ACM SIGCOMM Computer Communication Review
Special Issue on ATM, Volume 25, Number 2, April 1995

Topology Aggregation for Hierarchical Routing in ATM Networks

Whay C. Lee
Motorola ISG, 20 Cabot Boulevard, Mansfield, MA
Email: whay@dma.isg.mot.com

Abstract

This paper provides a tutorial for topology aggregation in large hierarchical communication networks. Specifically, it examines some issues of topology aggregation for hierarchical PNNI Routing in ATM networks, presents a complex node representation, and provides some guidelines for topology aggregation. There are two main reasons for topology aggregation. First, topology information must be compressed to avoid excessive complexity in topology advertisement. Second, the internal topology of a network may have to be hidden for security reasons. A desirable topology aggregation method must adequately represent the topology of a given network for efficient routing and network resource allocation, using a compact advertised topology.

1: Introduction

This paper provides a tutorial for topology aggregation in large hierarchical communication networks. Specifically, it examines some issues of topology aggregation for hierarchical PNNI (Private Network-to-Network Interface) Routing in ATM networks, presents a complex node representation, and provides some guidelines for topology aggregation [1], [2], [3], [4].

An ATM network consists of switching systems represented by nodes, and transmission facilities represented by links. PNNI Routing, designed for ATM networks, has a hierarchical structure. In hierarchical routing, nodes and links at a hierarchical level may be recursively aggregated into higher levels. At the lowest level of the hierarchy, each node represents a switching system. At a higher level, each node represents a collection of one or more nodes at a lower level. Similarly, each link at the lowest level represents a physical link, whereas at a higher level, each link represents a connectivity formed by one or more lower level links in series and/or in parallel.

In PNNI Routing, topology information contained in topology state parameters must be distributed throughout the network and maintained up-to-date for efficient path selection as well as allocation of network resources for establishing connections between end-users of the network. A topology state parameter is a generic term that refers to either a link state parameter or a nodal state parameter. A link state parameter provides information that captures an aspect or property of a link, and a nodal state parameter provides information that captures an aspect or property of a node.

The process of summarizing and compressing topology information at each hierarchical level to determine the topology information to be advertised at the level above is referred to as topology aggregation. There are two main reasons for topology aggregation. First, topology information must be compressed to avoid excessive complexity in topology advertisements. This is particularly necessary for large ATM networks since the number of links in a network can grow as an order of the square of the number of nodes in the network. Second, the internal topology of a network may have to be hidden for security reasons. A desirable topology aggregation method must adequately represent the topology of a given network for efficient routing and network resource allocation, using a compact advertised topology. In practice, topology aggregation, while necessary to achieve scaling in large networks, does not always allow perfect representation of the real underlying topology because of the need to avoid excessive complexity.

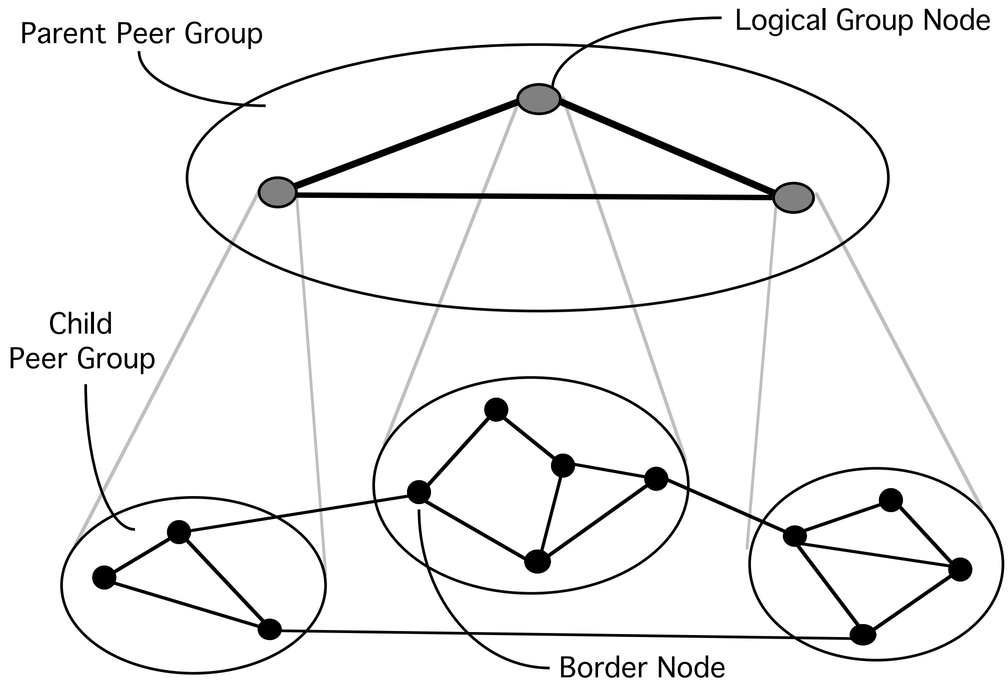

Figure 1: Hierarchical ATM Network

In a hierarchical network, a group of nodes may be abstractly represented by a single point known as a logical node. An aggregated topology associated with a logical node represents the internal connectivity of the logical node. We refer to a communication entity between a pair of reference points (e.g., ports) in a logical node as a logical connectivity, which is different from a logical link that is an abstract representation of the connectivity between two logical nodes.

