

Augmented Ring Networks

William Aiello, *Member, IEEE*, Sandeep N. Bhatt, Fan R.K. Chung,
Arnold L. Rosenberg, *Fellow, IEEE*, and Ramesh K. Sitaraman, *Member, IEEE*

Abstract—We study four augmentations of ring networks which are intended to enhance a ring’s efficiency as a communication medium significantly, while increasing its structural complexity only modestly. *Chordal rings* add “shortcut” edges, which can be viewed as chords, to the ring. *Express rings* are chordal rings whose chords are routed outside the ring. *Multirings* append subsidiary rings to edges of a ring and, recursively, to edges of appended subrings. *Hierarchical ring networks (HRNs)* append subsidiary rings to nodes of a ring and, recursively, to nodes of appended subrings. We show that these four modes of augmentation are very closely related: 1) Planar chordal rings, planar express rings, and multirings are topologically equivalent families of networks with the “cutwidth” of an express ring translating into the “tree depth” of its isomorphic multiring and vice versa. 2) Every depth- d HRN is a spanning subgraph of a depth- $(2d - 1)$ multiring. 3) Every depth- d multiring \mathcal{M} can be embedded into a d -dimensional mesh with dilation 3 in such a way that some node of \mathcal{M} resides at a corner of the mesh. 4) Every depth- d HRN \mathcal{H} can be embedded into a d -dimensional mesh with dilation 2 in such a way that some node of \mathcal{H} resides at a corner of the mesh. In addition to demonstrating that these four augmented ring networks are grid graphs, our embedding results afford us close bounds on how much decrease in diameter is achievable for a given increase in structural complexity for the networks. Specifically, we derive upper and lower bounds on the optimal diameters of N -node depth- d multirings and HRNs that are asymptotically tight for large N and d .

Index Terms—Ring networks, chordal rings, express rings, multirings, hierarchical ring networks, grid graphs, graph embedding, diameter trade-offs.



1 INTRODUCTION

EVER since the earliest uses of networks for communication and computation, there has been serious interest in ring networks because of their structural simplicity and (modest) fault tolerance. Of course, this interest has been moderated by the large diameter and small (bisection) bandwidth of rings. It was natural, therefore, to seek ways to augment a ring network in a manner that enhances the network’s efficiency as a communication medium—by increasing its bandwidth and decreasing its diameter significantly—while increasing its structural complexity only modestly. This paper is devoted to studying four avenues for such augmentation, each of which adds “shortcut” edges to a ring in a specific way. All four of the resulting network families are still the object of active study to this day—both by theoreticians and by system-designers.¹ We begin by exposing certain basic structural properties of the four families, proving that their constituent networks are “essentially” equivalent in communication power and are “almost” grid graphs, i.e., subgraphs of high-dimensional meshes. We then apply

the exposed structure to the question of how much decrease in diameter one can achieve via each type of ring augmentation, for a given increase in the structural complexity of the augmentation, deriving asymptotically tight upper and lower bounds on the diameter decrease as a function of augmentation complexity.

The remainder of this section is devoted to filling in the gaps of the preceding paragraph in a way that prepares the reader for our study. In Section 1.1, we formally describe the ring augmentations of interest. In Section 1.2, we preview our main results in a way that precisely defines the preceding paragraph’s terms “essentially” and “almost.” In Section 1.3, we suggest why a predominantly theoretical study, such as ours, is relevant to those who design and build systems. Finally, in Section 1.4, we discuss the literature on which our work builds.

1.1 Four Ring Augmentations

Our four modes of augmentation all start with the N -node ring network \mathcal{R}_N which has node-set $\mathbf{Z}_N \stackrel{\text{def}}{=} \{0, 1, \dots, N - 1\}$ and edge-set $\{(i, i + 1 \bmod N) \mid i \in \mathbf{Z}_N\}$. Of the names we use for our four families of augmented rings, “chordal ring” and “hierarchical ring network” appear in the literature while “express ring” and “multiring” are, to our knowledge, new here.

