Distrib. Comput. (2001) 14: 65–73

DISTRIBUTED COMPUTING

ATM layouts with bounded hop count and congestion

Michele Flammini[1], **Enrico Nardelli**[2], **Guido Proietti**[2]

[1] Dipartimento di Matematica Pura ed Applicata, University of L'Aquila, Via Vetoio loc. Coppito, 67100 L'Aquila, Italy (e-mail: flammini@univaq.it)
[2] Dipartimento di Matematica Pura ed Applicata, University of L'Aquila, Via Vetoio loc. Coppito, 67100 L'Aquila, Italy
[3] Istituto di Analisi dei Sistemi ed Informatica, Consiglio Nazionale delle Ricerche, Viale Manzoni 30, 00185 Roma, Italy (e-mail: {nardelli,proietti}@univaq.it)

Received: December 1997 / Accepted: August 2000

Summary. In this paper we introduce and analyze two new cost measures related to the communication overhead and the space requirements associated with virtual path layouts in ATM networks, that is *the edge congestion* and the *node congestion*. Informally, the edge congestion of a given edge e at an incident node u is defined as the number of VPs terminating at or starting from u and using e, while the node congestion of a node v is defined as the number of VPs having v as an endpoint. We investigate the problem of constructing virtual path layouts allowing to connect a specified root node to all the others in at most h hops and with maximum edge or node congestion c, for two given integers h and c. We first give tight results concerning the time complexity of the construction of such layouts for both the two congestion measures, that is we exactly determine all the tractable and intractable cases. Then, we provide some combinatorial bounds for arbitrary networks, together with optimal layouts for specific topologies such as chains, rings and grids.

Key words: Routing – ATM networks – Computational complexity – Edge and node congestion

1 Introduction

The *Asynchronous Transfer Mode* (ATM for short) is the most popular networking paradigm for Broadband ISDN [13,12,15]. It transfers data in the form of small fixed-size *cells*, and in order to achieve the stringent transfer rate requirements, is based on two types of predetermined routes in the network: *virtual paths* or VPs, constituted by a sequence of successive edges or physical links, and *virtual channels* or VCs, each given by the concatenation of a proper sequence of VPs. Routing in virtual paths can be performed very efficiently by dedicated hardware, while a cell passing from one virtual path to another one requires more complex and slower elaboration.

Given a network and a set of connections to be established, in order to provide the performance required by B-ISDN applications it is important that routing is performed in a hardware fashion in most of the nodes a cell traverses, at the same time limiting the number of virtual paths sharing a same physical link [16,3,17,1,11].

A graph theoretical model related to this ATM design problem has been first proposed in [11,4]. In such a framework, the VP layouts determined by the VPs constructed on the network are evaluated mainly with respect to two different cost measures: the *hop count*, that is the maximum number of VPs belonging to a VC, which represents the number of VP changes of cells along their route to the destination, and the *load*, given by the maximum number of virtual paths sharing an edge, that determines the size of the VP routing tables (see, e.g., [5]). Another relevant parameter is the *stretch factor*, i.e., the ratio between the length of the path that a VC takes in the physical graph and the shortest possible path between its endpoints. This parameter controls the efficiency of the utilization of the network. For further details and technical justifications of the model for ATM networks see for instance [1,11].

While the problem of determining VP layouts with bounded hop count and load is NP-hard under different assumptions [11,6], many optimal and nearly optimal constructions have been given for various interconnection networks such as chain, trees, grids and so forth [4,14,9,10,18,2]. A more detailed list of these results can be found in [19].

In this paper we introduce and analyze two new cost measures associated to virtual path layouts: the *edge congestion*, which is given by the maximum number of VPs terminating or starting from a given edge at a given node, and the *node congestion*, that is the maximum number of VPs having as an endpoint a given node. The main motivation behind these new cost measures is due to the fact that while along a VP routing is performed at hardware level with negligible delay, at the end of the VP a non hardware route selection must be taken at considerably higher cost. Even if in absence of resources contention due to other cells delays are proportional to the number of hops or traversed VPs, under the assumption of moderate till bursty traffic along each VP, the effect of non hardware routing decisions can seriously affect the delivering time if the number of cells at the terminals of the VPs is not kept low

Work supported by the EU TMR Research Training Grant N. ERBFM-BICT960861 and by the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT)

or balanced. If an autonomous routing capability is given to the input channels of ATM switches, this translates in a direct way to devising VP layouts with a low edge congestion, while if each switch is in charge of routing in a centralized way all the cells arriving from all its input channels, then the node congestion becomes the relevant parameter to be minimized. Another important effect of the edge and node congestions is that they directly influence the dimension of VC tables, as any VP that increases the congestion of an edge or of a node causes a number of entries in the corresponding VC table equal to the number of VCs such a VP belongs to. Finally, the node congestion allows to model the requirement of bounded degreeness of the nodes, that is a real constraint in many practical cases.

Although the new congestion measures are not completely unrelated with the edge load of [11,4], there are fundamental differences that in general make the results provided for the edge load not comparable to the ones for the edge and node congestions (see Sect. 3). For instance, while a layout with edge load l has also edge congestion $c \leq l$, the reverse in general is not true and one can find layouts with small edge congestion ($c \leq 2$) and edge load linear in the number of nodes ($l \geq N/2$). Moreover, as will be shown in the sequel, further differences hold when dealing with an unbounded stretch factor, like in the results for grids provided in [14,2]. Similar considerations apply also to the node congestion and as a consequence new layouts and methods are needed in order to achieve optimal solutions.

As in [11] and [6], in this paper we focus on layouts that enable the routing between all nodes and a single root node (rather than between any pair of nodes), under the assumption of a stretch factor equal to one, that is all the physical routed paths are the shortest. In fact, this restricted case can be seen as a building block for more complex routing problems and nevertheless its simplicity has not been fully understood yet. After a comparison with the existing performance measures and some general properties and results, we give tight results on the time complexity of constructing optimal rooted virtual path layouts. We then provide some optimal layouts for specific networks, such as chains, rings and grids.

As a comparison with the previous edge load results for grids, while our layouts yield an edge congestion $c < \frac{\sqrt[h]{N}}{2}$, where N is the number of nodes, in [2] it has been shown that any layout with an unbounded stretch factor requires an edge load $l = \Omega(\sqrt[h]{\frac{N^2}{h}})$, while it is possible to achieve $l \leq h\sqrt[h]{N^2}$. Therefore, under the reasonable assumption that $h = o(N)$, our edge congestion c is below the lower bound in [2], even if the stretch factor is equal to one. In fact, our layouts in general have an edge load l higher than c. Similar arguments apply also to the node congestion, thus giving a more precise evidence of how the different cost measures can be highly unrelated.

The paper is organized as follows. In Sect. 2 we define the preliminary notation and definitions. In Sect. 3 we discuss the relationship between the new cost measures and the previous parameters, together with some basic results. In Sect. 4 we provide the above-mentioned time complexity results. In Sect. 5 we present the optimal layouts for specific topologies and finally, in Sect. 6, we give some concluding remarks and list some open problems.

2 Preliminaries

We model the network as an undirected graph $G = (V, E)$, where nodes in V represent switches and edges in E are the point-to-point physical communication links.

