



On Hierarchical Extensions of Large-Scale 4-regular Grid Network Structures

Pedersen, Jens Myrup; Patel, A.; Knudsen, Thomas Phillip; Madsen, Ole Brun

Publication date:
2004

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Pedersen, J. M., Patel, A., Knudsen, T. P., & Madsen, O. B. (2004). *On Hierarchical Extensions of Large-Scale 4-regular Grid Network Structures*. <Forlag uden navn>.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

On Hierarchical Extensions of Large-Scale 4-regular Grid Network Structures

Jens Myrup Pedersen*, Ahmed Patel**, Thomas Phillip Knudsen*, Ole Brun Madsen*

* *Department of Control Engineering, Aalborg University, DK-9220, Aalborg East, Denmark*
Tel. +45 96358700, Fax +45 98151739, mail: {jens,thomas,obm}@control.aau.dk

***Computer Networks & Distributed Systems Research Group*
Department of Computer Science, University College Dublin, Belfield, Dublin 4, Ireland
Tel. +353-1-7162476, Fax +353-1-2697262, mail: apatel@cnds.ucd.ie

Abstract

It is studied how the introduction of ordered hierarchies in 4-regular grid network structures decreases distances remarkably, while at the same time allowing for simple topological routing schemes. Both meshes and tori are considered; In both cases non-hierarchical structures have power-law dependencies between the number of nodes and the distances in the structures. The perfect square mesh is introduced for hierarchies, and it is shown that applying ordered hierarchies in this way results in logarithmic dependencies between the number of nodes and the distances, resulting in better scaling structures. For example, in a mesh of 391876 nodes the average distance is reduced from 417.33 to 17.32 by adding hierarchical lines. This is gained by increasing the number of lines by 4.20% compared to the non-hierarchical structure. A similar hierarchical extension of the torus structure also results in logarithmic dependencies, the relative difference between performance of mesh and torus structures being less significant than for non-hierarchical structures, especially for large structures. The skew and extended meshes are introduced as variants of the perfect square mesh and their performances studied, and it is shown that while they allow for more flexibility in design and construction of structures supporting topological routing, their performances are comparable to the performance of the perfect square mesh. Finally suggestions for further research within the field is given.

Keywords: Topological Routing, Routing, Network Structures, Large-Scale Networks, Network Planning

1 Introduction

Routing schemes applied in large-scale networks such as the Internet have until recently operated on a best-effort

basis only. This has been sufficient for many applications such as email, news, web browsing and file transfers, but the needs for reliable connections are increasing as the dependency on communication networks grows: As a consequence of the convergence of communications, several different medias are becoming able to communicate over the same physical lines such as telephony, television, video and data traffic. At the same time an increasing amount of control and distributed applications are being developed to communicate over the Internet, such as the many teleoperation and telerobotics projects, e.g. [1][2][3]. This convergence of communications is expected to continue in near as well as far future[4].

Routing schemes based on large tables are under severe pressure in today's Internet. In 2001 experiments showed that it took on average three minutes to recover from inter-domain path failovers. For some path failovers it took up to 15 minutes before the routing tables were stabilised[5]. The dependency on exact and updated tables is a major problem for the routing schemes, since the size of the Internet makes it impossible to maintain updated tables of the complete Internet topology; Not only would this require huge tables, it would also be extremely bandwidth and resource demanding to keep them updated. Furthermore, the provisioning of two or more physical independent end-to-end paths is necessary in order to provide connections that are reliable even in case of equipment failures, broken cables etc, and even if routing tables are exact and updated, this is an extremely difficult task in today's complex Internet topology.