In PNNI Routing, a peer group is defined to be a group of nodes at a given hierarchical level, where each node exchanges topology information with all other nodes in the group, so that topology information may be synchronized among all nodes in the group for path selection and network resource allocation [1]. A peer group may also represent a collection of lower level peer groups. In this case, the lower level peer groups are referred to as child peer groups with respect to the upper level peer group that represents the collection, and the upper level peer group is referred to as a parent peer group. Each child peer group is represented by a logical group node inside its parent peer group. In other words, a logical group node is a logical node that represents an entire peer group. One of the logical group nodes associated with a peer group is elected to be a peer group leader that performs certain functions, including topology aggregation and advertisement, on behalf of all members of the peer group. Figure 1 shows an example of a two-level PNNI hierarchical network with three peer groups at the lower hierarchical level and one parent peer group at the upper hierarchical level.

A border node in a given peer group is a node having at least one link that crosses the peer group boundary. Typically, only a subset of the nodes in a peer group are border nodes through which transit paths that originate and terminate in external peer groups may enter and leave the peer group. The topology of a peer group is aggregated to form the state of the corresponding logical group node in its parent peer group. The points where links cross the peer group boundary appear as ports in the logical group node. Topology aggregation must retain sufficient topology information to characterize the quality of service available for information transfer between each pair of ports. PNNI Routing obviously requires a path selection algorithm that accepts both link and nodal state parameters as inputs.

Consider a topology formed by a group of N nodes, M of which being border nodes, interconnected by L links, where $M \leq N$. Topology aggregation maps the given topology into a compact topology with M nodes and S logical links, where $S < L$, such that enough topology information may be derived from the aggregated topology for efficient admission and routing of calls. In view of complexity, it is further desirable that S << L.

There are two steps in topology aggregation: full-mesh generation and graph reduction. The first step is to determine an appropriate full-mesh representation of the original topology such that the full-mesh topology consists only of border nodes, where a logical link connects each pair of them together, and an appropriate link state parameter is assigned to each logical link. The second step is to determine what summarized information on the full-mesh representation is to be advertised. This step of graph reduction determines an appropriate subgraph to "encode" the full-mesh, such that the smallest number of logical links are used without unreasonably compromising the accurate representation of the original topology.

2: Conventional Methods for Topology Aggregation

There are two conventional methods for topology aggregation: the Symmetric-Node approach and the Full-Mesh approach [1], [2], [5], [6]. These methods represent the extreme alternatives for a wide spectrum of possible topology aggregation methods representing different levels of trade-off between accuracy and compactness.

In the Symmetric-Node approach, a given topology with multiple nodes is collapsed into one virtual node. An all-in-one parameter (usually the "worst case" parameter that is often loosely known as the "diameter") is advertised. This approach clearly offers the greatest reduction of advertised information. Unfortunately, it does not adequately reflect any asymmetric topology information or capture any multiple connectivity in the original topology. For example, when the transition from one border node to another in a peer group is rendered impossible due to partitioning while transitions across the peer group in other ways remain available, this approach would simply result in a complete impasse in the aggregated topology for the peer group.

The Full-Mesh approach uses a logical link between each pair of border nodes to construct the aggregated topology (e.g., see Figure 2). Note that the Symmetric-Node approach is equivalent to the Full-Mesh approach if the logical links in the full-mesh representation are all identical. The topology information embedded in the full-mesh representation of a given topology appears to be adequate for efficient routing and network resource allocation. Unfortunately, the amount of information to be advertised increases as the square of the number of border nodes. The Full-Mesh approach may be subject to topology explosion that defeats the purpose of aggregation. For example, there is topology explosion when a peer group is a ring network with more than three nodes and every one of the nodes in the peer group is a border node. There may also be considerable redundancy in the aggregated topology determined with this approach when the topology of the original peer group is relatively symmetrical.

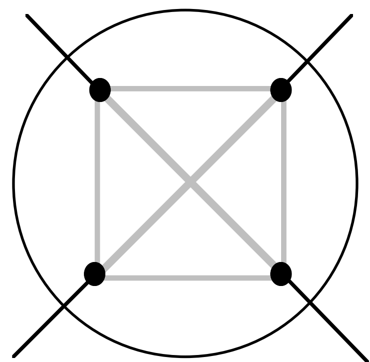

Figure 2: Full-Mesh Approach

A compromise between the two extreme approaches is a Star approach (e.g., see Figure 3), which may be considered as an extension of the Symmetric-Node approach. In the Star approach, the virtual node, that is the center of the "star", is explicitly connected to the border nodes via logical links which are not necessarily identical. This approach has the complexity of an order that is linear in the number of the border nodes. Like the Symmetric-Node approach, this approach does not adequately capture any multiple connectivity in the peer group. However, the approach offers a limited flexibility for reflecting asymmetric topology information.

