DIVIDING A NETWORK INTO PEER GROUPS TO BUILD A HIERARCHICAL STRUCTURE

Piet Van Mieghem

AReNA

Alcatel Corporate Research

Francis Wellesplein 1, B-2018 Antwerp (BELGIUM)

mieghemp@rc.bel.alcatel.be

ABSTRACT

A strategy of how to divide a network into peer groups is discussed. The work is inspired by the ATM Forum's PNNI specification. The proposed strategy relies on two pillars, a symmetry criterion and the concept of strongly connected components in graph theory.

The hierarchy, which offers network scalability, is without doubt the most advanced feature in PNNI. However, the construction of this hierarchy is still a topic of a current research. This article investigates a first phaze in building the hierarchy (how to combine nodes to form a peer group or subnet) while the second phaze, the information condensation, is discussed elsewhere (Van Mieghem, 1998a). Finally, for the routing in a hierarchy, we refer to Van Mieghem (1998a).

INTRODUCTION

The division of a large network into peer groups is necessary to build a hierarchical structure of the network. The hierarchy furnishes scalability (Van Mieghem, 1997). Scalability will become an increasingly attractive feature in future networks because the network intelligence functions, such as routing, flooding, ..., tend to grow in complexity as the number of nodes and links in the network increases. The complexity in network functions augments because the services are demanding more quality of service (QoS).

In particular, where best-effort, unreliable delivery (as in the current Internet) only requires a shortest path routing function with a complexity roughly quadratic in the number of nodes, QoS-aware networks are confronted with QoS routing (multi-parameter routing subject to QoS requirements). In principle, QoS routing has a non-polynomial complexity when the number of nodes and links increase (Wang and Crowcroft, 1996). Interestingly, also for QoS routing, very efficient algorithms can be designed such as TAMCRA (Tunable Accuracy Multiple Constraint Routing Algorithm) proposed in De Neve and Van Mieghem (1998).

There are many factors that may influence the division of a large network into peer groups. Without pretending to be complete, we list the more important ones. First, the topology itself is determined by geographic and demographic aspects. For instance, a large population on a small area may give rise to many strongly interconnected nodes. Further, a network may extend over country boundaries. Differences in legal and financial (taxes,...) regulations may encourage an operator to divide his network per country. But, even in a same country, a large network may be owned by various operators who want to control their part of the network. The separation of the network may be based on policy agreements.

In the sequel, we only consider how one, large topology, entirely owned and controlled by one single authority can be divided into peer groups. Hence, we assume that legal, policy and country factors have already been taken into account. We dispose, so to say, of a large homogeneous network where each node is equally important. Thus, only topological considerations are invoked to propose a strategy to divide the large network into peer groups (or subnets).

By the best of our knowledge, network planning aspects in hierarchies as discussed here have not been discussed earlier.

The division of a network into peer groups is the first step to build the hierarchy. The following process is information condensation (see e.g. Van Mieghem, 1998), i.e. node and link aggregation, where each peer group on a lower level in the hierarchy, say *k*, is represented by one complex (or logical) node on level *k+1*.

After the information condensation, we obtain a reduced network view on level *k+1*. At that level, again a division process is needed to combine these complex nodes in level *k+1* peer groups, that, in a next information condensation cycle lead to the network view on level *k+2*.

Recursively applying the topology condensation and 'division into peer groups' processes, yields the eventual hierarchical structure of the network, as shown in Figure 1.