1.1.1 Chordal Rings

A *chordal ring* augments a ring network by adding “shortcut” edges that can be viewed as chords of the ring. Chordal rings were among the earliest proposed cost-efficient interconnection networks for parallel architectures [3] and they continue to be studied in this context, especially in regard to the problem of diameter

1. Citations which bear out this claim accompany the networks’ formal definitions in Section 1.1.

- W. Aiello is with AT&T Labs—Research, Florham Park, NJ 07960. E-mail: aiello@research.att.com.
- S.N. Bhatt is with Akamai Technologies, 201 Broadway, Cambridge, MA 02139. E-mail: bhatt@akamai.com.
- F.R.K. Chung is with the Department of Computer Science, University of California at San Diego, La Jolla, CA 92093. E-mail: fan@ucsd.edu.
- A.L. Rosenberg and R.K. Sitaraman are with the Department of Computer Science at the University of Massachusetts, Amherst, MA 01003. E-mail: {rsnrbg, ramesh}@cs.umass.edu.

Manuscript received 23 Feb. 1999; accepted 10 Dec. 1999.

For information on obtaining reprints of this article, please send e-mail to: tpds@computer.org, and reference IEEECS Log Number 109288.




Fig. 1. A chordal augmentation of the 21-node ring \mathcal{R}_{21} .

reduction (see, e.g., [4], [8], [25]). Fig. 1 depicts a chordal augmentation of \mathcal{R}_{21} .

1.1.2 Express Rings

An *express ring* is not a new network topology but, rather, a *drawing* of a chordal ring which routes the “shortcuts” around the exterior of the ring (in either the clockwise or counterclockwise sense), thereby turning each chord into an arc of the ring. Express rings constitute the simplest instance of both the express interconnection networks of [11] and the “hop”-enhanced optical communication networks of [12], [15], [16], [17], [18], [23], [31]. In deference to the latter sources, we henceforth use the word *hop* to refer ambiguously to an arc of an express ring or an edge of the underlying ring. The express ring in Fig. 2 is a drawing of the chordal ring in Fig. 1.

An important structural characteristic of an express ring \mathcal{R} is its (*circular*) *cutwidth*, i.e., the maximum number of hops that “cross over” any ring edge, counting the edge itself. Formally, if \mathcal{R} has N nodes, then its cutwidth is




Fig. 2. An express ring that is isomorphic to the chordal ring of Fig. 1.




Fig. 3. An MR that is isomorphic to the express ring of Fig. 2.

$$\max_{i \in \mathbb{Z}_N} |\{(j, k) \in Hops(\mathcal{R}) \mid j \leq i \text{ and } k \geq i + 1 \text{ mod } N\}|. \tag{1.1}$$

In order to understand this definition, one must keep in mind that circular cutwidth is a characteristic of a *drawing* since an express ring is itself a drawing. Hence, as is reflected in (1.1), one measures the cutwidth of an express ring by scanning hops in their clockwise sense (whence the “mod N ” in (1.1)). Thus, for instance, the express ring in Fig. 2 has cutwidth 5 because, e.g., in the drawing, the hop that connects nodes 0 and 14 crosses over the ring-edges (14, 15), (15, 16), (16, 17), (17, 18), (18, 19), (19, 20), (20, 0) and *not* over any other ring-edges.

Note. Cutwidth models important network resources such as (depending on context): the bandwidth requirements of the underlying ring’s various links; the number of distinct frequencies that the underlying ring’s links must support; the sizes of the switches that control the wires that cross “above” the express ring’s nodes.

Of special interest in our study are express rings that admit cutwidth- c drawings in which hops never cross. For brevity, we call such networks *cutwidth- c noncrossing express rings*.

1.1.3 Multirings

A *multiring* (MR, for short) is obtained from a ring network by appending at most one subsidiary ring to each edge of the ring and, recursively, to each edge of each subsidiary ring. (Each edge of an MR thus belongs to either one or two subrings; in the latter case, one subring is the “child” of the other in the tree of rings created by the recursive appendages.) MRs thus are one possible realization of a “tree of rings.” MRs capture the essential structure of the SONET (Synchronous Optical NETWORK) MRs [10] that are important in the realm of communication (especially telephonic) networks. Fig. 3 depicts an MR that is isomorphic to the express ring of Fig. 2.