Definition 2.1 *[11] A* rooted virtual path layout *(or simply layout) Ψ is a collection of paths in G, termed* virtual paths *(VPs for short), and a node $r \in V$, termed the* root *of the layout.*

Definition 2.2 *[11] The* hop count $\mathcal{H}(v)$ *of a node $v \in V$ in a layout Ψ is the minimum number of VPs whose concatenation forms a shortest path in G from v to r. If no such VPs exist, define $\mathcal{H}(v) \equiv \infty$. The* maximal hop count *of a layout Ψ is $\mathcal{H}_{\max}(\Psi) \equiv \max_{v \in V}\{\mathcal{H}(v)\}$.*

Given $v \in V$, let us denote as $I(v)$ the set of the edges in E incident to v.

Definition 2.3 *Given $v \in V$ and $e \in I(v)$, the* edge congestion $\mathcal{E}(e, v)$ *of the edge e with respect to v in a layout Ψ is the number of VPs $\psi \in \Psi$ that include e and have v as an endpoint. The* maximal edge congestion $\mathcal{E}_{\max}(\Psi)$ *of a layout Ψ is $\max_{v \in V, e \in I(v)} \mathcal{E}(e, v)$.*

A layout Ψ with $\mathcal{H}_{\max}(\Psi) \leq h$ and $\mathcal{E}_{\max}(\Psi) \leq c$ is called a *$\langle h, c\rangle$-edge layout.*

At each node of the network, a more global congestion measure can be considered which takes into account the total cost required at the node.

Definition 2.4 *Given $v \in V$, the* node congestion $\mathcal{N}(v)$ *of v in a layout Ψ is the number of VPs $\psi \in \Psi$ such that v is an endpoint of ψ. The* maximal node congestion $\mathcal{N}_{\max}(\Psi)$ *of a layout Ψ is $\max_{v \in V} \mathcal{N}(v)$.*

A layout Ψ with $\mathcal{H}_{\max}(\Psi) \leq h$ and $\mathcal{N}_{\max}(\Psi) \leq c$ is called a *$\langle h, c\rangle$-node layout.*

Clearly, the hop count and the edge (or node) congestion are conflicting parameters, as in general a low hop count requires an high congestion and a low congestion causes a high hop count. Thus, a very natural problem arises in which one parameter is traded off for the other. Moreover, once fixed two bounds h and c respectively on the hop count and on the edge (or node) congestion, in a parametric family of graphs it makes sense to consider the problem of determining the highest order graph that admits a layout respecting such bounds.

Definition 2.5 *Let $\mathcal{G}$ be a family of graphs. For any two positive integers h and c, $E_{\mathcal{G}}(h, c)$ (resp. $N_{\mathcal{G}}(h, c)$) is defined as the maximum integer N such that there exists an N-node graph in $\mathcal{G}$ with a $\langle h, c\rangle$-edge layout (resp. a $\langle h, c\rangle$-node layout).*

For the sake of brevity, when clear from the context, we will denote $E_{\mathcal{G}}(h, c)$ and $N_{\mathcal{G}}(h, c)$ respectively as $E(h, c)$ and $N(h, c)$.

Notice that all the above definitions assume a stretch factor equal to one, i.e., all the physical routed paths are the shortest.

3 Comparisons and basic properties

The congestion measures introduced in the previous section and the edge load defined in [11] (maximum number of VPs that share a physical edge) are not completely unrelated. In fact, as it can be easily verified, an edge load l implies an edge congestion at most l and a node congestion at most equal $l \cdot \delta$, where δ is the maximum node degree.

On the other hand, apart from this relationship, it seems that there is no strong connection between these parameters. For instance, small edge and node congestions do not necessarily imply a small load, as one can easily find VP layouts with constant (edge or node) congestion and load linear in the number of nodes. As an example, consider a VP layout Ψ for a chain of N nodes $1, \ldots, N$ with VPs $\langle i, i+1 \rangle$, $1 \leq i < N$, and $\langle i, N-i \rangle$, $1 \leq i \leq N/2$. Then Ψ has $\mathcal{E}_{\max}(\Psi) = 2$, $\mathcal{N}_{\max}(\Psi) = 3$ and edge load $\lfloor N/2 \rfloor$.

Another basic difference is that for the new cost measures it does not make sense to consider layouts with unbounded physical routed lengths. In fact, optimal layout constructions for the node congestion case can always be determined when the physical routed length is unbounded as follows. Consider any ordering of the nodes, except the root. Then, the root reaches through a VP in one hop the first c nodes, and iteratively in the order each reached node is assigned a VP to all the next $(c-1)$ unreached nodes. This always gives $N(h,c) = \frac{c(c-1)^h - 2}{c-2}$ in a straightforward way. In the edge congestion case the construction is slightly more complicated, because nodes have to be ordered non increasingly with respect to their degrees. Since to the purpose of minimizing the edge congestion VPs have not necessarily to correspond to simple physical paths, at every node the incident VPs can be equally distributed among its incident edges. Thus an optimal layout can be easily determined.

Optimal layouts can be easily found, still assuming unbounded physical routed lengths, even in the all-to-all case in which, by respecting the bounds on the edge or node congestion, each node wants to reach every other node in at most a given number of hops. Here the construction becomes a pure combinatorial graph design problem. In fact, if the node congestion is bounded by c, there is a layout for a graph G within a given hop count h if and only if there exists a c-bounded degree graph with diameter h and the same number of nodes of G. Any embedding of such a graph on G gives the desired layout. A similar argument holds for the edge congestion, but here there is a layout respecting h and c if there exists a graph with the same number of nodes, diameter h and such that, if we denote as d_i the degree of node i in the initial graph G, the i-th node of the graph has degree at most $c \cdot d_i$.

We conclude the section with a simple counting argument that allows to establish upper bounds on the number of nodes in networks admitting $\langle h, c \rangle$-edge layouts or $\langle h, c \rangle$-node layouts.

Given a graph G with a specified root node r, we say that a non root node u has *branch parameter* d if it has exactly d incident edges $\{u, v_1\}, \ldots, \{u, v_d\}$ such that for each i, $1 \leq i \leq d$, u is on a shortest path from r to v_i. Let the branch parameter of a family of graphs $\mathcal{G}$ be the maximum branch parameter of a non root node of a graph in $\mathcal{G}$. Then, if d_r is the degree of the root r, in any layout with edge congestion c from r it is possible to reach in one hop at most $c \cdot d_r$ nodes; each node with hop count 1 can then reach in another hop at most cd other nodes (for a total of $cd_r(cd)$ nodes), and this holds for every node with hop count at least 1, since the physical routed paths have to be the shortest and thus each node can use at most d outgoing edges to reach other nodes. Hence, $E(h,c) \leq 1 + cd_r + cd_r(cd) + cd_r(cd)^2 + \ldots + cd_r(cd)^{h-1}$ and the following fact is proved.

Fact 3.1 *Let $\mathcal{G}$ a family of graphs with branch parameter d and root r of degree at most d_r. Then $E(h,c) \leq 1 + cd_r \frac{(cd)^h - 1}{cd - 1}$.*

Similarly, in a layout with node congestion c, starting from r it is possible to reach in one hop at most c nodes; each node with hop count 1 can then reach in another hop at most $c-1$ other nodes, and this holds for every node with hop count at least 1, since the VP through which a node is reached contributes 1 to its node congestion. Therefore, $N(h,c) \leq 1 + c + c(c-1) + c(c-1)^2 + \ldots + c(c-1)^{h-1}$ and the following fact holds.

