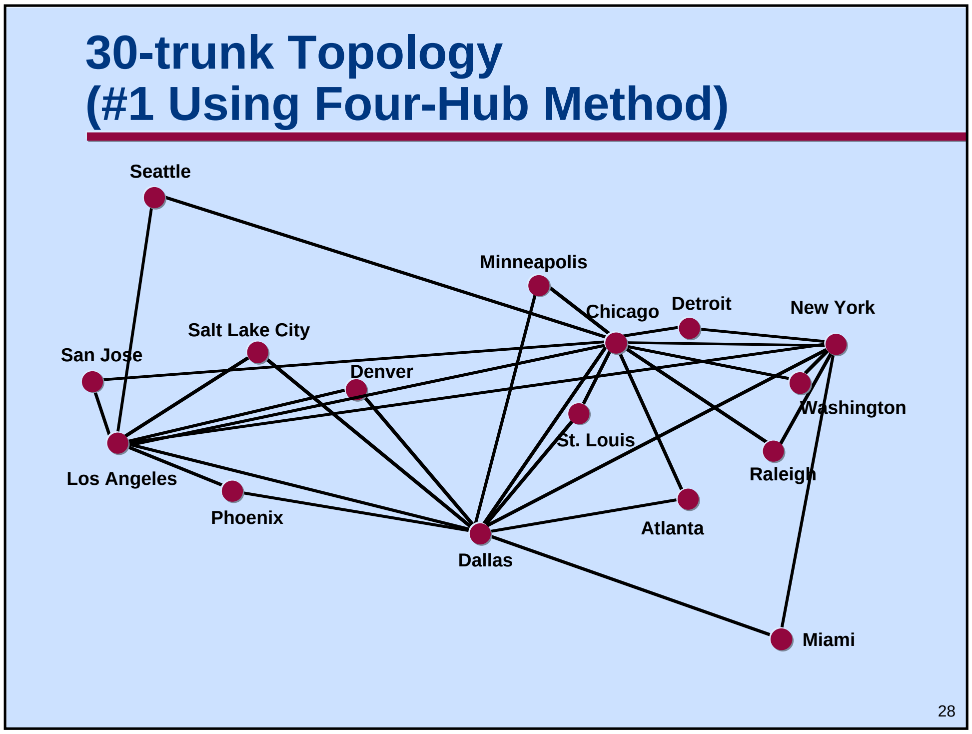

This is one of the two four-hub topologies done as part of the exercise. The hubs are Chicago, New York, Dallas and Los Angeles, which are directly connected to each other. The other cities are connected to the nearest two hubs, with one exception (San Jose).

Cost of 30-trunk Topology #1 Created with Four-hub Method

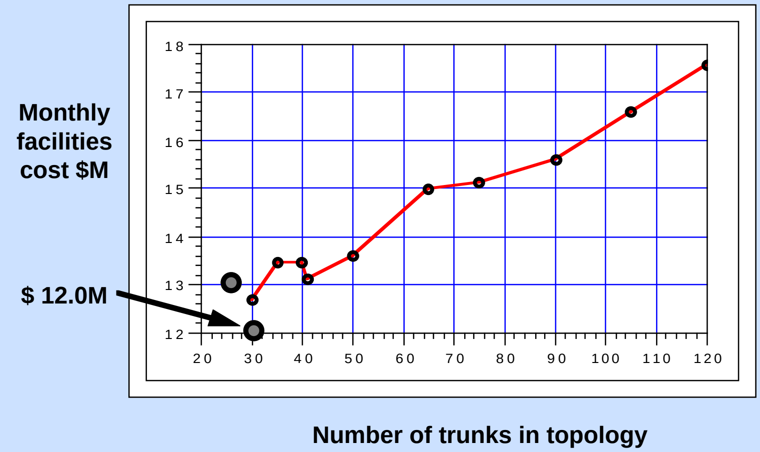

29

As has happened several times before in Nortel's design experience, a simple four-hub design prepared in a few minutes has seemingly come out of nowhere to win the contest. At the least, this is an approach one should use to keep other designs honest, because it can be done so easily.