Reducing the needs for large tables while at the same time offering easy ways to determine several independent paths, can be done by designing networks with global structural and topological properties such as the 4-regular grid network structure[6] (which is actually not 4-regular since

edge and corner nodes have degree lower than four). It is suggested a base for a future access network infrastructure, making the study of it highly actual: Many countries are expected to let fiber networks replace the existing copper-based infrastructures in near future, and for these infrastructures to benefit from global structural and topological properties it must be well planned. Changing the physical structure after implementation is costly due to the high duct costs, and the physical structures are expected to have long lifetimes as the networks are upgradeable by changing end equipment only. This paper shows how hierarchical extensions of the 4-regular grid network structure can be used for reducing distances, especially in large-scale structures.

2 Terminology

Abstractions of networks, called structures, are studied. A structure consists of a set of nodes and a set of lines, such that each line connects two nodes. Lines are considered undirected: If a pair of nodes (u, v) is connected by a line, so is (v, u) . A path between a source node u and a destination node v is a set of nodes and lines $(u = u_0), e_1, u_1, e_2, u_2, e_{n-1}, u_{n-1}, e_n, (u_n = v)$, where each line e_i connects the nodes u_{i-1} and u_i . The path length is determined by the number of lines between the source and destination node; In the case above, the path is of length n . The distance between two nodes u and v is written $d(u, v)$ and is equal to the length of the shortest path between u and v . Note that $d(u, v) = 0$ if and only if $u = v$. The size of a structure is equal to the number of nodes it contains. The degree of a node is the number of lines joined to it.

3 Background

Structural QoS[4] and Sustainable QoS[7] were introduced recently and deal with parameters related to architecture and structural properties of networks. This motivates an approach of designing physical network infrastructures that are able to support reliability and QoS demanding applications. [6] investigates an approach of designing network structures which by their high degree of regularity support a simple and essentially table-free routing scheme known as topological routing. It is shown how topological routing can solve a large number of the problems faced in today's routing schemes, but the networks must be well structured and organized for such a scheme to work. It suggests that the 4-regular grid structure is used for topological routing: The nodes are addressed from a cartesian coordinate system, such that every node has a coordinate address (x, y) associated to it: Let dim_x and dim_y be positive integers. Then every integer coordinate set (x, y) , such that $0 \leq x \leq dim_x$ and $0 \leq y \leq dim_y$, is the address of a node. Every node has associated to it such an address, and no two nodes have the

same address. Two variants are used, the mesh and torus. In the mesh, a line between two nodes (x_1, y_1) and (x_2, y_2) exists if and only if $|x_2 - x_1| + |y_2 - y_1| = 1$, in the torus additional lines exist connecting nodes where either $|x_2 - x_1| = dim_x$ and $y_1 = y_2$ or $|y_2 - y_1| = dim_y$ and $x_1 = x_2$.

A packet is routed from source to destination in such a structure on a hop-by-hop basis: In every node it is forwarded to a node with the smallest possible distance to the destination, which is easily determined. The packet only needs to carry its destination address, and every node only needs to know the addresses of its neighbours. In the torus, every node must also know dim_x and dim_y .

While such structures are used for multiprocessor systems, and much research in this area has dealt with them (e.g. [8][9]), they suffer from severe scalability problems in order to be used for large-scale networks: There is a power-law dependency between the number of nodes and distances in the structures. For example the average distance in a square mesh structure of 10000 nodes is 66.67. This should be compared to the following: In 1998 average path lengths of the Internet was measured, and it was shown that the Internet at that time had an average path length of around 11-24, depending on location[10]. In 1999 appr. 88000 different nodes (routers) were found in the Internet[11].

Hierarchical extensions of the 4-regular grid structure were introduced in [6]: A set of hierarchical lines is added to a structure, such that a revised topological routing scheme is supported, while at the same time the distances are shortened. The revision of the routing scheme is based on the fact that deciding whether routing should be done in higher hierarchies is done in each node, provided knowledge of a few global parameters. If routing is to be done through a higher hierarchical layer, a shortest path can always be found using the closest higher hierarchy node, which is easily determined. When the highest layer to be used is reached by the packet, the basic routing scheme is used. As a result, a shortest path between any two nodes is easily determined, taking the hierarchical lines into account.