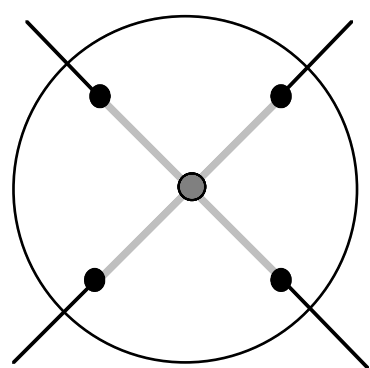

Figure 3: Star Approach

3: Topology Metrics and Attributes

There are two different classes of link or nodal state parameters: metrics and attributes [7], [8], [9]. A link metric is a link state parameter that requires the values of the parameter for all links along a given path to be combined to determine whether the path is acceptable and/or desirable for carrying a given connection, whereas a link attribute is a link state parameter that is considered individually to determine whether a given link is acceptable and/or desirable for carrying a given connection. Nodal metric and nodal attribute are similarly defined. A topology metric is a generic term that refers to either a link metric or a nodal metric, whereas a topology attribute is a generic term that refers to either a link attribute or a nodal attribute. Whenever a qualifier, "link" or "nodal", is left out, reference is assumed to be made to topology state parameters. We discuss in this section the difficulty in topology aggregation with respect to metrics and attributes.

3.1: Metric Aggregation

For metric aggregation, we assume that the full-mesh representation for the metric is obtained such that a logical link between each pair of border nodes is derived from the shortest path, with respect to the given metric, between the same pair of border nodes. The length of the shortest path between a pair of nodes is known as the distance between the nodes. The maximum ratio of the distance in the subgraph to the distance in the original graph between any pair of border nodes is known as the stretch factor associated with the subgraph [10], [11]. It is desirable to encode the full-mesh representation with a compact subgraph whose stretch factor is close to one.

The simple example shown in Figure 4 suffices to demonstrate the difficulty in metric aggregation [12]. Consider a full-mesh network with 4 nodes and 6 links. Suppose that each link has a unit metric weight. The minimum number of links needed to aggregate this topology is 3, and the stretch factor is 2 in this case. If 4 links are used to aggregate the topology, the best approach gives a stretch factor of 2. Even when 5 links are used, the smallest possible stretch factor is still 2. Accurate representation is possible only when all 6 links are used.

In another example as shown in Figure 5, we allow links to be assigned different metric weights. The most compact subgraph that provides an exact encoding (i.e., stretch factor of 1) of the given full-mesh is the 4-link aggregation shown in the figure. Since the subgraph is connected, there exists at least a path connecting any pair of nodes between which there is no logical link. The metric value of each logical link that is not included in the subgraph can be derived by accumulating the metric values along the "shortest" alternate path to the logical link. The 3-link aggregation shown in the figure is actually a minimum weight spanning tree. Although it is the most compact connected subgraph possible for the full-mesh in the figure, it has a stretch factor of 4/3.

For the aggregation of a topology with a symmetric metric, a Spanner Graph approach is found in the literature [11], [13], [14]. While [13] and [14] consider a metric with a unit weight, [11] allows arbitrary metric weights. A subgraph derived from a given graph is said to be a t-Spanner if, between any pair of nodes, the distance in the subgraph is at most t times longer than the distance in the original graph. The value of t is thus the stretch factor associated with the subgraph. It is shown in [11] that, for an undirected graph, there exists a polynomially constructible $(2t+1)$-Spanner such that the number of links on the spanner is smaller than $M*C(M^{1/t})$, where M is the number of nodes and $C(x)$ denotes the integral ceiling of x.