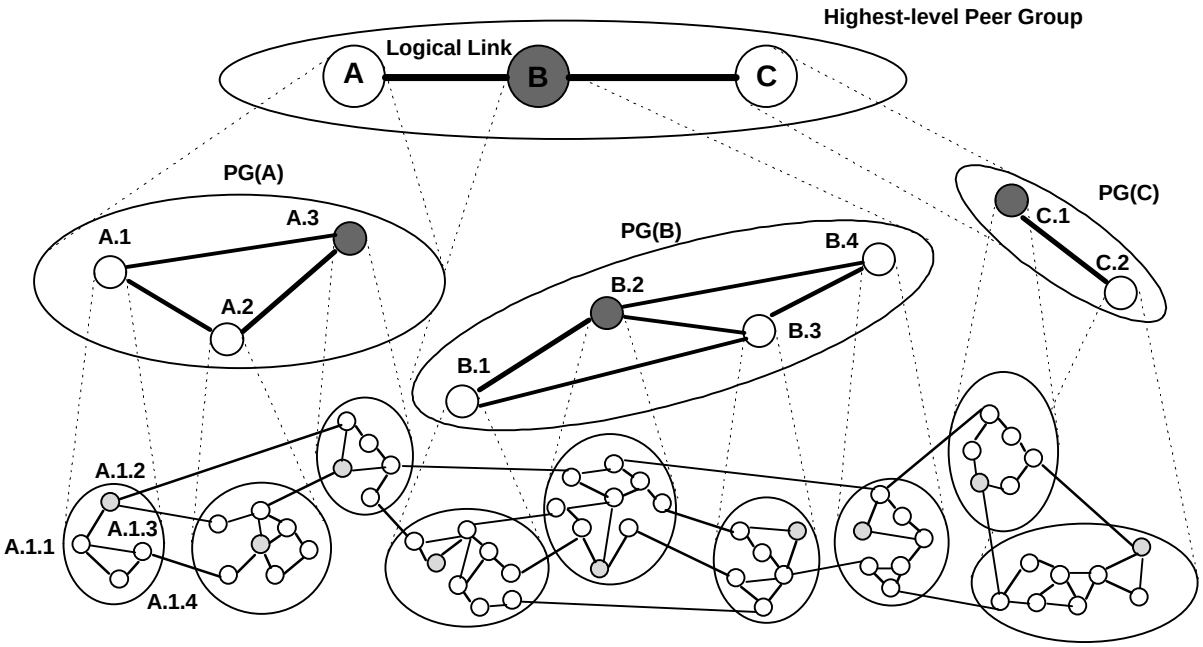

Figure 1. A hierarchical structure based on the specifications of PNNI where a subnet is coined a peer group, denoted as PG(.).

STRATEGY

The proposed strategy to separate a network into peer groups draws on two pillars.

On the one side, our earlier dimensioning study (Van Mieghem, 1997) shows the optimal hierarchical structure of a PNNI topology and gives clear indications how to obtain it. In particular, the results demonstrate that symmetry should be aimed at in the sense that all peer groups on a certain level *k* should contain a same number of nodes. We refer to this first guideline as the *symmetry criterion*.

In addition, the analysis demonstrated that, for an increasing hierarchical level *k*, the number of (logical) nodes on that level *k* (which equals the number of peer groups on level *k-1*) seriously grows, a result somewhat unexpected and opposed to what is shown in Figure 1, where, on the average, a lower level contains more nodes than a higher one.

Another useful result is the explicit expression of the number of hierarchical levels *N*, when the lowest level peer groups all contain $x > 2$ (physical) nodes and the total number of physical nodes is *M* (i.e. the total number of nodes on the lowest level of the hierarchy $k = 0$), as

$$N = [\log_2 (1 + \log_2 M/\log_2 (x/2))]$$

where $[y]$ denotes the integer smaller than (or equal to) y. However, these results are based on the following assumptions:

(1) the optimization criterion is the worst case routing complexity *C*, (2) the routing complexity is additive and (3) the worst case complexity only depends on the number of nodes (as is, e.g., the case for the Dijkstra algorithm where $C = O(M^2)$). These assumptions are quite valid if we can construct the original network (and its hierarchy) from scratch on. In most cases, a network is given and defined by a topology consisting of *M* nodes and *E* links. In order to take the existing interconnectivity into account, modifications of our previous analysis (Van Mieghem, 1997) are needed, reflected by a second pillar.

This second pillar relies on *strongly connected components* in graph theory (Cormen *et al.*, 1990, sec. 23.5). A strongly connected component of a directed graph *G(V,E)* with *E* links (edges) and *V* nodes (vertices) is a maximal set of vertices *U* such that every pair of vertices *u* and *v* in *U* are reachable from each other. The problem of finding the strongly connected components of a directed graph is related to the problem of determining the ergodic subchains and transient states of a Markov chain. There exists an algorithm due to Tarjan (1972), based on depth first search, that computes the strongly connected components in linear-time, i.e. O(V+E). In other words, no matter how large the original topology, a separation in strongly connected components is always best feasible.