In this paper such a way of constructing hierarchical lines is evaluated by studying how the addition of such lines decreases distances.

4 Methods

Structures are evaluated by calculating and comparing distances. Let u_1, \dots, u_n be the nodes of a structure. Then for every node u_j , the average distance from u_j to all other

nodes is calculated as $d_{avg}(u_j) = \frac{\sum_{i: i \neq j, 1 \leq i \leq n} d(u_i, u_j)}{n-1}$. Three measures are used for evaluation: Average distance, worst-case average distance and diameter. The average distance




Figure 3. Avg. distances in hierarchical torus.

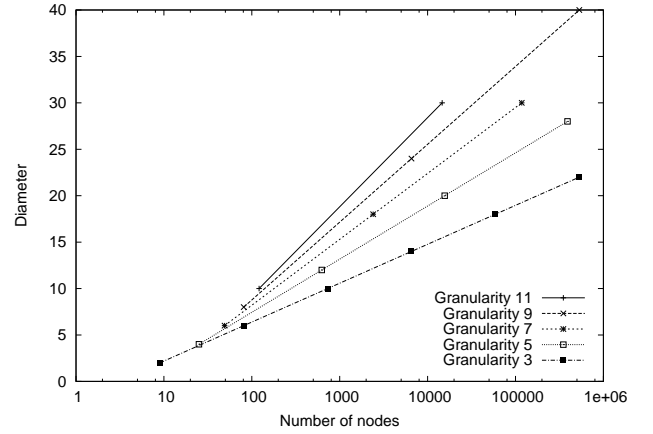


Figure 5. Diameters in hierarchical torus.




Figure 4. Worst-case average distances in hierarchical torus.




Figure 6. Average distances in perfect square mesh.

and they all show virtually perfect logarithmic dependencies between structure sizes and the various distance measures. Second, the diameter does clearly increase linear with n_H , and therefore close to logarithmic with the number of nodes. This is true for both the hierarchical torus, the perfect square mesh and the variants for all values of g_x and g_y . Some measurements of very small structures seem to differ slightly from the rest, but all structures soon approximate a logarithmic dependency.

Fig. 2 shows how the non-hierarchical approach results in power-law dependencies between the number of nodes and distance measures. Even though distances are shorter in the torus than in the mesh, none of the structures scale well, and in large structures distances are huge.

These results should be compared to the distances in the hierarchical torus shown in fig. 3-5 and perfect square mesh shown in fig. 6-8. The logarithmic dependencies between the number of nodes and distances is interesting, but it is also worth noting that the distances are kept reasonable low, even for very large structures. For example, the hierarchical

torus with $g = 5$ and $n_H = 3$ is a structure of size 15625 with average distance 11.76, worst-case average distance 14.70 and diameter 20. The non-hierarchical torus of same size has average distance and worst-case average distance 62.50 and diameter 124. This gain is obtained by adding 1300 lines to the original 31250, an increase of 4.16%.

The perfect square mesh shows the same pattern as the hierarchical torus. Among the non-hierarchical structures, the torus performs considerably better than the mesh, but for the hierarchical extensions the relative differences are smaller, especially for large structures.

The perfect square mesh with $g = 5$ and $n_H = 3$ contains 15876 nodes with average distance 12.48, worst-case average distance 15.96 and diameter 22. A square mesh structure of this size, without the hierarchical extension, has corresponding distances 84.00, 125.01 and 250. The number of lines is 32864, while it is 31500 in the non-hierarchical structure, an increase of 4.33% or 0.086 lines per node. $g = 5$ and $n_H = 4$ gives a perfect square mesh of size 391876. In this case, average distance, worst-




Figure 7. Worst-case average distances in perfect square mesh.




Figure 9. Skew mesh with granularity 3/11 and perfect square mesh with granularity 9.




Figure 8. Diameters in perfect square mesh.