Since our complexity-preserving transformations of MRs to and from express rings (in Section 2.1) have some subtlety, we now define MRs in some formal detail. The node-set of a *depth-d* MR \mathcal{M} consists of d pairwise disjoint sets, V_1, V_2, \dots, V_d ; each V_i comprises the *level-i* nodes of \mathcal{M} . The edges of \mathcal{M} are specified implicitly as follows:

- The induced subgraph on the level-1 node-set V_1 is the (unique) *level-1 ring* of \mathcal{M} . In Fig. 3, $V_1 = \{0, 7, 14\}$.
- The level-2 node-set V_2 is a disjoint union of some number $k_2 \leq |V_1|$ disjoint sets, $V_{2,1}, V_{2,2}, \dots, V_{2,k_2}$.
 - Each set $V_{2,j}$ is associated with a distinct pair of nodes $u_{2,j}, v_{2,j}$ that are adjacent in the level-1 ring of \mathcal{M} .
 - The induced subgraph of \mathcal{M} on each node-set $V_{2,j} \cup \{u_{2,j}, v_{2,j}\}$ is a *level-2 ring* of \mathcal{M} .

In Fig. 3, $V_{2,1} = \{4, 6\}$; $V_{2,2} = \{11, 13\}$; $V_{2,3} = \{18, 20\}$.

- For each level $i > 2$, the level- i node-set V_i is a disjoint union of some number $k_i \leq |V_{i-1}| + k_{i-1}$ disjoint sets, $V_{i,1}, V_{i,2}, \dots, V_{i,k_i}$.
 - Each set $V_{i,j}$ is associated with a distinct pair of adjacent nodes $u_{i,j}, v_{i,j}$ from a level- $(i-1)$ ring of \mathcal{M} . At least one of $u_{i,j}, v_{i,j}$ must be an element of V_{i-1} ; the other could belong to a lower-index V_k .
 - The induced subgraph of \mathcal{M} on each node-set $V_{i,j} \cup \{u_{i,j}, v_{i,j}\}$ is a *level- i ring* of \mathcal{M} .

In Fig. 3,

$$\begin{aligned} V_{3,1} &= \{1\}; V_{3,2} = \{5\}; V_{3,3} = \{8\}; V_{3,4} = \{12\}; \\ V_{3,5} &= \{15\}; V_{3,6} = \{19\}. \\ V_{4,1} &= \{3\}; V_{4,2} = \{10\}; V_{4,3} = \{17\}. \\ V_{5,1} &= \{2\}; V_{5,2} = \{9\}; V_{5,3} = \{16\}. \end{aligned}$$

An important structural characteristic of an MR is its (*tree-*)*depth*, i.e., the number of levels of the recursive appending of subsidiary rings. Easily, the depth of an MR is its number of levels, i.e., the number of sets V_i .

Note. The depth of an MR most obviously represents the number of decision points needed to route messages within the network. Less obviously, it is also a measure of the network's tolerance to edge-faults. It thus represents a trade-off of routing ease for fault tolerance and routing distance.

1.1.4 Hierarchical Ring Networks

A *hierarchical ring network* (HRN, for short) is obtained from a ring network by appending at most one subsidiary ring to each *node* of the ring and, recursively, to each node of each subsidiary ring. (Each node of an HRN thus belongs to either one or two rings; in the latter case, one ring is the "child" of the other.) HRNs are thus an alternative to MRs as a possible realization of a "tree of rings." As with MRs, the (*tree-*)*depth* of an HRN, i.e., the number of levels of the recursive appending of subsidiary rings, is an important structural characteristic. HRNs have been studied extensively in the context of shared-memory multiprocessors and have actually been implemented in such machines [9], [26],




Fig. 4. A hierarchical ring network that is a spanning subgraph of the MR of Fig. 3.

[28], [30]. Fig. 4 depicts a depth-3 HRN which is (isomorphic to) a spanning subgraph of the MR of Fig. 3.

The reader can easily adapt the formal description of MRs in Section 1.1.3 to a formal description of HRNs. The key difference is that, with an HRN \mathcal{H} :

- Each set $V_{i,j}$ is associated with a distinct node $u_{i,j} \in V_{i-1}$ and
- the induced subgraph of \mathcal{H} on each node-set $V_{i,j} \cup \{u_{i,j}\}$ is a *level- i ring* of \mathcal{H} .

