Estimation of an Optimal PNNI Topology

Piet Van Mieghem
AReNA

Alcatel Research Division
Francis Wellesplein 1, B-2018 Antwerp, Belgium¤

Abstract

Optimality of the PNNI topology is investigated analytically based on a routing complexity criterion that only depends on the number of nodes in a (sub)network. Results are presented for a complexity that is quadratic in the number of nodes because this case corresponds to the complexity of the well-known Dijkstra routing algorithm.

The resulting optimal PNNI topology is highly symmetrical (all peer groups on a same hierarchical level are of equal size). Moreover, the lowest (physical) peer group size consists only of a few nodes, while, surprisingly perhaps, just a couple of hierarchical levels are needed to span a global network of 10^{10} nodes. Routing algorithms with a higher nodal complexity than Dijkstra's bene…t from a higher complexity reduction using a hierarchical PNNI optimal topology.

1 Introduction.

The ATM Forum's PNNI topology favours a hierarchical structure of peer groups on di¤erent levels [2]. It is of interest to know how to structure a given network of M nodes into a PNNI topology with N hierarchical levels so that a minimal routing complexity is achieved.

In the analysis below, the complexity C to compute some measure in this PNNI network is assumed to be dependent only on the number of nodes, hence $C = f(M)$. For example, to calculate the shortest path between node a and b, we can use the Dijkstra algorithm that has a complexity of $C = O(M^2)$. In addition, we assume that the complexity is additive meaning that the complexity of a network consisting of two subnetworks equals the sum of the complexities in each subnetwork. These simplifying assumptions allow an exact solution of the problem. On the other hand, the results obtained are only indicative because the additivity assumption ignores direct interaction between subgroups on one level. Speci…cally, PNNI uplinks are not taken into account. Finally, only the routing complexity in a static PNNI topology is considered. This means that computational (or complexity) aspects of information condensation (such as node and link aggregation) are not given any attention.

Before presenting the analysis, a few words on the notation used are appropriate. The number of nodes in a certain peer group on hierarchical level k is denoted as $m_{r_1 r_2 r_3 \cdots r_k}$. This presentation reveals all ancestrial information of this peer group. For example, it demonstrates

¤Proceedings of the IEEE ATM Workshop, Lisabon, May 25-28, 1997, pp. 570-577
tel.: (32)-3-240 91 62, fax: (32)-3-240 99 32, email: pvmi@rc.bel.alcatel.be

that the peer group in question has as parent the peer group with index $r_1\, r_2\, r_3 \cdots r_{k-1}$. The highest level ancestor is indexed by $r_1 = 1$, because at that level, there can be only one peer group enveloping the whole network. However, that peer group possesses m_{r_1} nodes. A sketch of a PNNI topology using this terminology is drawn in Fig. 1.

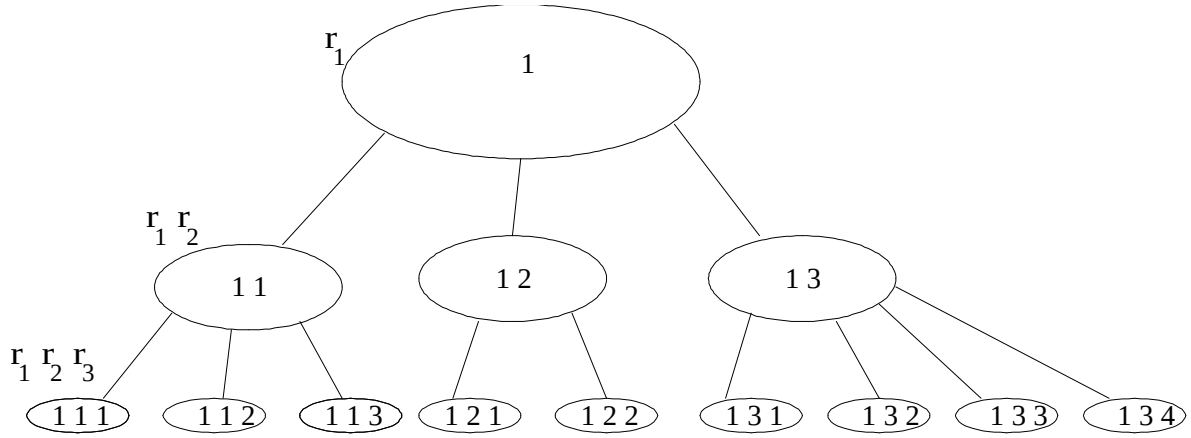

Figure 1: A sketch of the here considered simplified PNNI hierarchical topology.