Cost and Performance of the Twelve Topologies

Design Method	Trunks	Cost/Mo ($M)*	Delay (ms)**
Pruning	120	17.585	10.2
	105	16.574	10.2
	90	15.592	10.3
	75	15.130	10.5
	65	14.978	10.6
	50	13.607	11.3
	41	13.115	11.5
	40	13.460	11.6
	35	13.478	12.0
	30	12.692	12.3
Additive	26	12.931	13.2
Four-hub	30 (#1)	11.999	11.4 ← Best one
	30 (#2)	12.485	11.4

* Trunk cost per month

** Excluding access delays

30

This shows the facilities cost and delay performance of the 12 topologies. Note that delay across the network (using great-circle calculations for distance and therefore propagation delay) is largely determined by propagation delay at the high trunking speeds used here. Nodal delays, transmission (clocking) delays, and trunk queueing delays are minimal by comparison. The controlled utilizations on trunks in the designs are low enough to keep queueing delays very low and also low enough that QOS is not an issue at all. The networks have been designed to carry all the traffic specified, and other calculations, done as a spot-check, would indicate that the lower layers in the diagram on page 4 can meet any reasonable QOS—i.e. the facility network layer will have done its job, given the VP-level traffic demand as specified.

So, the winning topology is the one shown, since cost is the only issue of significance.

How did the 30-trunk Four-hub Topology #1 Win?

- The four cities with the greatest traffic levels to or from other cities are fully connected to each other
- The topology makes good use of the OC-12 tariff for long runs
  – a different tariff could easily result in a different winner
- The average hop count is reasonably low, and the maximum number of hops is only three
  – tandem traffic is still fairly low

The topology has achieved the best balance in cost among:

- Number of trunks (fewer bigger ones are cheaper)
- Tandem traffic (which increases trunk bandwidth and cost)
- Spare bandwidth for failure protection (which increases cost)

31

This chart summarizes how the winner won.

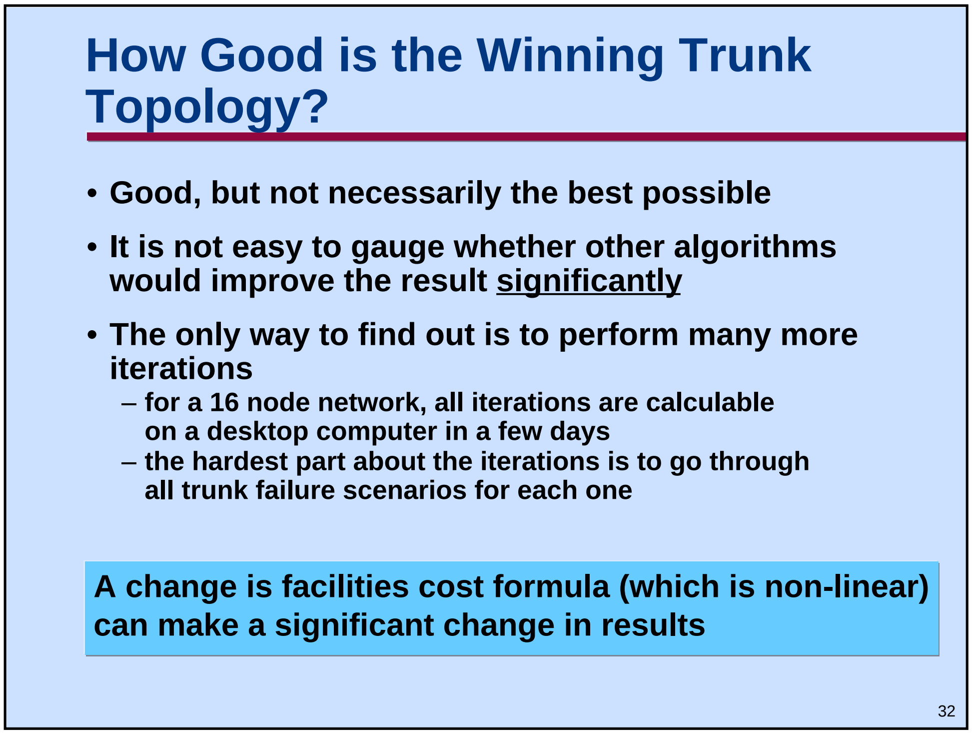

There are probably many other topologies that are better than the winner here for the inputs as specified. But, it is not likely that they will be much better.

A secondary, but important issue, is the cost of the interfaces in the equipments. They will tend to favor the low-trunk topologies on one hand, but, when termination equipment reliability is considered, as it really should be, the OC-12 reliability issue would have to be revisited. That could alter the winning sequence again, but not by much. The best topologies are fairly closely bunched.