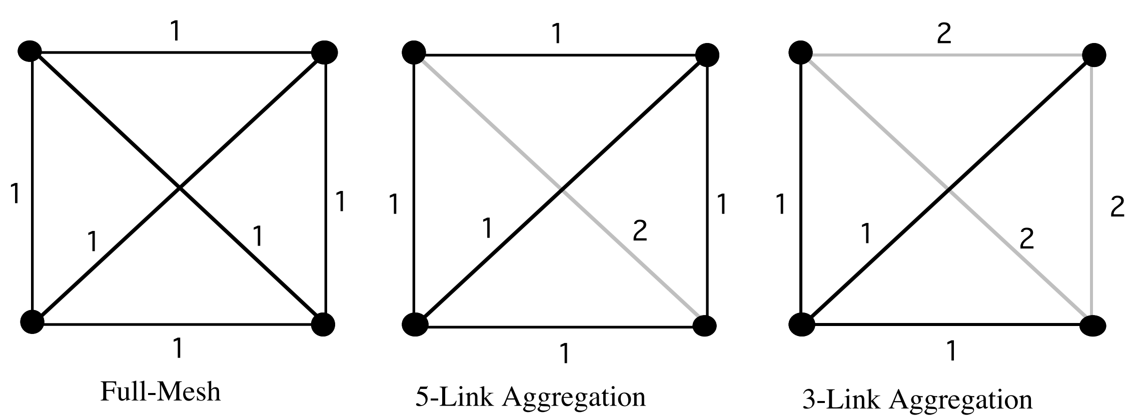

Figure 4: Aggregation of a Topology with a Constant Metric Weight

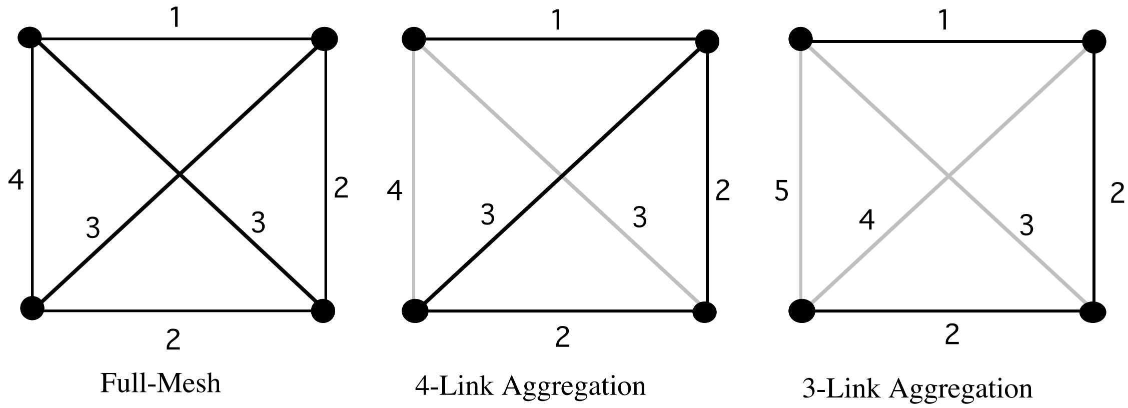

Figure 5: Aggregation of a Topology with Arbitrary Metric Weights

The approach proposed in [11] makes use of an extension of a minimum weight spanning tree algorithm. It scans the links of a given graph in the order of non-decreasing weights. At each iteration, the scanned link is selected if there exists no alternate path in the subgraph of selected links such that the alternate path is shorter than or equal to a fixed number, say r, times the metric weight of the scanned link. When all links have been scanned, the subgraph of selected links gives an r-Spanner for the original graph. Note the trade-off between accuracy and compactness when the above algorithm is used for metric aggregation. For O(M) complexity, the stretch factor is of order log M. Obviously, when the link metric is asymmetrical, the trade-off is less acceptable.

3.2: Attribute Aggregation

For attribute aggregation, we assume that the full-mesh representation for the attribute is obtained such that a logical link between each pair of border nodes is derived from the "optimal path", with respect to the given attribute, between the same pair of nodes. The optimality criterion depends on the attribute. There are two types of attributes. In one type, the minimum link attribute value along a tandem of links determines the attribute value for the path formed by the links. An example of this is Available Cell Rate. In the other type, the maximum link attribute value along a tandem of links determines the attribute value for the path formed by the links. An example of this is Cell Loss Ratio, as it has been decided in the ATM Forum to treat this parameter as an attribute [15]. Unless explicitly noted, we discuss only the first type. It is straightforward to extend our discussion to the other type.

A simple Spanning Tree method for aggregating a topology with a symmetric attribute is available [16], [17]. The method constructs a spanning tree and advertises the attribute values of logical links on the spanning tree, such that the attribute values of logical links that are not on the spanning tree may be derived from the advertised attribute values. There is a unique path connecting each pair of nodes on a spanning tree. If the unique path consists of only one logical link, then the logical link connecting the nodes has obviously been selected for the construction of the spanning tree. If the unique path consists of more than one logical link, then the logical link connecting the nodes has not been selected for the construction of the spanning tree.

For topology aggregation with respect to the first type of attributes (i.e., minimum is the bottleneck), a maximum weight spanning tree with respect to the attribute value as the weight is used for aggregation. An example is shown in Figure 6. In this case, the attribute value of a logical link that is not on the spanning tree is bounded from above by the minimum weight among the logical links along the corresponding unique path. Since the full-mesh representation is determined based on maximizing the attribute value of each logical link, one observes that the attribute value of the unique path between any pair of nodes on the spanning tree cannot be larger than the weight of the logical link connecting the same pair of nodes. We thus conclude that the full-mesh representation derived from maximizing the attribute value of each logical link can be perfectly encoded by a maximum weight spanning tree.

For topology aggregation with respect to the second type of attributes (i.e., maximum is the bottleneck), a minimum weight spanning tree with respect to the attribute value as the weight is used for aggregation.