2 The optimal solution.

On the assumption of the additivity of the complexity, we simply need to sum the complexities in all peer groups. Clearly, this implies a worse case approach because we assume that a path traverses the total hierarchy of peer groups. Thus, on the highest level the complexity is $C_1 = f(m_{r_1})$ while on the subsequent lower level, consisting of m_{r_1} different peer groups, we have $C_2 = \sum_{r_2=1}^{m_{r_1}} f(m_{r_1 r_2})$. Proceeding in this way, we verify that

$$C_k = \sum_{r_2=1}^{m_{r_1}} \sum_{r_3=1}^{m_{r_1 r_2}} \cdots \sum_{r_k=1}^{m_{r_1 r_2 r_3 \cdots r_{k-1}}} f(m_{r_1 r_2 r_3 \cdots r_k}) \qquad (1)$$

Hence, the total complexity reads $C_{pnni} = \sum_{k=1}^{N} C_k$ or

$$C_{pnni} = \sum_{k=1}^{N} \sum_{r_2=1}^{m_{r_1}} \sum_{r_3=1}^{m_{r_1 r_2}} \cdots \sum_{r_k=1}^{m_{r_1 r_2 r_3 \cdots r_{k-1}}} f(m_{r_1 r_2 r_3 \cdots r_k}) \qquad (2)$$

Our goal is to minimize the complexity. We invoke the method of Lagrange multipliers that optimizes (2) subject to the constraint that the total number of nodes on the lowest peer group level equals the number of nodes M in the original network. Only this lowest level consists of physical nodes, the higher levels are merely logical nodes. Thus, we have as Lagrange function L,

$$L = C_{pnni} + \lambda \left(\sum_{r_2=1}^{m_{r_1}} \sum_{r_3=1}^{m_{r_1 r_2}} \cdots \sum_{r_N=1}^{m_{r_1 r_2 r_3 \cdots r_{N-1}}} m_{r_1 r_2 r_3 \cdots r_N} - M \right) \qquad (3)$$

In order to fully exploit the strengths of functional analysis, we extend the integers m_i to real numbers. The system of extremal equations then becomes

$$\frac{\partial L}{\partial m_{r_1 r_2 r_3 \cdots r_k}} = 0 \qquad \text{for all } r_k \text{ with } 1 \le k \le N \tag{4}$$

$$\frac{\partial L}{\partial \lambda} = 0 \tag{5}$$

For each hierarchical level k, one may verify by induction from (3) that

$$\frac{\partial L}{\partial m_{r_1 r_2 r_3 \cdots r_{k-1} i}} = \frac{\partial L}{\partial m_{r_1 r_2 r_3 \cdots r_{k-1} j}} \qquad (\text{for all } 1 \le i, j \le r_k) \tag{6}$$

because each peer group on a same hierarchical level plays a similar role leading to formally identical partial derivatives.

This relation (6) indicates that an optimal solution exhibits symmetry. Moreover, this symmetry substantially simplifies the analysis because it means that on each hierarchical level, the peer groups have a same number of nodes. Consequently, guided by this symmetry, the notational burden can be significantly reduced: there is no longer any reason to specify the ancestrial dependence in the notation. Hence, in the optimal topology, we write m_k instead of $m_{r_1 r_2 r_3 \cdots r_k}$. Using this simpler notation, the Langrange function (3) becomes

$$L = \sum_{i=1}^{N} f(m_i) \prod_{j=1}^{i-1} m_j + \lambda \left(\prod_{j=1}^{N} m_j - M \right) \tag{7}$$

while the set of extremal equations simplifies to

$$\frac{\partial L}{\partial m_k} = f'(m_k) \prod_{j=1}^{k-1} m_j + \sum_{i=k+1}^{N} f(m_i) \prod_{j=1; j \neq k}^{i-1} m_j + \lambda \prod_{j=1; j \neq k}^{N} m_j = 0 \tag{8}$$