Fact 3.2 *For any family of graphs $\mathcal{G}$, $N(h,c) \leq \frac{c(c-1)^h - 2}{c-2}$.*

As we will see in Sect. 5, nevertheless their conceptual simplicity, Facts 3.1 and 3.2 allow to establish tight upper bounds in all the considered topologies.

4 Time complexity results

In this section we show that constructing optimal layouts is in general an NP-hard problem for both the two congestion measures. We first show that deciding the existence of an $\langle h, c \rangle$-edge layout is an NP-complete problem, even for $h = 3$ and $c = 1$.

Theorem 4.1 *Given a network $G = (V, E)$ and a root $r \in V$, deciding the existence of a $\langle 3, 1 \rangle$-edge layout for G with root r is an NP-complete problem.*

Proof First of all, observe that for any h and c the problem of deciding the existence of a $\langle h, c \rangle$-edge layout is in NP, as given a layout Ψ for $G = (V, E)$ with a given root $r \in V$, one can easily check whether $\mathcal{E}(e, v) \leq c$ for every node v and incident edge $e \in E$ and whether $\mathcal{H}_{\max}(\Psi) \leq h$ (see [6]).

We prove the claim by providing a polynomial time reduction from *3-SAT* (known to be NP-complete; see [8]). An instance of this problem is constituted by a boolean formula f over m variables $x_1, \ldots, x_m$, where f is in conjunctive normal form, i.e., f is the conjunction of g clauses $c_1, \ldots, c_g$, each of which is the disjunction of three literals. We want to determine whether there exists a truth assignment for $x_1, \ldots, x_m$ which satisfies f.

Starting from an instance of 3-SAT, we construct a graph G that admits a $\langle 3, 1 \rangle$-edge layout if and only if f is satisfiable.

Let $G = (V, E)$, where $V = \{r\} \cup V_1 \cup V_2 \cup V_3 \cup V_4$, and $E = E_1 \cup E_2 \cup E_3 \cup E_4 \cup E_5$ (see Fig. 1), with:

$V_1 = \{\overline{u}_a, u_a \mid a = 1, \ldots, m\}$, $V_2 = \{v_a \mid a = 1, \ldots, m\}$,
$V_3 = \{\overline{q}_a, q_a \mid a = 1, \ldots, m\}$, $V_4 = \{z_b \mid b = 1, \ldots, g\}$,

and

$E_1 = \{\{r, \overline{u}_a\}, \{r, u_a\} \mid a = 1, \ldots, m\}$,
$E_2 = \{\{\overline{u}_a, v_a\}, \{u_a, v_a\} \mid a = 1, \ldots, m\}$,

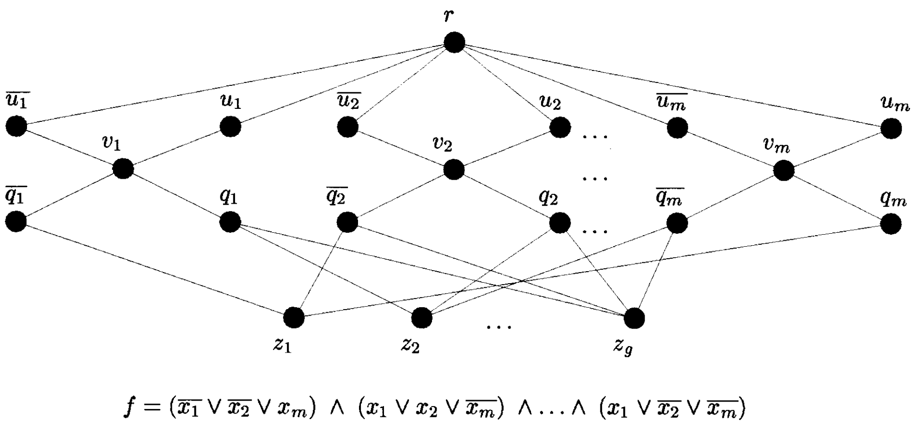

Fig. 1. The reduction graph for the $\langle 3,1\rangle$-edge layout problem

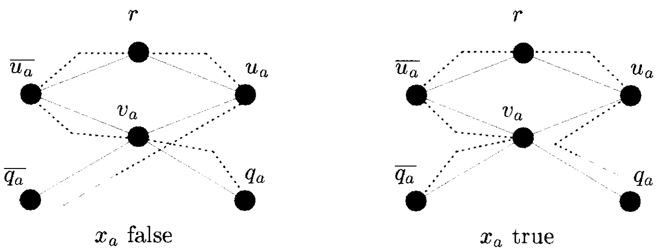

Fig. 2. Truth setting component and path layout for $\langle 3,1\rangle$-edge layout problem

$E_3 = \{\{v_a, \overline{q}_a\}, \{v_a, q_a\} \mid a = 1, \ldots, m\}$,
$E_4 = \{\{\overline{q}_a, z_b\} \mid a = 1, \ldots, m\,,\, b = 1, \ldots, g\,,\, \overline{x}_a \in c_b\}$,
$E_5 = \{\{q_a, z_b\} \mid a = 1, \ldots, m\,,\, b = 1, \ldots, g\,,\, x_a \in c_b\}$.

We call the subgraph of G induced by nodes $r, \overline{u}_a, u_a, v_a, \overline{q}_a, q_a$ the a-th truth setting component of G, as the restriction of any $\langle 3,1\rangle$-edge layout Ψ on this subgraph can be associated in a very natural way to a truth assignment for x_a (see Fig. 2). In fact, edge $\{r, \overline{u}_a\}$ (resp. $\{r, u_a\}$) must belong to Ψ, otherwise $\mathcal{H}(\overline{u_a}) = \infty$ (resp. $\mathcal{H}(u_a) = \infty$), as we have to route a physical shortest path from r to $\overline{u}_a$ (resp. u_a). Since $\mathcal{E}_{\max}(\Psi) = 1$, no other VP can start from r through edges $\{r, \overline{u}_a\}$ and $\{r, u_a\}$. Moreover, one of the two edges $\{\overline{u}_a, v_a\}$ or $\{u_a, v_a\}$, say $\{\overline{u}_a, v_a\}$, must form a VP (otherwise $\mathcal{H}(v_a) = \infty$) and again, since $\mathcal{E}_{\max}(\Psi) = 1$, there is one VP starting from u_a (if not we simply add it respecting the bound on the edge congestion and without increasing the hop count of the other nodes) which steps through or terminates at either $\overline{q}_a$ or q_a. In the first case the truth assignment associated to x_a is false, otherwise it is true. Notice that, if x_a is false (resp. true), then $\langle v_a, q_a\rangle$ (resp. $\langle v_a, \overline{q}_a\rangle$) must be a VP, so that $\mathcal{H}(q_a) = 3 < \infty$ (resp. $\mathcal{H}(\overline{q}_a) = 3 < \infty$).