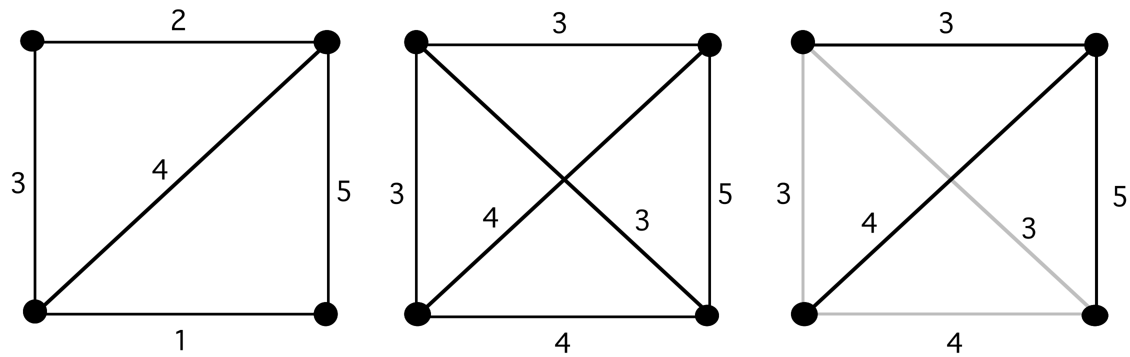

Original Topology Full-Mesh Spanning Tree

Figure 6: Aggregation of a Topology with a Symmetric Attribute

Using the above Spanning Tree approach, one can aggregate a given topology with a symmetric attribute by a compact spanning tree. When there are M border nodes in the given topology, the spanning tree consists of exactly (M-1) logical links. If a given topology has an asymmetric attribute, accurate representation of the topology is often much more complicated. An example with an asymmetric attribute is shown in Figure 7. In this example, there are 12 directed links in the original topology, and only three of them could be excluded from advertisement if accurate representation of the original topology is required. This example can be generalized to one with M nodes [2]. Suppose that the nodes in a given full-mesh topology are numbered 1 through M in a clockwise manner. Let (x,y) denote a directed link from node x to node y. We first arrange the directed links in the following order:

(2,1), (3,1), . . ., (M-1,1), (M,1),
(3,2), (4,2), . . . (M,2),
. . . .
(M,M-1),
(1,2), (2,3), . . ., (M-1,M).

We then assign weights in a strictly decreasing order to each of these links and assume that the weights on all other links are strictly smaller. It can be verified that this topology requires at least L(M) links in order to construct an accurate aggregate representation, where

$$L(M) = \frac{M(M-1)}{2} + (M-1)$$

In conclusion, if accurate aggregate representation is required, the worst case complexity of attribute aggregation is inevitably in the order of M^2. In a typical case, the complexity varies between O(M) and $O(M^2)$, depending on the amount of redundancy in the topology.

4: Asymmetric Topology Parameters

Symmetry in a given topology implies redundancy [18]. It permits efficient topology aggregation without sacrificing accuracy. There are two kinds of symmetry: link-wise symmetry and node-wise symmetry. With link-wise symmetry, the values of a link state parameter in both directions of each link are the same. With node-wise symmetry, there is also link-wise symmetry, and in addition, the values of a link state parameter on all logical links in a full-mesh representation are the same.

Link-wise symmetry is not as well defined for logical links as for physical links. A logical link may be symmetrical because all physical links in a given topology have link-wise symmetry and any selected path always makes use of the same link resources in both directions. A logical link may also be symmetrical because it is made up of two directional paths having the same aggregated parameter values but not sharing the same link resources. In subsequent discussions, we assume the first model.

As we have shown in the previous section, topology aggregation is generally difficult when link metrics and attributes are not symmetrical in opposite directions of a link. Any approach with $O(M^2)$ complexity is clearly unacceptable. An approach with O(1) complexity is desirable, but its capability for supporting efficient routing and network resource allocation is inadequate. An approach with O(M) complexity is reasonable, provided it offers a significant advantage over any approach with O(1) complexity. Given that any approach that permits total asymmetry in the original topology necessarily has a worst case complexity of $O(M^2)$, the only obvious option is to resort to some kind of symmetry.

Traditionally, node-wise symmetry is almost always assumed. Both the Symmetric-Node approach and Full-Mesh approach described in Section 2 take advantage of node-wise symmetry. Although the approaches offer considerable reduction of advertised information, they do not adequately reflect any asymmetric topology information or capture any multiple connectivity in the original topology. Link-wise symmetry is indeed one step towards more accurate topology aggregation. Unfortunately, asymmetric link parameters are rather common in practice. It behooves us to continue looking for better methods of aggregating topology with asymmetric metrics and attributes.

5: Multi-Criteria Path Selection

To derive a full-mesh representation for a given topology, one requires a path selection criterion. PNNI routing requirements imply that there are multiple path selection criteria [1], [19]. Topology metrics are usually associated with their corresponding path selection criteria. Topology attributes are usually not subject to any path selection optimization. The following alternative approaches are available for dealing with multiple topology state parameters and path selection criteria [18].

One approach is to choose the "best" path based on a combination of path selection criteria. This approach is impractical because of the lack of a good understanding of the tradeoff among the different criteria. For example, one might consider selecting paths based on the minimization of a linear combination of Cell Transfer Delay and Administrative Weight. Then, it is not clear how one might choose the relative weights for combining the two objective functions.

Another approach is to stick to one full-mesh representation based on a predetermined path selection criterion for all topology state parameters. This approach may lead to conservative advertisements. It is also not clear which path selection criterion should be used.

Yet another approach is to use a different full-mesh representation based on a given path selection criterion for each topology state parameter associated with the same path selection criterion. This approach may lead to aggressive advertisements. Moreover, it is not clear what to do with topology state parameters that are not associated with any path selection criterion.

In view of insufficient understanding of the above approaches, none of them has been adopted as a requirement for PNNI Routing. If a topology advertisement is conservative, available resources cannot be efficiently tapped, and calls are blocked at their respective sources during path computation. On the other hand, if a topology advertisement is aggressive, calls are blocked during connection setup, call setup time may be excessive due to crankback, and signaling overhead may be high.