$$\frac{\partial L}{\partial \lambda} = \prod_{j=1}^{N} m_j - M = 0 \tag{9}$$

This set (8,9) should be solved for $\{m_k\}$. Concentrating on the first set of N equations in (8), we observe recursion,

$$f'(m_k) = -\left(\sum_{i=k+1}^{N} f(m_i) \prod_{j=1+k}^{i-1} m_j + \lambda \prod_{j=k+1}^{N} m_j \right) \tag{10}$$

$$f'(m_N) = -\lambda,$$

A brief inspection reveals a quite simple relationship between m_{k+1} and m_k, for we can rewrite the first relation as

$$\begin{aligned} f'(m_k) &= -f(m_{k+1}) - \left(\sum_{i=k+2}^{N} f(m_i) \prod_{j=k+1}^{i-1} m_j + \lambda\, m_{k+1} \prod_{j=k+2}^{N} m_j \right) \\ &= -f(m_{k+1}) - m_{k+1} \left(\sum_{i=k+2}^{N} f(m_i) \prod_{j=k+2}^{i-1} m_j + \lambda \prod_{j=k+2}^{N} m_j \right) \end{aligned}$$

With (10) applied to hierachical level k + 1, the last term between brackets is identi…ed as $-f'(m_{k+1})$. Hence, we …nd the important recursive relationship

$$f'(m_k) = -f(m_{k+1}) + m_{k+1} f'(m_{k+1}) \tag{11}$$

or, summarizing the relevant equations,

$$m_k = f'^{-1}[-f(m_{k+1}) + m_{k+1} f'(m_{k+1})] \tag{12}$$

$$m_N = f'^{-1}(-\gamma) \tag{13}$$

$$M = \prod_{j=1}^{N} m_j \tag{14}$$

where we have denoted $f'^{-1}(x)$ as the inverse function of $f'(x)$.

However, to proceed further, $f(n)$ needs to be speci…ed. Although the recursion (12) starting from (13) can be computed formally, only a very small class of functions allows an exact solution in terms of M using (14). The class of arbitrary powers, $f(n) = n^{\alpha}$, is such an example.

Fortunately, just a power law complexity $O(n^{\alpha})$ is generally found to express complexity of tractable algorithms. In the next section, we will investigate in greater detail the case where $f(n) = n^2$ which corresponds to the complexity of the Dijkstra algorithm [1] that is known to be of the order[1] of M^2 where M are the number of nodes in the network.

We still did not prove that the solution of the Lagrange equations corresponds to a minimum. However, rather than embarking on a mathematical track to prove this, we derive below in sec. 3 the Lagrangian (with $\gamma = 0$) evaluated at the extremal solution. Comparison with the complexity of the original network clearly demonstrates that the extremal solution is indeed a minimum.

3 The complexity of Dijkstra's algorithm.

Although we concentrate in this section primarily on the n^2-complexity of Dijkstra's routing algorithm, the mathematical framework is general enough to treat an arbitrary power law complexity. In order not to obscure the understanding by tedious mathematical manipulations, we merely will state the results for $f(n) = n^{\alpha}$ (where $\alpha > 1$ and real) without detailing the derivation.

3.1 Analytical solution.

Using $f(n) = n^2$ in (11), we have

$$m_k = \frac{m_{k+1}^2}{2} \tag{15}$$

$$m_k \overset{\alpha}{=} \left(1 - \frac{1}{\alpha}\right)^{\frac{1}{\alpha-1}} m_{k+1}^{\frac{\alpha}{\alpha-1}} \tag{16}$$

which means that, in the optimal con…guration, the number of nodes on a higher hierarchical level k is half the square of the number of nodes of one peer group on a lower level $k+1$. Further,

[1] The order notation, the Landau big O notation, can be rigorously de…ned. A function $f(x) = O(g(x)$ for large x means that $\lim_{x \to \infty} \frac{f(x)}{g(x)} = 1$.