Assume first that there is a truth assignment satisfying f. We show that there exists a $\langle 3,1\rangle$-edge layout Ψ for G. The VPs of Ψ are constituted by all edges in $E_1 \cup E_4 \cup E_5$ plus for each a, $1 \le a \le m$, the VP $\langle \overline{u}_a, v_a\rangle$, and if the truth assignment satisfies $\overline{x}_a$ (resp. x_a), the VPs $\langle u_a, v_a, \overline{q}_a\rangle$ and $\langle v_a, q_a\rangle$ (resp. $\langle u_a, v_a, q_a\rangle$ and $\langle v_a, \overline{q}_a\rangle$). Then, all nodes $\overline{u}_a, u_a \in V_1$ have hop count 1, all vertices $v_a \in V_2$ hop count 2, all vertices $\overline{q}_a \in V_3$ such that x_a is false (resp. true) have hop count 2 (resp. 3), and all vertices $q_a \in V_3$ such that x_a is false (resp. true) hop count 3 (resp. 2). Finally, each node $z_b \in V_4$ has hop count 3, as there is at least one literal in c_b, say $\overline{x}_a$ (resp. x_a), which is satisfied, i.e. such that $\{\overline{q}_a, z_b\}$ (resp. $\{q_a, z_b\}$) is a VP and $\mathcal{H}(\overline{q}_a) = 2$ (resp. $\mathcal{H}(q_a) = 2$).

Assume now that there is a $\langle 3,1\rangle$-edge layout Ψ for G and let us show that f is satisfiable. Consider the truth assignment induced by Ψ on the variables $x_1, \ldots, x_m$. For every clause c_b, by hypothesis node z_b has hop count at most 3. If a literal x_a or $\overline{x}_a$ belonging to c_b is not satisfied, then it is not possible to reach z_b in 3 hops through VPs coming from the a-th truth setting component. Therefore, there must exist at least one literal of c_b which is satisfied and since this is true for every clause, f is satisfiable. □

Even if for the sake of brevity in this paper we do not give any further complexity result for the edge-congestion case, by using proof techniques similar to those in [6] it is possible to give an exact characterization of all the tractable and intractable cases. In fact, the problem is NP-complete for any h and c, except for the cases $h = 1$ (any c), and $h = 2, c = 1$, for which a solution can be obtained in polynomial time by means of suitable flow constructions. A detailed proof can be found in the technical report associated to this paper [7].

A result analogous to Theorem 4.1 holds also for the node congestion case.

Theorem 4.2 *Given a network $G = (V, E)$, a root $r \in V$ and a positive integer c, deciding the existence of a $\langle 2,c\rangle$-node layout for G with root r is an NP-complete problem.*

Proof We show that the problem is NP-complete by providing a polynomial time transformation from the *Dominating Set* problem (DS) (known to be NP-complete; see [8]). In DS we have a universe set $U = \{u_1, \ldots, u_m\}$ of m elements, a family $\{A_1, \ldots, A_f\}$ of f subsets of U and an integer $k \le f$; we want to decide if there exist k subsets $A_{j_1}, \ldots, A_{j_k}$ which cover U, i.e., such that $\bigcup_{i=1}^{k} A_{j_i} = U$.

Starting from an instance I_{DS} of DS, we construct a graph G that admits a $\langle 2,c\rangle$-node layout with $c = m$ if and only if I_{DS} admits a cover.

Let $G = (V, E)$, where $V = \{r\} \cup V_1 \cup V_2 \cup V_3 \cup V_4$ and $E = E_1 \cup E_2 \cup E_3 \cup E_4$ (see Fig. 3), with:

$V_1 = \{v_a \mid a = 1, \ldots, m - k\}$,
$V_2 = \{q_b \mid b = 1, \ldots, m - (f - k) - 1\}$,
$V_3 = \{w_d \mid d = 1, \ldots, f\}$,
$V_4 = \{z_e \mid e = 1, \ldots, m\}$,

and

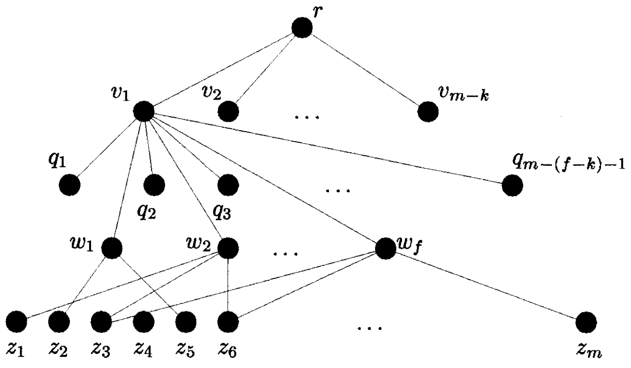

Fig. 3. The reduction graph for the $\langle h, c\rangle$-node layout problem

$E_1 = \{\{r, v_a\} \mid a = 1, \ldots, m-k\}$,
$E_2 = \{\{v_1, q_b\} \mid b = 1, \ldots, m-(f-k)-1\}$,
$E_3 = \{\{v_1, w_d\} \mid d = 1, \ldots, f\}$,
$E_4 = \{\{w_d, z_e\} \mid u_e \in A_d\}$.

Informally, in the reduction graph each subset A_d corresponds to the subgraph induced by node w_d and all nodes z_e such that $u_e \in A_d$, which are all connected to w_d. The idea underlying our construction is that, since at most k of the nodes w_d can be reached from r in one hop, if there are k dominating sets in I_{DS}, then all nodes of G can be reached from r in at most 2 hops.

Assume that there are k dominating sets $A_{j_1}, \ldots, A_{j_k}$. We show that there exists a $\langle 2, c\rangle$-node layout for G. The VPs of Ψ are constituted by all edges in $E_1 \cup E_2 \cup E_4$, the edges $\{v_1, w_d\} \in E_3$ such that A_d is not one of the dominating sets, i.e. $d \neq j_i$, $i = 1, \ldots, k$, and finally the VPs $\langle r, v_1, w_{j_i}\rangle$ for $i = 1, \ldots, k$ (which correspond to the k dominating sets). By construction, $\mathcal{N}(v) \leq m = c$ for each node $v \in V$. In order to check whether $\mathcal{H}_{\max}(\Psi) \leq 2$, it suffices to observe that all nodes $v_a \in V_1$ are reached in one hop, nodes $q_b \in V_2$ are reached in two hops, nodes $w_d \in V_3$ not corresponding to dominating sets are reached in two hops, nodes $w_{j_i} \in V_2$ corresponding to dominating sets are reached in one hop (through the VP $\langle r, v_1, w_{j_i}\rangle$), and as nodes $w_{j_1}, \ldots, w_{j_k}$ correspond to the k dominating sets, all nodes $z_e \in V_4$ are reached in two hops, since each of them is connected to at least one w_{j_i}.