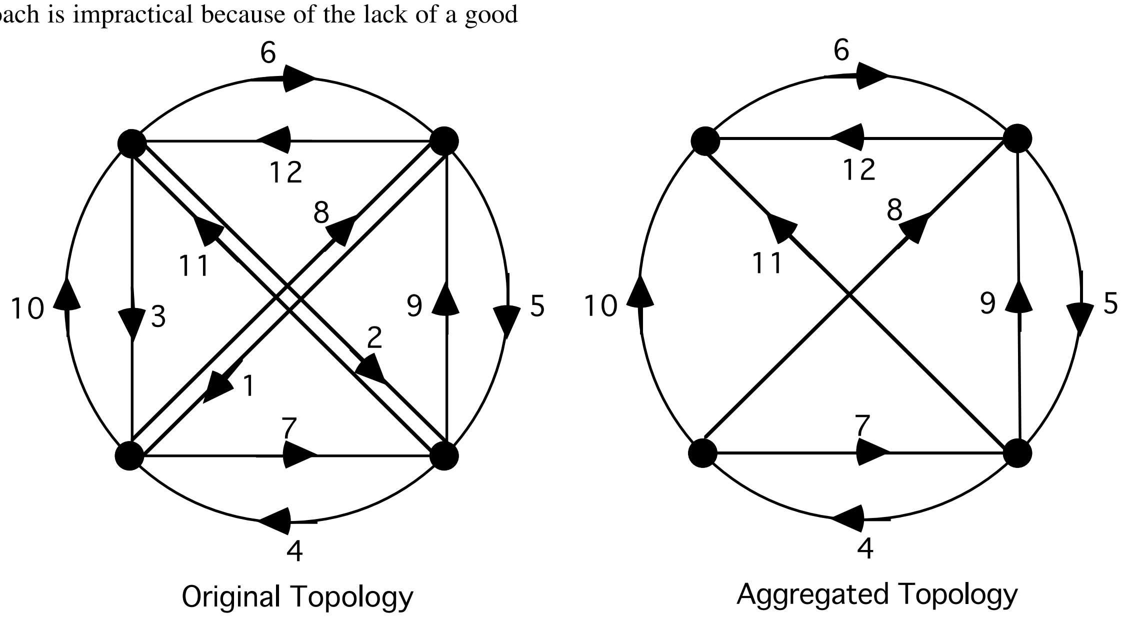

Figure 7: Aggregation of a Topology with an Asymmetric Attribute

6: Complex Node Representation

A complex node representation is a collection of nodal state parameters that provide detailed state information associated with a logical node. This representation provides aggregated topology information not only for traversing the logical node, but also routing to and from the "inside" of the logical node. The interior reference point of a logical node is referred to as a nucleus. A border reference point of a logical node is known as a port. A logical connectivity between the nucleus and a port of the logical node is referred to as a spoke. The complex node representation is a flexible representation for describing the connectivity among border nodes in a peer group. There is plenty of flexibility for choosing between conservative and aggressive advertisements.

PNNI Routing supports a default node representation that consists of a single value for each nodal state parameter giving a presumed value between any port and the nucleus of the logical node (See Figure 8). In the default node representation, there is no advertisement of any port-to-port connectivity. This default, based on the Symmetric-Node model described in Section 2, has a complexity of O(1). Unfortunately, it does not provide much accuracy as mentioned before.

Given that a logical node is in general not "perfectly round", the next level of sophistication is to allow any of the spokes in the star topology be different from the default "radius" to reflect some asymmetry in the logical node. More often than not, topology aggregation is a lossy process, such that one may not be able to recover the original topology information from its aggregated version. There are, however, situations where one might want to selectively expose a logical connectivity so that its parameter values are not lost in the aggregation. For such flexibility, PNNI Routing allows logical connectivities between ports to be designated as exceptions. An exception is a connectivity advertisement that represents something other than its default.

The complex node representation for PNNI Routing is summarized below [3], [18], [20], [21]:

1) Conceptually overlay on each logical node a star topology with a nucleus representing the "inside" of the logical node, and spokes connecting the ports of the logical node to the nucleus.
2) For each nodal state parameter, advertise a "radius" to be used as the default value for the spokes.
3) Any spoke or any logical connectivity between a pair of ports may be designated as an exception.
4) For each such exception, advertise the entire set of nodal state parameters associated with it.
5) For each spoke advertised as an exception, the advertised value of each nodal state parameter supersedes its corresponding default value.
6) A path through the logical node is obtained from the "best" concatenation of exceptions and/or default spokes in the complex node representation.