In Fig. 4,

$$\begin{aligned} V_1 &= \{0, 7, 14\}. \\ V_{2,1} &= \{1, 4\}; V_{2,2} = \{8, 11\}; V_{2,3} = \{15, 18\}. \\ V_{3,1} &= \{2, 3\}; V_{3,2} = \{5, 6\}; V_{3,3} = \{9, 10\}; \\ V_{3,4} &= \{12, 13\}; V_{3,5} = \{16, 17\}; V_{3,6} = \{19, 20\}. \end{aligned}$$

1.2 Our Main Results

1.2.1 Qualitative Results

Our study of augmented ring networks begins in Section 2 with two sets of structural results.

Topological equivalences. In Section 2.1, we establish strong relationships among our four families of augmented ring networks, which show them to be essentially equivalent in communication power.

1. Let \mathcal{R} be an augmented ring within one of the following three families: the family **C** of noncrossing chordal rings; the family **E** of noncrossing express rings²; the family **M** of MRs. In each of the other two families, there exist augmented rings that are (graph-theoretically) isomorphic to \mathcal{R} . More importantly, key structural parameters can be preserved under this isomorphism: For each cutwidth- c noncrossing express ring there is an isomorphic depth- c MR and vice versa.
2. Every depth- d HRN is (isomorphic to) a spanning³ subgraph of a depth- $(2d-1)$ MR.

² Easily, the families **C** and **E** comprise all and only *outerplanar* graphs [20].

³ A *spanning* subgraph of a connected graph \mathcal{G} is a connected subgraph on the nodes of \mathcal{G} .

Embeddability results. In Section 2.2, we prove that every depth- d MR \mathcal{M} —hence, any cutwidth- d noncrossing express ring—can be embedded into the positive orthant of the d -dimensional lattice⁴ (the d -lattice, for short) with dilation 3, in such a way that some node of \mathcal{M} resides at the origin of the orthant. Using an even simpler embedding, we prove that every depth- d HRN \mathcal{H} can be embedded into the positive orthant of the d -lattice with dilation 2, in such a way that some node of \mathcal{H} resides at the origin of the orthant.

1.2.2 Quantitative Results

In Section 3, we study diameter-structure trade-offs for augmented rings. In Section 3.1, we build on the results of Section 2 to derive asymptotically tight bounds on the diameters of HRNs and MRs, as functions of their depths.

- Let $\rho_d(N)$ denote the radius of the smallest d -dimensional ℓ_1 -sphere that contains $\geq N$ lattice points. (Equivalently, $\rho_d(N)$ is the smallest integer k such that $\geq N$ points of the d -lattice reside within distance k of the origin.)
- Let $\Delta_d^{(\text{HRN})}(N)$ (resp., $\Delta_d^{(\text{MR})}(N)$) denote the optimal diameter of an N -node depth- d HRN (resp., an N -node depth- d MR).

Note that, when $d = 1$, an N -node MR or HRN consists of a single ring so that

$$\Delta_1^{(\text{MR})}(N) = \Delta_1^{(\text{HRN})}(N) = \lfloor N/2 \rfloor.$$

For each fixed dimensionality $d > 1$, we derive the following upper and lower bounds that are asymptotically tight for large d and N :

$$2^{1-2/d} \rho_d(N)(1 - o(1)) \leq \Delta_d^{(\text{HRN})}(N), \Delta_d^{(\text{MR})}(N) \leq 2\rho_d(N).$$

Further, we show that, for any fixed d ,

$$\rho_d(N) = \frac{1}{2} N^{1/d} (d!)^{1/d} (1 \pm o(1)),$$

whence, for any fixed d ,

$$2^{-2/d} N^{1/d} (d!)^{1/d} (1 - o(1)) \leq \Delta_d^{(\text{HRN})}(N), \Delta_d^{(\text{MR})}(N) \leq N^{1/d} (d!)^{1/d} (1 + o(1)).$$

Our equivalence results readily extend these bounds to the diameters of cutwidth- d noncrossing express rings and chordal rings.

In Section 3.2, we settle the natural question of how our diameter-structure trade-off—worded in terms of express rings—changes when we allow hops to cross. We prove that the bounds of the trade-off weaken, at most, by a factor of two, i.e., allowing hops to cross can at best halve the diameters of cutwidth- c express rings. Simple examples show that the factor 2 is the best possible.