It remains to show that if there are not k dominating sets, then no $\langle 2, c\rangle$-node layout Ψ for G exists. Consider any layout Ψ for G. Notice first that each of the edges $\{r, v_a\}$ must belong to Ψ, otherwise $\mathcal{H}(v_a) = \infty$. Similarly, since each node q_b must be reached through a shortest path, either the edge $\{v_1, q_b\}$ or the path $\langle r, v_1, q_b\rangle$ must be a VP of Ψ. Without loss of generality we can assume that the first case holds, as otherwise inserting $\{v_1, q_b\}$ in the set of the VPs of Ψ and replacing another VP starting from v_1 with a longer one directly from r, $\mathcal{H}(q_b) = 2$ and the hop count and node congestion of all the other nodes can only be decreased. Then, there are f nodes $w_d \in V_3$ to be reached along shortest paths and this can be done only through the remaining f VPs, of which k can start from the root and $f-k$ from v_1, yielding respectively hop count 1 and 2. Hence, no node in V_4 can be reached in two hops without exploiting a VP starting from a node $w_d \in V_3$. Let $w_{j_1}, \ldots, w_{j_k}$ be the k nodes in V_3 such that $\mathcal{H}(w_{j_i}) = 1$, $i = 1, \ldots, k$. Since there are not k dominating sets, then at least one node z_e is not connected to any of the nodes $w_{j_1}, \ldots, w_{j_k}$, and therefore $\mathcal{H}(z_e) \geq 3$. □

Like for the edge-congestion case, also here it is possible to give an exact characterization of all the tractable and intractable cases. In particular, observe that in the node congestion case, once fixed h and c, the problem of determining the existence of a $\langle h, c\rangle$-node layout for any graph G has a polynomial time-complexity, since from Fact 3.2 we know that the number of nodes in G has to be $N(h, c) \leq \frac{c(c-1)^h - 2}{c-2}$, i.e., it is always bounded by a constant. Hence, in all the intractable cases either h or c or both are not constant, i.e. they are part of the instance of the problem. Then, it is possible to show that the node layout problem is NP-complete for any fixed $h \geq 2$ (c not constant) and for any $c \geq 3$ (h not constant), while it can be solved in polynomial time in all the remaining cases. Again, for a detailed description see the technical report [7].

5 Results for specific topologies

In this section we give optimum layouts for specific topologies.

Let us consider first a chain or path of nodes with node set $V = \{1, \ldots N\}$ and edge set $E = \{\{i, i+1\} | 1 \leq i < N\}$. In order to give worst case estimations on the longest chain admitting a $\langle h, c\rangle$-edge or $\langle h, c\rangle$-node layout, we assume $r = 1$ as the root node.

Theorem 5.1 *Let $\mathcal{P}$ be the family of chain (or path) graphs. Then $E(h, c) = \frac{c^{h+1}-1}{c-1}$ and $N(h, c) = \frac{c(c-1)^h - 2}{c-2}$.*

Proof By Fact 3.1 $E(h, c) \leq 1 + c\frac{c^h - 1}{c-1} = \frac{c^{h+1}-1}{c-1}$ and by Fact 3.2 $N(h, c) \leq \frac{c(c-1)^h - 2}{c-2}$.

The lower bound on $E(h, c)$ (resp. $N(h, c)$) follows by observing that from the root of any chain it is possible to reach the next c nodes in one hop, and from each node with hop count at least one again the first next unreached c nodes (resp. $c-1$ nodes), thus yielding $E(h, c) \geq 1 + c + c^2 + \ldots + c^h = \frac{c^{h+1}-1}{c-1}$ and $N(h, c) \geq 1 + c + c(c-1) + \ldots + c(c-1)^{h-1} = \frac{c(c-1)^h - 2}{c-2}$. □

A ring graph consists of a node set $V = \{0, \ldots, N-1\}$ and an edge set $E = \{\{i, (i+1)_{mod\,N}\} | 0 \leq i < N\}$. As a ring is node-symmetric, without loss of generality it is possible to choose any node as the root. By arguments similar to those ones for chain graphs it is possible to prove the following theorem.

Theorem 5.2 *Let $\mathcal{R}$ be the family of ring graphs, then $E(h, c) = 2\frac{c^{h+1}-1}{c-1} - 1$ and $N(h, c) = \frac{c(c-1)^h - 2}{c-2}$ if c is even, otherwise $N(h, c) = 1 + \frac{(c-1)^{h+1} - 1}{c-2}$.*

Proof Again, by Fact 3.1, taking $d_r = 2$ and $d = 1$, $E(h, c) \leq 1 + 2c\frac{c^h - 1}{c-1} = 2\frac{c^{h+1}-1}{c-1} - 1$ and such an upper bound can always be attained by observing that, similarly as for chains, from the root it is possible to reach in h hops $c + c^2 + \ldots + c^h = \frac{c^{h+1}-1}{c-1} - 1$ nodes clockwise and $\frac{c^{h+1}-1}{c-1} - 1$ nodes anti-clockwise, thus forming a $\langle h, c\rangle$-edge layout for a ring of $2(\frac{c^{h+1}-1}{c-1} - 1) + 1 = 2\frac{c^{h+1}-1}{c-1} - 1$ nodes.

Concerning the node congestion case, we distinguish between the case in which c is even and the case in which c is odd.

In the former, by Fact 3.2, $N(h,c) \leq \frac{c(c-1)^h-2}{c-2}$, and the lower bound on $N(h,c)$ follows by observing that from the root it is possible to reach in one hop the closest c nodes, of which $c/2$ clockwise and the other $c/2$ anti-clockwise, and from each node reached clockwise (resp. anti-clockwise) in at least one hop again the first next unreached clockwise (resp. anti-clockwise) $c-1$ nodes, thus yielding $N(h,c) \geq 1+c+c(c-1)+\ldots+c(c-1)^{h-1} = \frac{c(c-1)^h-2}{c-2}$.

If c is odd, then either clockwise or anti-clockwise, say clockwise, it is possible to reach at most $\frac{c-1}{2}$ nodes. From each of these nodes again clockwise it is possible to reach at most other $c-1$ unreached nodes and so forth, till reaching within h hops a total of at most $\frac{c-1}{2}(1+(c-1)+\ldots+(c-1)^{h-1}) = \frac{c-1}{2}\frac{(c-1)^h-1}{c-2}$ nodes. Since all the nodes must be reached along shortest paths and the last two nodes reached respectively clockwise and anti-clockwise must be adjacent, starting from the root in the anti-clockwise direction it is possible to reach at most $\frac{c-1}{2}\frac{(c-1)^h-1}{c-2}+1$ nodes. This yields $N(h,c) \leq 1 + \frac{c-1}{2}\frac{(c-1)^h-1}{c-2} + \frac{c-1}{2}\frac{(c-1)^h-1}{c-2} + 1 = 1 + \frac{(c-1)^{h+1}-1}{c-2}$.

The lower bound on $N(h,c)$ follows by observing that the above construction can always be done in a ring with such a number of nodes. In fact, reaching from the root in one hop the closest $\frac{c-1}{2}$ nodes clockwise and the closest $\frac{c-1}{2}$ nodes anti-clockwise, similarly as for chains, $\frac{c-1}{2}\frac{(c-1)^h-1}{c-2}$ nodes can always be reached clockwise and $\frac{c-1}{2}\frac{(c-1)^h-1}{c-2}$ anti-clockwise, plus another one anti-clockwise through the remaining available VP from the root. □

We now turn our attention to the 2-dimensional extension of chains, that is to grids.

Given a square grid $G_{n\times n}$ of $N = n^2$ nodes, with node set $V = \{(i,j) | 1 \leq i \leq n, 1 \leq j \leq n\}$ and edge set $E = \{\{(i,j),(i+1,j)\} | 1 \leq i < n, 1 \leq j \leq n\} \cup \{\{(i,j),(i,j+1)\} | 1 \leq i \leq n, 1 \leq j < n\}$, again in order to give worst case estimations on the largest grid admitting a $\langle h,c\rangle$-edge or $\langle h,c\rangle$-node layout, we assume $r = (1,1)$ as the root node.

For the case of edge congestion $c = 1$, as stated by the following theorem the dimension of the largest grid admitting a $\langle h,c\rangle$-edge layout is dominated by the maximum number of nodes reachable in h hops along the first row or column.