equation (15) exhibits self-similarity: precisely the same structure appears on all levels, however, scaled according to (15). This property stems from both the hierarchical PNNI topology and the polynomial complexity assumption. In addition, (15) is a simple difference equation with initial value given by (13). The solution for $k > 0$ is

$$m_{N-k} = \frac{\beta^{2^k}}{2^{2^{k+1}-1}} \tag{17}$$

$$m_{N-k} \overset{\alpha}{=} \left(1 - \frac{1}{\alpha}\right)^{\left(\frac{\alpha}{\alpha-1}\right)^{k-1}} \left(-\frac{\beta}{\alpha}\right)^{\frac{\alpha^k}{(\alpha-1)^{k+1}}} \tag{18}$$

We now turn to the elimination of the Lagrange multiplier β using (14), which we rewrite as $M = m_N \prod_{k=1}^{N-1} m_{N-k}$. Introducing (17) leads to

$$M = \left(-\frac{\beta}{2}\right)^{\sum_{k=1}^{N-1} \frac{\beta^{2^k}}{2^{2^{k+1}-1}}} = \left(-\frac{\beta}{2}\right)^{\frac{\beta^{\sum_{k=1}^{N-1}(2^{k+1}-1)}}{2^{\sum_{k=1}^{N-1}(2^{k+1}-1)}}} = \left(-\frac{\beta}{2}\right)^{\frac{\beta^{2^N-1}}{2^{2^N-2N}}}$$

Solving for β yields

$$\beta = -2^{\frac{2^N-2N}{2^N-1}} M^{\frac{1}{2^N-1}} \tag{19}$$

$$\left(-\frac{\beta}{\alpha}\right)^{\frac{\alpha^k}{(\alpha-1)^{k+1}}} \overset{\alpha}{=} \left[M\left(1-\frac{1}{\alpha}\right)^{\alpha\left(\frac{\alpha}{\alpha-1}\right)^{N-1}-1+N-1}\right]^{\frac{1}{\left(\frac{\alpha}{\alpha-1}\right)^{N-1}}} \tag{20}$$

Using (19) to specify the m_k leads to

$$m_N = 2^{-2^{-N}} M^{\frac{1}{2^N-1}} \tag{21}$$

$$m_{N-k} = 2^{-2^{-N}} M^{\frac{2^k}{2^N-1}} \tag{22}$$

$$m_N \overset{\alpha}{=} \left[M\left(1-\frac{1}{\alpha}\right)^{\alpha\left(\frac{\alpha}{\alpha-1}\right)^{N-1}-1+N-1}\right]^{\frac{\alpha}{(\alpha-1)\left(\frac{\alpha}{\alpha-1}\right)^{N-1}}} \tag{23}$$

$$m_{N-k} \overset{\alpha}{=} \left(1-\frac{1}{\alpha}\right)^{\left(\frac{\alpha}{\alpha-1}\right)^{k-1}} \left[M\left(1-\frac{1}{\alpha}\right)^{\alpha\left(\frac{\alpha}{\alpha-1}\right)^{N-1}-1+N-1}\right]^{\frac{\alpha^k}{(\alpha-1)^{k+1}\left(\frac{\alpha}{\alpha-1}\right)^{N-1}}} \tag{24}$$

Substituting these relations into (2) gives the minimal complexity we can achieve employing the PNNI hierarchical structure based on the complexity of Dijkstra's algorithm. After some algebraic manipulations, we arrive at

$$C_{pnni} = O\left(2M^{-2^{-N}} M^{\frac{1}{2^N-1}} \left(1 + \frac{(N-1)\beta}{2^{N-1}}\right)\right)$$

which roughly means that

$$C_{pnni} = O\left(M^{1+\frac{1}{2^N-1}}\right) \tag{25}$$

$$C_{pnni} \overset{\alpha}{=} O\left(M^{1+\frac{1}{\left(\frac{\alpha}{\alpha-1}\right)^{N-1}}}\right) \tag{26}$$

Compared to the original complexity which is $O(M^2)$, the optimal PNNI topology definitely results in huge savings. Specifically, the amount of saving can be estimated for $\alpha > 2$. Define $g(\alpha)$ as the exponent of M in (26). Then, we have

$$g(\alpha) = 1 + \frac{\frac{1}{N}}{\frac{\alpha}{\alpha - 1} - 1} = \frac{1}{1 - \left(1 - \frac{1}{\alpha}\right)^N}$$