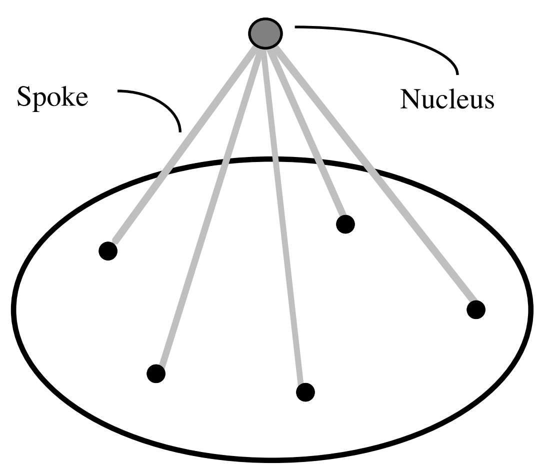

Figure 8: Default Node Representation

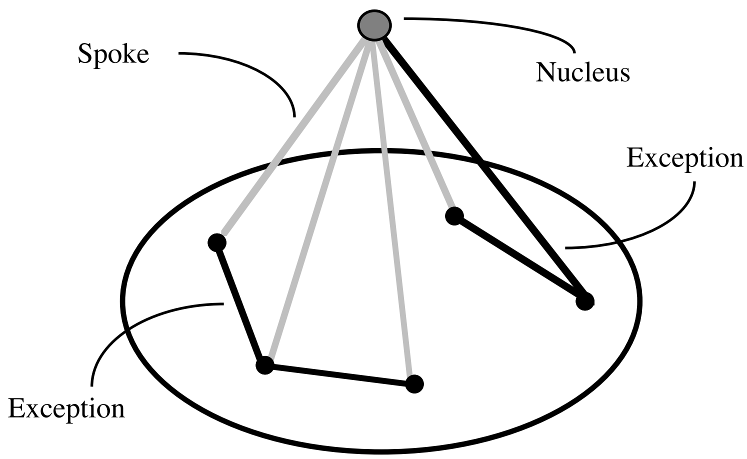

Figure 9: Complex Node Representation

A typical configuration of a PNNI complex node representation is shown in Figure 9. The complex node representation allows one to advertise practically any aggregated topology ranging from Symmetric-Node to Full-Mesh, and there are hooks for supporting a variety of other aggregation methods.

7: Topology Aggregation Guidelines

To preserve the flexibility of the complex node representation for PNNI Routing, no requirements are specified for when and how the nodal state parameters constituting the complex node representation should be generated. Nevertheless, several guidelines, as shown below, have been adopted for the normal operation of PNNI topology aggregation [4].

7.1: Default Node Representation

For each nodal state parameter, the "diameter" of a logical node is defined as an aggregation of all parameter values in a full-mesh representation of the logical node. Each "diameter" must be converted to a "radius". For a nodal metric, the "radius" is simply half the "diameter". For a nodal attribute, the "radius" is the same as the "diameter".

7.2: Multiple Nodal State Parameters

A conservative approach determines the entire set of parameters based on a single path selection criterion, whereas an aggressive approach determines each of them based on a possibly different path selection criterion. The degree of aggressiveness is left to the discretion of the advertisers.

7.3: Number of Exceptions

While PNNI Routing does not specify a hard limit on the number of exceptions, it is recommended that the number of exceptions used to configure a complex node representation of a logical node be normally kept smaller than 3 times the number of ports in the logical node.

7.4: Significant Exceptions

Exceptions are useful if they provide "significantly different" topology information than that revealed by the default node representation. For example, one could designate a spoke on a default node representation of a peer group to be an exception to expose an outlier, which is a node whose exclusion from its containing peer group would significantly improve the accuracy and simplicity of the aggregation of the remainder of the peer group topology.

7.5: Useless Exceptions

It is useful to identify conditions under which a logical connectivity within a logical node is not to be designated as an exception. For example, we may consider the following condition. A logical connectivity is said to be dominated by a path if at least one parameter value associated with the logical connectivity is less desirable than the corresponding parameter value of the path, and all other parameter values are pair-wise equal. Since PNNI Routing selects internal paths through a logical node from the "best" concatenation of exceptions and/or default spokes in the complex node representation, it is not useful to designate a logical connectivity to be an exception if it would be dominated by at least one alternate path in the advertised topology.

8: Conclusion

Topology aggregation is a mechanism for complexity reduction and information hiding in hierarchical ATM networks. This paper has provided an overview of the problem and the difficulty in searching for a satisfactory solution. There are many challenging open issues that await further investigations.

In PNNI Routing, topology advertisements may be aggressive such that an advertised set of metrics and attributes is actually not feasible even though the parameters are individually available. With aggressive advertisements, a call setup is subject to a "budgeting problem", whereby a cumulative metric value increases faster than anticipated by the originating source of a call during the connection establishment for the call, and control messages are generated as the connection setup is subsequently aborted due to insufficient network resources [22]. On the other hand, with conservative advertisements, network resources may be under-utilized as calls may be unnecessarily blocked by a remote source because the source does not know that there are actually enough network resources to support the calls. The implications of and trade-off between aggressive and conservative advertisements are open for further studies.

Aggregation of a given topology with a symmetric attribute requires only O(M) complexity, where M is the number of border nodes in the topology. For example, the Spanning Tree approach produces a compact aggregated topology with (M-1) logical links. The Spanning Tree method can be extended for aggregating a topology with an asymmetric attribute. Some preliminary work on this was reported in [16]. Unfortunately, in this case, the logical links whose attribute values are advertised do not necessarily form a spanning tree, and the complexity varies between O(M) and $O(M^2)$.

Topology aggregation with respect to a metric, symmetric or not, is a complicated and lossy process. The Spanner Graph approach for metric aggregation appears promising if accuracy is not critical. Hierarchical topology aggregation, while attractive at the bottom hierarchical level, would result in compounding errors at higher hierarchical levels. If conservative advertisements derived from the Spanner Graph approach are permitted, especially if the stretch factor achievable with an aggregated topology of reasonable complexity is much greater than one, the efficiency of routing and network resource allocation can be severely compromised.