Theorem 5.3 *Let $\mathcal{P}^2$ be the family of square grid graphs, then $E(h,1) = h^2$ if $h \leq 3$, otherwise $E(h,1) = (h+1)^2$.*

Proof Observe first that $E(h,1) \leq \lfloor\sqrt{N}\rfloor^2$, where $N = 2^{h+1}-1$. In fact, by Fact 3.1, taking $d_r = 2$ and $d = 2$, the maximum number of nodes reachable in h hops from r is $N \leq \frac{(2c)^{h+1}-1}{2c-1} = 2^{h+1}-1$, and since every grid has a quadratic number of nodes, that is n^2 for a given integer $n \geq 1$, $n = \lfloor\sqrt{N}\rfloor$ is the maximum integer such that $n^2 \leq N$. On the other hand, the maximum number of nodes along the first row or column reachable from r in h hops is $h+1$ (r included), so that $E(h,1) \leq (h+1)^2$. Since $\lfloor\sqrt{N}\rfloor^2 = \lfloor\sqrt{2^{h+1}-1}\rfloor^2 \leq h^2$ for $h \leq 3$, by combining the above constraints we obtain $E(h,1) \leq h^2$ for $h \leq 3$ and $E(h,1) \leq (h+1)^2$ for $h > 3$. Layouts matching these upper bounds are explicitly shown in Fig. 4 for $h \leq 4$, together with the hop distance of each node from the root. Note that for $h = 4$ we are able to build a layout for a 5×5 square grid such that at each node $(i,5)$ (resp. $(5,i)$), $i < 5$, belonging to the right (resp. bottom) border of the grid it results $\mathcal{E}((i,5),\{(i,5),(i+1,5)\}) = 0$ (resp. $\mathcal{E}((5,i),\{(5,i),(5,i+1)\}) = 0$). Therefore, we can build a $\langle 5,1\rangle$-edge layout for a 6×6 square grids by using these edges and edges $\{(i,5),(i,6)\}$ and $\{(5,i),(6,i)\}$, still maintaining the same property at the nodes $(i,6)$ and $(6,i)$ with $i < 6$ belonging to the right or bottom border of the 6×6 grid. In general, this gives an inductive construction to obtain from a $\langle h-1,1\rangle$-edge layout for a $h\times h$ square grid a $\langle h,1\rangle$-edge layout for a $(h+1)\times(h+1)$ square grid, for any $h \geq 5$. □

Theorem 5.4 *Let $\mathcal{P}^2$ be the family of square grid graphs. Then, for $c \geq 4$, $E(h,c) = \lfloor\sqrt{N_{h,c}}\rfloor^2$, where $N_{h,c} = \frac{(2c)^{h+1}-1}{2c-1}$.*

Proof Again by Fact 3.1 with $d_r = 2$ and $d = 2$, the maximum number of nodes reachable in h hops is $N_{h,c} \leq \frac{(2c)^{h+1}-1}{2c-1}$, and since every grid has a quadratic number of nodes, the upper bound on $E(h,c)$ derives directly by observing that $n = \lfloor\sqrt{N_{h,c}}\rfloor$ is the maximum integer such that $n^2 \leq N_{h,c}$.

In order to provide an optimal layout, given a square grid G with at least $N_{h,c}$ nodes, we define a gridoid G_h as the subgrid of G induced by nodes (i,j) with $i \leq \lfloor\sqrt{N_{h,c}}\rfloor$ and $j \leq \lfloor\sqrt{N_{h,c}}\rfloor$, i.e., the $\lfloor\sqrt{N_{h,c}}\rfloor \times \lfloor\sqrt{N_{h,c}}\rfloor$ subgrid induced by the first $\lfloor\sqrt{N_{h,c}}\rfloor$ rows and columns, plus the $N_{h,c} - \lfloor\sqrt{N_{h,c}}\rfloor^2$ nodes starting from node $(\lceil\sqrt{N_{h,c}}\rceil, 1)$, going toward node $(\lceil\sqrt{N_{h,c}}\rceil, \lceil\sqrt{N_{h,c}}\rceil)$ along row $\lceil\sqrt{N_{h,c}}\rceil$ and then, if $N_{h,c} - \lfloor\sqrt{N_{h,c}}\rfloor^2 > \lceil\sqrt{N_{h,c}}\rceil$, up along column $\lceil\sqrt{N_{h,c}}\rceil$ taking nodes $(\lceil\sqrt{N_{h,c}}\rceil-1, \lceil\sqrt{N_{h,c}}\rceil)$, $(\lceil\sqrt{N_{h,c}}\rceil-2, \lceil\sqrt{N_{h,c}}\rceil)$, and so forth.

Let the order of G_h be $n_h = \lfloor\sqrt{N_{h,c}}\rfloor$, that is the number of rows or columns of the largest subgrid contained in G_h. We now show an incremental construction for layouts with edge congestion at most c such that, for any positive integer h, the subgraph induced by all the nodes with hop count at most h is G_h (see Fig. 5). The theorem then follows by considering the restriction of the layout on the $n_h \times n_h$ subgrid of G_h containing $n_h^2 = E(h,c)$ nodes.

Clearly G_0 contains only the root $(1,1)$ and a $\langle 1,c\rangle$-edge layout for G_1 can be easily constructed by putting a suitable VP from the root to each node in G_1. Let us now show when $h \geq 1$ how to construct from a $\langle h,c\rangle$-edge layout for G_h a $\langle h+1,c\rangle$-edge layout for G_{h+1}. Notice that, for any node (i,j), all the nodes (i',j') with $i' \leq i$ and $j' \leq j$ belong to a shortest path from (i,j) to the root $(1,1)$. Then, we first have a set of *expanding* VPs that, for each row i (resp. column i) with $1 \leq i \leq n_h$, are between the nodes in row i (resp. column i) with hop count h (that is belonging to G_h but not to G_{h-1}) and the nodes in row i (resp. column i) belonging to $G_{h+1} - G_h$, so that each of them is reached in $h+1$ hops. All the remaining available VPs from the nodes in $G_h - G_{h-1}$ are used to reach the remaining not considered nodes (i,j) of $G_{h+1} - G_h$ with $i > n_h$ and $j > n_h$, i.e. in the right-down corner.

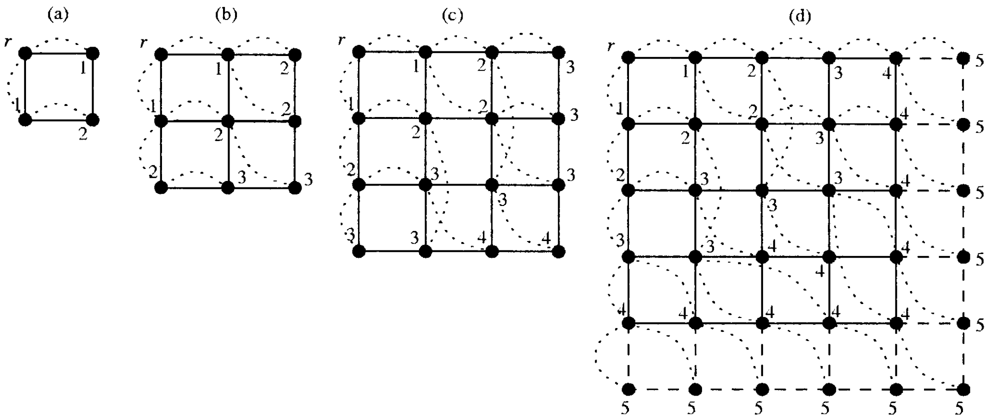

Fig. 4. (a) $\langle 2,1\rangle$-edge layout (dotted) for a 2×2 square grid graph; (b) $\langle 3,1\rangle$-edge layout for a 3×3 grid; (c) $\langle 4,1\rangle$-edge layout for a 4×4 grid; (d) $\langle 4,1\rangle$-edge layout for a 5×5 grid with the inductive step to get a $\langle 5,1\rangle$-edge layout for a 6×6 grid