The denominator can be rewritten as

$$1 - \left(1 - \frac{1}{\alpha}\right)^N = -\sum_{k=1}^{N} \binom{N}{k} (-1)^k (\alpha)^{-k} = \sum_{j=0}^{N-1} \binom{N}{N-j} (-1)^{N-j+1} (\alpha)^{-N+j}$$
$$= \alpha^{-N} \sum_{j=0}^{N-1} \binom{N}{j} (-1)^{N-j+1} \alpha^j$$

so that

$$g(\alpha) = \frac{\alpha^N}{N\alpha^{N-1} - \binom{N}{2}\alpha^{N-2} + \sum_{j=0}^{N-3} \binom{N}{j} (-1)^{N-j+1} \alpha^j} \tag{27}$$

For sufficiently large α, we observe from (27) that $g(\alpha) = O(\alpha/N)$. Due to the fact that the denominator is an alternating series, $g(\alpha)$ rapidly behaves linear in α. Indeed, numerical results indicate that this is already the case for $\alpha > 2$ and moderate values of N ($2 \leq N < 5$). This result means that the savings in complexity, $O(M^\alpha - M^{\alpha/N})$, increase with the complexity measure α and with the number of hierarchical levels N.

We are now sufficiently prepared to investigate how much hierarchical levels N are required. From the previous equations (21,22), we observe that the limiting behaviour where $N \to \infty$ results in the somewhat singular case where $m_k = 2$ ($m_k \stackrel{\alpha}{=} \frac{\alpha}{\alpha - 1}$) for all the infinitely many hierarchical levels. But then, the number of nodes M also tends to infinity. This 'optimal' limit situation learns that in each peer group (even on the lowest level) there must be at least 2 (or $\frac{\alpha}{\alpha - 1}$) nodes.

Consider now the situation where all lowest peer groups consist of x physical nodes. This means that on (21), we have

$$m_N = 2^{2^{-N}} M^{\frac{1}{2^N - 1}} = x$$

from which we can estimate N as a function of the total number of nodes. Indeed, this equation reduces to

$$2^N \log_2 \frac{x}{2} + N = \log_2 \left(\frac{x}{2} M\right) \tag{28}$$

However, no exact solution exists apart from infinite expansions. But, assuming that N is sufficiently large, we may neglect the linear term in N and solve for N resulting in

$$N < \log_2 \left(1 + \frac{\log_2 M}{\log_2 \frac{x}{2}}\right) \tag{29}$$

$$N \stackrel{\alpha}{<} \log_{\frac{\alpha}{\alpha - 1}} \left(1 + \frac{\log M}{(\alpha - 1)\log(1 - \frac{1}{\alpha})\, x}\right) \tag{30}$$

where we have written the inequality sign to account for the neglect of the linear term. Hence, we have found an upper bound for the number of hierarchical levels as a function of the total

number of nodes M and the number of physical nodes in each lowest level peer group. The assumption of large N is to some extent compensated by the integer nature of N. Even if we could solve the equation exactly, we would …nd a non-integer value of N, almost surely. Hence, con…nement to integers by taking the integral part would be necessary.

3.2 Numerical examples.

Some numerical examples may illustrate the above analysis. Suppose we have in each lowest level peer group x = 5 nodes and a total of $M = 10^{10}$ physical nodes in the original network. The upper bound (29) gives N < 4:7, thus N = 4. Suppose we use N = 5 in the exact equation (28), the left hand side amounts to 47.24 while for N = 4 we …nd for the left hand side 25.15. The number closest to the right hand side (which equals 34.54) is evidently the better, hence, N = 4. This example illustrates that the upper bound (29) is quite useful although the assumption of large N is seemingly not supported.