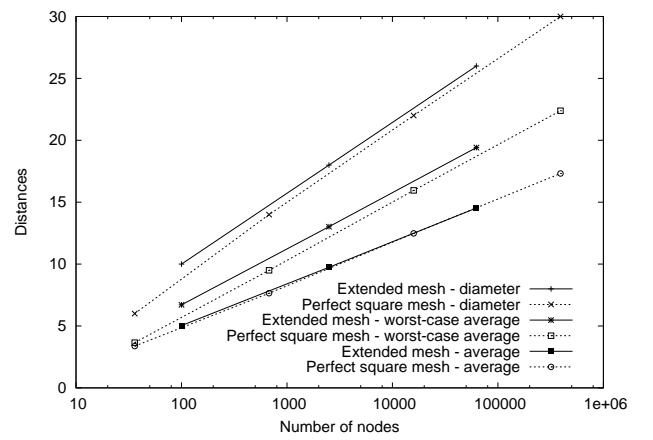


Figure 10. Extended mesh compared to perfect square mesh, both with granularity 5.

case average distance and diameter are 17.32, 22.39 and 30 respectively. These should be compared to the non-hierarchical structure of the same size, with corresponding distances 417.33, 625.00 and 1250 respectively. The number of lines in the non-hierarchical structure is 782500, while it is 815364 in the hierarchical extension, an increase of 4.20%, or 0.084 lines per node.

As expected, the smaller g is chosen, the shorter the distances are. On the other hand, for small choices of g more lines are added. $g = 3$, the smallest possible, and $n_H = 5$ gives a structure of size 59356, with average distance, worst-case average distance and diameter 12.21, 15.13 and 20 respectively. 15004 lines are added to the non-hierarchical structure of 118584 lines, an increase of 12.65%, or 0.25 lines per node. Without hierarchies the corresponding distances are 162.67, 243.00 and 486.

Fig. 9 and fig. 10 show the performances of the variants of the perfect square mesh, compared to the perfect square mesh with performance closest to those of the variants.

The skew mesh with $g_x = 3$ and $g_y = 11$ performs

similarly to the perfect square mesh with $g = 9$ in terms of both average distance, worst-case average distance and diameter. The skew mesh does, however, require a larger number of hierarchical lines. Since no structures of comparable size exist this can be seen by comparing the number of hierarchical lines relative to the number of nodes. The perfect square mesh of size 6724 and 532900 have 0.027 respectively 0.025 hierarchical lines per node, while the skew mesh of size 37296 has 0.064 hierarchical lines per node.

The extended mesh with $g = 5$ performs comparable to the perfect square mesh with $g = 5$, with only slightly larger diameters and worst-case average distances. All distance measures are approximating those of the perfect square mesh for large numbers of nodes. As with the skew mesh, it is not possible to compare two structures of the same size, because only structures of different sizes are supported. Since the granularity is the same for both variants, the number of hierarchical lines are close to each other, being only a bit smaller for the extended mesh. The extended mesh of size 62500 has 0.081 hierarchical lines per node.

6 Discussion and conclusion

It was shown that ordered hierarchies as introduced are useful for reducing distances in 4-regular grid network structures significantly, especially for large-scale structures. The main model used was the perfect square mesh and its torus counterpart. While there is a power-law dependency between the number of nodes and average as well as worst-case average distances in the non-hierarchical structures, the corresponding dependencies for the hierarchical structures are logarithmic and leads to noticeable smaller distances. For example, the perfect square mesh with $g = 5$ and $n_H = 4$ is a structure of size 391876 with average distance only 17.32, compared to an average distance of 417.33 for a similar sized mesh structure without hierarchies. The performance difference between mesh and torus, which is noticeable for non-hierarchical structures, is considerably smaller for the hierarchical structures, especially for large structures. Even for small values of g the number of hierarchical lines added to the structure is quite small compared to the total number of lines. Even though these hierarchical lines are longer than the non-hierarchical lines in the sense that they connect nodes with larger distances in the basic structure, they can follow the same ducts, minimizing the costs. The small number of hierarchical lines implies a similar limited increase in the node degrees, which is another important cost factor.