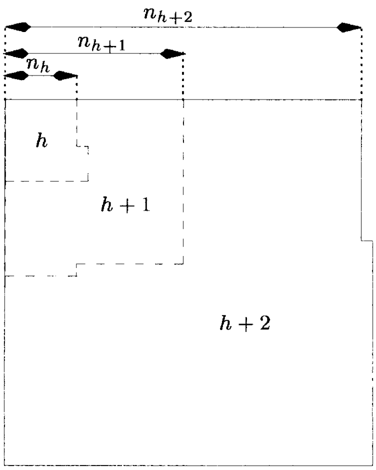

Fig. 5. The incremental layout for grids

For any given row i (resp. column i) with $1 \le i \le n_h$, let d_h be the number of nodes in row i (resp. column i) belonging to $G_{h+1} - G_h$. Since each edge can have congestion at most c, in order to guarantee the correctness of the above incremental construction we have to prove that $c \cdot d_h \ge d_{h+1}$.

By construction,

$$
\begin{aligned}
d_h &\le n_{h+1} - n_h + 1 = \lfloor\sqrt{N_{h+1,c}}\rfloor - \lfloor\sqrt{N_{h,c}}\rfloor + 1 \\
&\le \sqrt{\frac{(2c)^{h+2}-1}{2c-1}} - \sqrt{\frac{(2c)^{h+1}-1}{2c-1}} + 2 \\
&= \sqrt{\frac{(2c)^{h+1}}{2c-1}}\left(\sqrt{2c-\frac{1}{(2c)^{h+1}}} - \sqrt{1-\frac{1}{(2c)^{h+1}}}\right) + 2 \\
&\le \sqrt{\frac{(2c)^{h+1}}{2c-1}}\left(\sqrt{2c-\frac{1}{2c}} - \sqrt{1-\frac{1}{2c}}\right) + 2 \\
&= \sqrt{\frac{(2c)^{h+1}}{2c-1}}\left(\sqrt{\frac{(2c-1)(2c+1)}{2c}} - \sqrt{\frac{2c-1}{2c}}\right) + 2 \\
&= \sqrt{(2c)^{h+1}}\left(\frac{\sqrt{2c+1}-1}{\sqrt{2c}}\right) + 2.
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
d_h &\ge n_{h+1} - n_h - 1 = \lfloor\sqrt{N_{h+1,c}}\rfloor - \lfloor\sqrt{N_{h,c}}\rfloor - 1 \\
&\ge \sqrt{\frac{(2c)^{h+2}-1}{2c-1}} - \sqrt{\frac{(2c)^{h+1}-1}{2c-1}} - 2 \\
&> \sqrt{(2c)^{h+1}}\left(\frac{\sqrt{2c}-1}{\sqrt{2c-1}}\right) - 2.
\end{aligned}
$$

It is possible to verify that if $4 \le c \le 11$, $h \ge 2$ or if $c \ge 12$, any h it is

$$
\begin{aligned}
cd_h &> c\sqrt{(2c)^{h+1}}\left(\frac{\sqrt{2c}-1}{\sqrt{2c-1}}\right) - 2c \\
&\ge \sqrt{(2c)^{h+2}}\left(\frac{\sqrt{2c+1}-1}{\sqrt{2c}}\right) + 2 \ge d_{h+1}.
\end{aligned}
$$

A case analysis shows that the construction works also for $4 \le c \le 11$ and $0 \le h < 2$. This completes the proof of the theorem. □

Tighter results can be determined for the node congestion case. Notice first that no $\langle h,c\rangle$-node layout with $c \le 2$ can exist for a grid larger than 2×2. In fact, for $c = 1$ it is not possible to have a VP from the root $(1,1)$ to one of its two neighbors $(1,2)$ and $(2,1)$, that in turn cannot be reached through a shortest path. If $c = 2$, as the edges from $(1,1)$ respectively to $(1,2)$ and $(2,1)$ must form two VPs, one of the 3 nodes $(1,3)$, $(3,1)$ and $(2,2)$ cannot be reached through a shortest path.

For $c \ge 3$ the following theorem holds.

Theorem 5.5 *Let $\mathcal{P}^2$ be the family of square grid graphs. Then, for $c \ge 3$, $N(h,c) = \lfloor\sqrt{\frac{c(c-1)^h-2}{c-2}}\rfloor^2$.*

Proof By Fact 3.2 the maximum number of nodes reachable in h hops is $\frac{c(c-1)^h-2}{c-2}$ and the $\lfloor\sqrt{\frac{c(c-1)^h-2}{c-2}}\rfloor^2$ upper bound on $N(h,c)$ follows by observing that every grid has a quadratic number of nodes.

In order to provide a matching lower bound, again we apply the gridoid method of Theorem 5.4, and the proof proceeds exactly as in Theorem 5.4 by considering the new value $\frac{c(c-1)^h-2}{c-2}$ for $N_{h,c}$.

Clearly the gridoid G_0 contains only the root $(1,1)$ and a $\langle 1,c\rangle$-node layout for G_1 can be easily constructed by putting a suitable VP from the root to each node in G_1.

Again, let $n_h = \lfloor\sqrt{N_{h,c}}\rfloor$ be the order of G_h, that is the number of rows or columns of the largest subgrid contained in G_h, and for any given row i (resp. column i) with $1 \le i \le n_h$, let d_h be the number of nodes in row i (resp. column i) belonging to $G_{h+1} - G_h$. Since each node has congestion at most c, in order to guarantee the correctness of the incremental construction allowing to determine gridoids G_h for $h > 1$, we have to prove that $c \cdot d_h \ge d_{h+1}$.

By construction,

$$\begin{aligned} d_h &\le n_{h+1} - n_h + 1 = \lfloor\sqrt{N_{h+1,c}}\rfloor - \lfloor\sqrt{N_{h,c}}\rfloor + 1 \\ &\le \sqrt{\frac{c(c-1)^{h+1}-2}{c-2}} - \sqrt{\frac{c(c-1)^h-2}{c-2}} + 2 \\ &= \sqrt{\frac{c(c-1)^h}{c-2}}\left(\sqrt{(c-1) - \frac{2}{c(c-1)^h}} \right. \\ &\quad \left. - \sqrt{1 - \frac{2}{c(c-1)^h}}\right) + 2. \end{aligned}$$

Similarly,

$$\begin{aligned} d_h &\ge n_{h+1} - n_h - 1 = \lfloor\sqrt{N_{h+1,c}}\rfloor - \lfloor\sqrt{N_{h,c}}\rfloor - 1 \\ &\ge \sqrt{\frac{c(c-1)^{h+1}-2}{c-2}} - \sqrt{\frac{c(c-1)^h-2}{c-2}} - 2 \\ &> \sqrt{\frac{c(c-1)^{h+1}}{c-2}} - \sqrt{\frac{c(c-1)^h}{c-2}} - 2 \\ &= \sqrt{\frac{c(c-1)^h}{c-2}}\left(\sqrt{c-1} - 1\right) - 2. \end{aligned}$$