As mentioned earlier, a change in tariffs of facilities costs, due to their non-linear impact, can also resequence the contenders very easily.

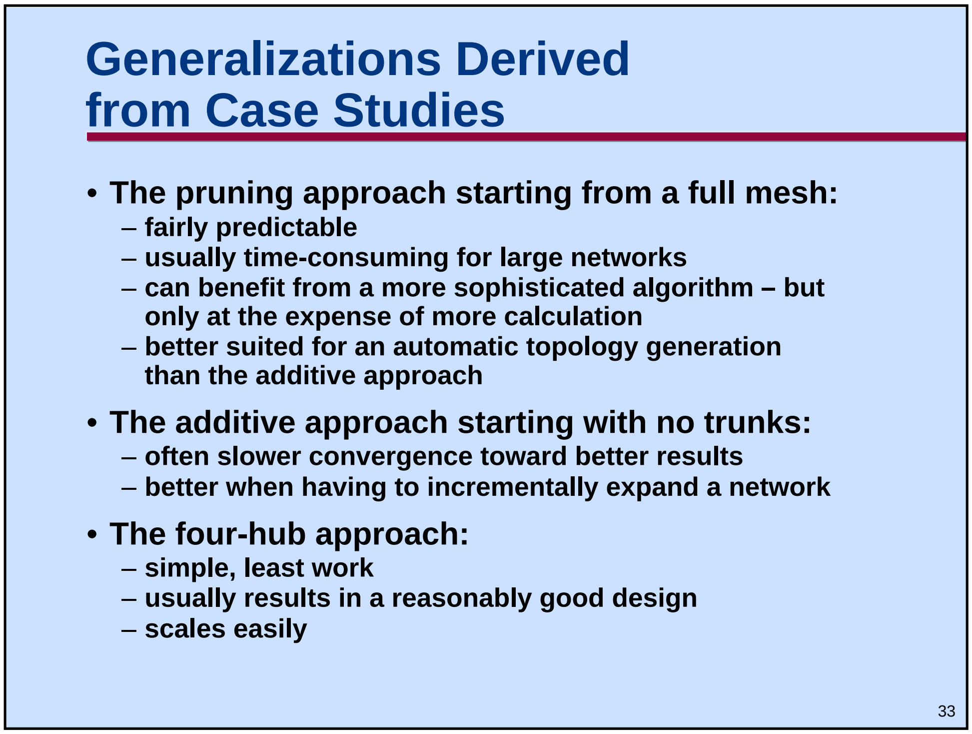

All three approaches have value.

The pruning approach is well suited for automation.

The additive approach is also often used in incremental designs, for obvious reasons.

The hub approach is a good check on the others, and sometimes wins, as it did here.

As a recommendation:

- Use the pruning approach and the four-hub approach for new topology designs.
- Use the additive approach for network upgrade or expansion.

Closing Remarks

- This presentation has focused on the network design part of the facility network layer
- A case study of ATM network design was described, illustrating:
  - three topology design approaches used at Nortel
  - Nortel's focus on designing to meet network performance requirements at minimum cost
  - Nortel's ability to quickly create, compare and evaluate many possible topologies using computer-aided design
- Other Inform '96 presentations discuss Magellan product attributes which impact the three lower layers in the model
- Another presentation provides valuable insights into ATM network engineering

34

Inform '96 presentations which discuss Magellan product attributes at the virtual path layer, call layer and cell layer are :

- Magellan Concorde Infrastructure Solutions, by Quyen Bo
- Magellan Passport: ATM Applications, by David Smith

This presentation has focused on broadband ATM network topology design, and designing to minimum cost. Another Inform '96 presentation, called ATM Network Engineering, by Hosein Badran, provides valuable insights on network engineering, quality of service attributes, classes of service, ATM connection types, connection admission control, network bandwidth allocation, usage parameter control, traffic shaping and the Magellan multiple priority system. An international network design case study is also provided in that presentation.

References

[1] "A layered broadband switching architecture with physical and virtual path configurations", by J.Y. Hui et al., IEEE Journal on Selected Areas of Communications, Vol. 9, No. 9, pp. 1416-1426, Dec 1991

[2] "Virtual Path Routing for Survivable ATM Networks", by K. Murakami and H.S. Kim, IEEE/ACM Transactions on Networking, Vol. 4, No. 1, Feb 1996