Two variants were introduced as alternatives to the perfect square mesh, the skew mesh and the extended mesh. As both show logarithmic dependencies between size and distances, they can be used in concrete network planning when no suitable perfect square mesh exists.

However, the skew mesh has a drawback since more hierarchical lines are used for this structure than for a perfect square mesh of the same size and performance. An alternative to this structure, which may prove to be better, is obtained by placing a number of perfect square mesh structures adjacent to each other, connected at the edges. This allows for more different shapes than the skew mesh, but further research is required for evaluation of such an approach.

On the other hand, the extended mesh has performance almost as good as the perfect square mesh, using even slightly fewer hierarchical lines per node. It may be useful for constructing structures of sizes different from what is supported by the perfect square mesh. It may also be possible to use it as a base for constructing cheaper and better performing hierarchical structures than the perfect square mesh, where more nodes are placed closer to higher hierarchical nodes, reducing especially diameters and worst-case distances.

With this contribution, the 4-regular grid structures have been shown to form a suitable base for fiber-based access

networks, offering high connectivity and reasonably small distances. However, efficient protection and restoration schemes still have to be developed for hierarchical extensions of the structures.

In order to create a cost-efficient alternative to the ring structures widely used today it should also be considered if the number of lines, and consequently the average node degree, could be reduced. As the number of hierarchical lines is small compared to the total number of lines even for small values of g , focus should be on the lines of the basic structure. A scheme reducing this number, and thus the connectivity, while maintaining the support of topological routing and relatively short distances would be a major contribution here.

Acknowledgements

The authors would like to thank University College Dublin for hosting J. M. Pedersen and T. P. Knudsen as visiting researchers during the work of this paper.

References

- [1] Kuk-Hyun Han, Yong-Jae Kim, Jong-Hwan Kim, and Steve Hsia. Internet control architecture for internet-based personal robot. *Autonomous Robots*, 10, Issue 2:135–147, 2001.
- [2] N. Xi and T. J. Tarn. Planning and control of internet-based teleoperation. *Proceedings of the SPIE - The International Society for Optical Engineering*, 3524:189–195, 1998.
- [3] Li Xiao-ming, Yang Can-jun, Chen Ying, and Hu Xu-dong. Hybrid event based control architecture for tele-robotic systems controlled through internet. *Journal of Zhejiang University SCIENCE*, 2004 5(3), pages 296–302, 2004.
- [4] Ole Brun Madsen, Jens Dalsgaard Nielsen, and Henrik Schiøler. Convergence. *RTLIA 2002*, Wien, 2002.
- [5] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. Delayed internet routing convergence. *IEEE/ACM Transactions On Networking*, 9(3):293–306, June 2001.
- [6] Jens Myrup Pedersen, Thomas Phillip Knudsen, and Ole Brun Madsen. Topological routing in large-scale networks. *ICACT 2004*, Korea, 2004.
- [7] K. Murakami and H. S. Kim. Virtual path routing for survivable atm networks. *IEEE/ACM Transactions On Networking*, February 1996.
- [8] Jie Wu. A fault-tolerant and deadlock-free routing protocol in 2d meshes based on odd-even turn model. *IEEE Transactions on Computers*, 39(9), September 2003.
- [9] Ming-Jer Tsai and Sheng-De Wang. Adaptive and deadlock-free routing for irregular faulty patterns in mesh multicomputers. *IEEE Transactions in Parallel and Distributed Systems*, 11(1), January 2000.
- [10] Aiguo Fei, Guangyu Pei, Roy Liu, and Lixia Zhang. Measurements on delay and hop-count of the internet. *Proceedings of Globecom 1998*, Sydney, Australia, 1998.
- [11] Hal Burch and Bill Cheswick. Mapping the internet. *IEEE Computer*, 32, Issue 4:97–98,102, 1999.