It is possible to verify that if $3 \le c \le 6$, $h \ge 6$ or if $c \ge 7$, any h it is

$$\begin{aligned} cd_h &> c\sqrt{\frac{c(c-1)^h}{c-2}}\left(\sqrt{c-1} - 1\right) - 2c \\ &\ge \sqrt{\frac{c(c-1)^{h+1}}{c-2}}\left(\sqrt{(c-1) - \frac{2}{c(c-1)^{h+1}}} \right. \\ &\quad \left. - \sqrt{(1 - \frac{2}{c(c-1)^{h+1}}}\right) + 2 \\ &\ge d_{h+1}. \end{aligned}$$

A case analysis shows that the construction works also for $3 \le c \le 6$ and $0 \le h < 6$, hence the theorem. □

6 Conclusion and open problems

In this paper we have introduced and analyzed two new cost measures related to the communication overhead and the space requirements associated to virtual path layouts in ATM networks, that is the *edge congestion* and the *node congestion*.

All the provided time complexity results are tight, and the same holds for the layout constructions for specific topologies, except for the cases of an edge congestion $c = 2$ and $c = 3$ in grids. We are very close to the determination of these layouts, but they are not incremental, that is the subset of the nodes with hop count at most equal to a given integer h in general does not form a gridoid. In fact, it is possible to see that in these cases the incremental solution does not work, as there are values of h such that from the gridoid of the nodes with hop count at most h it is not possible to build the successive one corresponding to a hop count at most equal to $h + 1$.

An interesting issue to be pursued is the determination of optimal path layouts for other network topologies. Moreover, it would be interesting to extend all the results to all-to-all layouts, where communication must be guaranteed between any two pairs of nodes.

Besides the relationships discussed in Sect. 3, another open question concerns the determination of further connections between the new congestion measures and the load parameter of [11,4].

Finally, while we have remarked that in this context it does not make sense to consider an unbounded stretch factor, a case worth to investigate is when the stretch factor is bounded by a given real number greater than one.

References

1. S. Ahn, R.P. Tsang, S.R. Tong, D.H.C. Du. Virtual path layout design on ATM networks. In: *INFOCOM'94*, pp. 192–200, 1994
2. L. Becchetti, P. Bertolazzi, C. Gaibisso, G. Gambosi. On the design of efficient ATM schemes. In *SOFSEM '97, Lecture Notes in Computer Science*, Vol. 1338, pp. 375–382. Berlin Heidelberg New York: Springer 1997
3. J. Burgin, D. Dorman: Broadband ISDN resource management: The role of virtual paths. *IEEE Communications Magazine* 29 (1991)
4. I. Cidon, O. Gerstel, S. Zaks: A scalable approach to routing in ATM networks. In: G. Tel, P.M.B. Vitányi (eds): *8th International Workshop on Distributed Algorithms (LNCS 857)*, pp. 209–222, Terschelling, The Netherlands, October 1994. Submitted for publication in IEEE/ACM Trans. on Networking
5. R. Cohen, A. Segall: Connection management and rerouting in ATM networks. In: *INFOCOM'94*, pp. 184–191, 1994
6. T. Eilam, M. Flammini, S. Zaks: A complete characterization of the path layout construction problem for ATM networks with given hop count and load. In: *24th International Colloquium on Automata, Languages and Programming (ICALP), Lecture Notes in Computer Science*, Vol. 1256, pp. 527–537. Berlin Heidelberg New York: Springer 1997
7. M. Flammini, E. Nardelli, G. Proietti: ATM layouts with bounded hop count and congestion. Technical Report 97-163, Dipartimento di Matematica Pura ed Applicata, Università di L'Aquila, October 1997
8. M.R. Garey, D.S. Johnson: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman 1979
9. O. Gerstel, I. Cidon, S. Zaks: The layout of virtual paths in ATM networks. *IEEE/ACM Transactions on Networking* 4(6), 873–884 (1996)
10. O. Gerstel, A. Wool, S. Zaks: Optimal layouts on a chain ATM network. In *3rd Annual European Symposium on Algorithms (ESA), (LNCS 979)*, Corfu, Greece, September 1995, pp. 508-522. To appear in *Discrete Applied Mathematics*
11. O. Gerstel, S. Zaks: The virtual path layout problem in fast networks. In: *13th ACM Symp. on Principles of Distributed Computing*, pp. 235–243, Los Angeles, USA, August 1994

12. R. Händler, M.N. Huber: *Integrated Broadband Networks: an introduction to ATM-based networks*. Reading, MA: Addison-Wesley 1991
13. ITU recommendation: I series (B-ISDN), Blue Book, November 1990
14. E. Kranakis, D. Krizanc, A. Pelc: Hop-congestion tradeoffs in ATM networks. In *9th IEEE Symposium on Parallel and Distributed Processing*, pp. 662–668, 1995
15. C. Partridge: *Gigabit Networking*. Reading, MA: Addison-Wesley 1994
16. K.I. Sato, S. Ohta, I. Tokizawa: Broad-band ATM network architecture based on virtual paths. *IEEE Transactions on Communications* 38(8), 1212–1222 (1990)
17. Y. Sato, K.I. Sato: Virtual path and link capacity design for ATM networks. *IEEE Journal of Selected Areas in Communications* 9(1) (1991)
18. L. Stacho, I. V'rto: Virtual path layouts for some bounded degree networks. In *3rd Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pp. 269–278. Carleton University Press, 1996
19. S. Zaks: Path layouts in ATM networks. In: *SOFSEM '97, Lecture Notes in Computer Science*, Vol. 1338, pp. 144–160. Berlin Heidelberg New York: Springer 1997

Michele Flammini is associate professor of Computer Science at the University of L'Aquila (Italy). He received a degree in Computer Science at the University of L'Aquila in 1990 and a Ph.D. degree in Computer Science at the University of Rome "La Sapienza" in 1995. His research interests include algorithms, computational complexity, communication problems in interconnection networks and routing. He recently conducted research also on multicriteria optimization problems with special emphasis in approximability and competitive on-line solutions. He has authored and co-authored several papers in his fields of interest published in the most reputed international conferences and journals.

Enrico Nardelli is full professor of Computer Science at the University of L'Aquila and invited research scientist at the Institute for System Analysis and Computer Science of the Italian National Research Council in Rome. He carries out research in various fields of theoretical and applied Computer Science, namely the design and analysis of algorithms and data structures, the definition and analysis of data models, the definition and design of tools and environments for interaction with information systems. He has authored more than 50 refereed papers in the fields of interest, published in the most reputed international scientific journals. He has been Team Leader and Principal Investigator since 1989 in more than 10 national and international Research and Development Projects in Computer Science.

Guido Proietti is associate professor of Computer Science at the University of L'Aquila and invited research scientist at the Institute for System Analysis and Computer Science of the Italian National Research Council in Rome. In 1998 he was a visiting scientist at the Department of Computer Science of Carnegie Mellon University, Pittsburgh. His current activity includes the study of efficient spatial access methods, query processing and evaluation, and the design of algorithms for fault-tolerance in networks. He is also interested in graph algorithms, computational geometry and computer vision. He has published more than 30 refereed articles.