Let us consider the sensitivity of N. Using (21), we list below m_N and the order of the complexity C_N as a function of $M = 10^{10}$ and N:

N	=	1	m_1	=	10^{10}	C_1	= $O(10^{20})$
N	=	2	m_2	=	2714:42	C_2	= $O(2{:}15\ 10^{13})$
N	=	3	m_3	=	39:86	C_3	= $O(2{:}68\ 10^{11})$
N	=	4	m_4	=	7:71	C_4	= $O(4{:}64\ 10^{10})$
N	=	5	m_5	=	3:75	C_5	= $O(2{:}10\ 10^{10})$
N	=	6	m_6	=	2:69	C_6	= $O(1{:}44\ 10^{10})$
N	=	7	m_7	=	2:30	C_7	= $O(1{:}19\ 10^{10})$
N	=	8	m_8	=	2:14	C_8	= $O(1{:}09\ 10^{10})$
N	=	9	m_9	=	2:06	C_9	= $O(1{:}04\ 10^{10})$
N	=	10	m_{10}	=	2:03	C_{10}	= $O(1{:}02\ 10^{10})$

These numbers illustrate that m_N very rapidly tends to its limiting behaviour $m_N = 2$ as $N \to 1$. This supports the assumption of large N in (29). Moreover, we observe that the optimal restructuring of a complete network of $M = 10^{10}$ nodes is highly e¤ective only when using a relatively small number of hierarchical levels. The savings in complexity by de…ning more than N > 5 hierarchical levels is clearly not worth the e¤ort. Again, this is in favour of (29).

An illustration of the number of nodes on all levels as a function of N and k is given in the table below that lists the values of m_k for all $1 \quad k \quad N$ with $M = 10^{10}$.

k=N	1	2	3	4	5	6	7
1	10^{10}	3684030	315693	98212:3	48469:5	29038	18988:9
2		2714:4	794:6	443:2	311:3	241:0	194:9
3			39:9	29:8	24:9	21:9	19:7
4				7:7	7:1	6:6	6:3
5					3:7	3:6	3:5
6						2:7	2:7
7							2:3

4 Conclusions.

The optimal PNNI topology has been determined by minimizing the routing complexity. The results are speci…ed for Dijkstra's algorithm, although extensions to arbitrary polynomial complexity functions, $f(n) = n^{®}$ with real, positive ®, are straighforward. We have shown that the more complex the algorithm is (thus, the larger ®) the higher the saving we achieve with a hierarchical, optimal PNNI topology. The interest of this result is that in real PNNI implementations, the routing complexity will be likely higher than that of Dijkstra (® = 2), mainly due to the requirements for quality of service (QoS). Just in that case (® > 2), the PNNI hierarchical structure o¤ers more advantageous complexity savings.

The optimal topology is highly symmetrical and even self-similar, characterized by the fundamental equation (15). In this optimal PNNI topology, the number of peer group nodes on level k decreases by going to a lower level k + 1 whereas the number of peer groups on level k, given by $m_{k\,|\,1}$, just increases in that direction. This means that on the physical (or lowest) level, the peer group size is the smallest. Just at this level, QoS and tra¢c issues are to be considered in detail. The small size of the lowest level peer group substantially simpli…es the latter analysis even to such an extent that queueing analyses may be feasible. Successful dynamic control on the queueing level can further lead to overall optimality of the PNNI-protocol.

Only a very small number of hierarchical levels, $N = O(\log_2 \log_2 M)$, are needed in realistic network con…gurations. Earlier studies by Kleinrock and Kamoun[3] on general hierarchical structures found $N = O(\ln M)$ by optimizing the length of the routing table.

In spite of the enthusiasm, some caution would seem to be appropriate. The analysis only relies on the number of nodes, discarding the di¤erences in links, neglecting QOS-related parameters or supplementary nodal information that may impact on the topology (such as node and links aggregation) and omitting the in‡uences of the ‡ooding mechanism. However, ‡ooding is not believed to be problematic since the basic peer group is quite small implying fast convergence. Admitting the simpli…cations, we believe this study suggests useful design rules for a hierarchical PNNI structuring of the topology.

Acknowledgements.

We are grateful to Y. T'Joens and M. Vandenhoute for their useful comments.

References

[1] Wai-Kai Chen. Theory of Nets: Flows in Networks. Wiley & Sons, New York, 1990.

[2] The ATM Forum. Private Network-Network Interface. af-pnni-0055.000, pnni 1.0 edition, March 1996.

[3] L. Kleinrock and F. Kamoun. Hierarchical routing for large networks: Performance evaluation and optimization. Computer Networks, 1:155–174, 1977.