The determination of strongly connected components relies on the underlying (given) topology, while the symmetry criterion is deduced from optimizing the worst case routing complexity. The latter makes abstraction of the underlying topology and only takes the number of nodes into account. Clearly, the pursue of both objectives, may lead to conflicts. Indeed, based on the original topology, the different strongly connected components may widely vary in the number of nodes so that the peer groups are asymmetrical with respect to the number of nodes. Hence, on the first optimality directive, the resulting hierarchical structure will be suboptimal. In case of conflicts, there are two choices.

Either change the underlying topology, which is rather drastic and usually impossible or combine smaller strongly connected components into one peer group so that the symmetry criterion is more or less satisfied. Hence, we may summarize the strategy as follows:

1. Compute the strongly connected components of the directed graph *G(V,E)*.
2. If the number of nodes in the strongly connected components are about the same, call these strongly connected components the peer groups and we are done.
3. else, assign smaller strongly connected components (adjacent to each other) to one peer group so that the symmetry criterion is obeyed as closely as possible.

There remains to explain how step 3 in the strategy must be performed. Replace each strongly connected component by one node, coined a “strongly connected component node” or *scc-node* and reduce the number of links connecting a pair of scc-nodes to one. Omit those scc-nodes corresponding to large strongly connected components of about the same size. Determine in this condensed and reduced topology again the strongly connected components. The resulting “second order” strongly connected components (i.e. strongly connected components of scc-nodes) surely contain more *original* (or physical) nodes. Check whether these “second order” strongly connected components are of the same size (in terms of number of physical nodes) as the “first order” components that we have omitted. If so, assign the latter strongly connected components to peer groups else, repeat the process. We believe it very unlikely that the process needs a third iteration.

How suboptimal the hierarchical structure is when neglecting the symmetry criterion is a yet unanswered question. It is an interesting issue for further study. In case symmetry is only a weak requirement, the strategy above will simplify considerably (only step 1 is needed).

ADDITIONAL CONSIDERATIONS

The proposed link and node strategy (Van Mieghem, 1998) is fine-tuned for this strategy because on each level of the hierarchy two complex nodes are connected by only one link. This makes that the strongly connected components on higher hierarchical levels are similarly computed than on the physical level. On the other hand, assigning strongly connected components to peer groups has the advantage that the peer groups are internally strongly connected, but less with other peer groups, reflecting the original purpose of the notion of "peer group".

This also means that the number of border nodes is likely to be smaller than the number of internal nodes. As only these border nodes have connectivity with other peer groups, they are crucial in the node and link aggregation.

The lower the number of border nodes, the simpler both node and link aggregation (also called topology condensation) is. Hence, the concept of strongly connected components seems both advantageous for topology condensation and for the 'division into peer groups' process.

REFERENCES

ATMF, 1996, Private Network Network Interface (PNNI), Specification Version 1.0, March 1996.

Cormen, T. H., C. E. Leiserson and R. L. Rivest, 1990, *Introduction to Algorithms*, The MIT Press, Cambridge, Massachusetts.

De Neve, H. and P. Van Mieghem, 1998, "A Multiple Quality of Service Routing Algorithm for PNNI", to appear in Proceedings of the IEEE ATM Workshop, May 26-29, Fairfax, USA.

Tarjan, 1972, "Depth-first Search and Linear Graph Algorithms", SIAM J. Comput., Vol. 1, No. 2, June, pp. 146-160.

Van Mieghem, P., 1997, "Estimation of an Optimal PNNI Topology", Proceedings of the IEEE ATM'97 Workshop, May 26-28, Lisboa, Portugal, pp. 570-577.

Van Mieghem, P., 1998, "Node and Link Aggregation in a Hierarchy", submitted to IEEE/ACM Trans. on Networking.

Van Mieghem, P., 1998a, "Routing in a hierarchical structure", to appear in the Proceedings of First IEEE International Conference on ATM (ICATM'98), June 22-24, Colmar, France.

Wang Z. and J. Crowcroft, 1996, "Quality-of-Service Routing for Supporting Multimedia Applications", IEEE Journal of Selected Areas in Communications, vol.14, pp.1228-1234.