Although the traditional assumption of node-wise symmetry permits considerable simplicity in topology aggregation, it is inadequate in today's ATM environment. Link-wise symmetry is a step towards more accurate topology aggregation, with a marginal increase in complexity. However, since asymmetric link parameters are common in practice, better methods of aggregating topology with asymmetric metrics and attributes are desirable. Barring any complexity that is greater than O(M), there is little hope for accuracy in topology aggregation. Should accuracy be desired, it is likely that either the aggregation algorithm is complicated, and/or the complexity of the aggregated representation is considerably greater than the order of the number of border nodes.

The PNNI complex node representation is a flexible representation for describing the connectivity among border nodes in a peer group. The use of exceptions provides hooks for existing as well as future topology aggregation methods. The default node representation offers a significant trade-off of accuracy for simplicity of topology aggregation. Outliers can be exposed by designating appropriate spokes to be exceptions. The nucleus may be used to represent a collection of reachable end-systems within a peer group, provided the reachable end-systems are indeed "close" to the "inside" of the peer group. The current version of the PNNI complex node representation is not adequate in capturing a topology where some reachable end-systems are "far" from the "inside" of the peer group [23]. An extension of the complex node representation to address this issue is yet to be determined.

Acknowledgment

The author thanks his colleagues who participate in the PNNI Subworking Group in the ATM Forum for valuable contributions to the understanding of the issues of topology aggregation and the formulation of the PNNI complex node representation. In particular, the author acknowledges Kumar Sivarajan for providing the example in Figure 7, and its generalization.

References

[1] Doug Dykeman (Editor), "PNNI Draft Specification," ATM Forum, Contribution 94-0471 (latest version), March 1995.

[2] Kumar Sivarajan & Whay Lee, "Issues in the Aggregation of Link State Parameters in Hierarchical PNNI Networks," ATM Forum, Contribution 94-0646, July 1994.

[3] Whay Lee & Ted Tedijanto, "Proposed Baseline Text for Topology Aggregation in PNNI Routing," ATM Forum, 94-1154, November/December 1994.

[4] Whay Lee, "Some Guidelines for PNNI Topology Aggregation," ATM Forum, 95-0153, February 1995.

[5] Joel Halpern & Ross Callon, "Link and Node Metrics," ATM Forum, Contribution 94-0590, July 1994.

[6] Vasant Sahay, Jonathan Davar, & Farokh Deboo, "Metrics for Representing Internal Topology of a Logical Group Node," ATM Forum, 94-0449, May 1994.

[7] Ted Tedijanto & Whay Lee, "Proposed Baseline Text for Link State Parameters and Update Protocol," ATM Forum, 94-0645, July 1994.

[8] Whay Lee, "Classification of Link State Parameters," ATM Forum, 94-0672, July 1994.

[9] Whay Lee, "Classification of Nodal State Parameters," ATM Forum, 95-0134, February 1995.

[10] David Peleg & Eli Upfal, "A Trade-Off between Space and Efficiency for Routing Tables," Journal of the ACM, Vol. 36, No. 3, pp. 510-530, July 1989.

[11] Ingo Althofer, Gautam Das, David Dobkin, Deborah Joseph, & Jose Soares, "On Sparse Spanners of Weighted Graphs," Discrete Computational Geometry, Vol.9, pp 81-100, 1993.

[12] Whay Lee, "Complexity of Metric Aggregation," ATM Forum, 95-0200, February 1995.

[13] David Peleg & Alejandro A. Schaffer, "Graph Spanners," Journal of Graph Theory, Vol. 13, No. 1, pp. 99-116, 1989.

[14] Baruch Awerbuch, "Complexity of Network Synchronization," Journal of the ACM, Vol. 32, No. 4, pp. 804-823, 1985.

[15] Whay Lee & Joel Halpern, "Cell Loss Ratio — Link Attribute or Link Metric?" ATM Forum, 95-0133, February 1995.

[16] Whay Lee, "A Minimal Spanning Connectivity Method for Link State Aggregation," ATM Forum, 94-0673, July 1994.

[17] Whay Lee, "Spanning Tree Method for Link State Aggregation in Large Communication Networks," IEEE INFOCOM '95, April 1995.

[18] Whay Lee, "Topology Aggregation in PNNI Routing," ATM Forum, 94-0874, September 1994.

[19] Whay Lee & Michael Hluchyj, "Multi-Criteria Routing subject to Resource and Performance Constraints," ATM Forum, 94-0280, March 1994.

[20] Joel Halpern, "Nodal Metrics with Pairwise Values," ATM Forum, 94-0862, September 1994.

[21] Ethan Spiegel, "Representation of Node Parameters Using Exceptions," ATM Forum, 94-0907, September 1994.

[22] Joel Halpern, "Allocating Resources in Transit Peer Groups," ATM Forum, 95-0100, February 1995.

[23] Whay Lee & Ted Tedijanto, "Topology Aggregation and Path Selection in PNNI Routing," ATM Forum, 94-1155, November 1994.