4. As usual, the d -dimensional lattice is the infinite analogue of the d -dimensional mesh, having nodes (lattice points) that are d -tuples of (positive and negative) integers and adjacencies between nodes that differ by precisely ± 1 in precisely one coordinate. The *positive orthant* of the lattice is the induced subgraph on the set of lattice points whose entries are all nonnegative.

1.3 The “Practical” Import of Our Results

Topological equivalences. The overriding significance of the topological equivalences we present in Section 2.1 is that they establish fundamental ties among the four essentially disjoint, active, contemporaneous bodies of literature that are devoted to our four ring augmentations. Our bibliography gives one but a peek into these literatures. These results are, thus, kindred to the several known structural equivalence results⁵ for butterfly-related networks (see, e.g., [5], [13], [24]). To our knowledge, our results are the first of this type for ring-related networks. More importantly than just obviating wasted work, the merging of these literatures can give one access to nonobvious efficient algorithms for important tasks. We exemplify this claim via two simple, closely related scenarios:

1. We are given an MR in its real-life incarnation as a SONET network. We wish to devise an efficient circuit-switching algorithm that establishes point-to-point connections in the network in a way that makes optimal use of available bandwidth. In contrast to “real” SONET networks, as described in [10], we allow both clockwise and counterclockwise routings of messages.
2. We wish to route messages within an optical HRN using as few frequencies as possible.

Although it is not obvious a priori, the algorithm developed in [14] for (in our terminology) converting chordal rings to express rings of minimum circular cutwidth is readily adapted (e.g., using our equivalence results) to efficiently solve both of these problems optimally.

Embeddability results. The most obvious lesson from our embedding results is that MRs and their kindred networks can be “packed” tightly into high-dimensional meshes. This tight packing leads directly (but not obviously) to the diameter-structure trade-offs that we establish for augmented rings in Section 3. More subtly, both the tightness of our embeddings and the dimensionalities of their target meshes expose and elucidate certain behavioral relationships between augmented rings and meshes along the lines of those reported in [19]. In a different direction, our embeddings lead to efficient emulations of augmented rings on other networks directly via mesh-embeddings, such as those in [21], and indirectly via complicated emulation algorithms, such as those in [22]. Finally, the embeddings expose a nonobvious grid-like structure in augmented rings which implies a variety of unexpected algorithmic and structural properties, as described in [7] and [32].

1.4 A Survey of Related Work

We have already cited numerous sources that relate to our study either in the object of study (augmented ring networks) or in the general focus (proving the “equivalence” of network families). However, we know of no prior technical results that relate directly to our new qualitative results. Further, we present the first upper and

5. Behavioral equivalence results are known for many other network families (see, e.g., [2], [22], [29]).

lower bounds for the diameter-structure trade-off for MRs and HRNs that are asymptotically tight. The following sources relate to our diameter-structure trade-offs for express rings (hence, for MRs):

Several papers study diameter-cutwidth trade-offs for hop-augmented networks. Notable among these are [23], whose journal version inspired our study of such trade-offs, and [18], [12], which were done contemporaneously to and independently of our study. The former two sources study *unidirectional express paths*: Start with a path and add hops that decrease the maximum left-to-right hopping distance from the left end of the path to any other node. We find in [18] an exact determination of the minimum of such a distance for any cutwidth- c N -node unidirectional express path. Although the bounds of [23] are less definitive than those of [18], their regular upper-bound construction is of interest when express links (hops) are weighted with a delay function which increases with the length of the hop.

In an invited talk in March 1997 [1], we presented an early version of the current results, providing the first nontrivial diameter-structure trade-offs for augmented ring networks. We showed there that:

$$\frac{1}{4e} dN^{1/d} - d/2 \leq \frac{1}{2} \rho_d(N) \leq \Delta_d^{(\text{MR})}(N) \leq 2^{-1/d} dN^{1/d}. \quad (1.2)$$

These bounds are within constant factors when d is “small,” say, $d \leq \frac{1}{6} \log_2 N$. Importantly, this work introduced a new lower-bounding technique—an early version of Theorem 2.4—based on embedding depth- d MRs into the d -lattice. Subsequently, the bounds in (1.2) were improved in [12]: The upper bound was decreased by a constant factor and the lower bound was increased by adapting the grid-embedding technique of [1] so that $\Delta_d^{(\text{MR})}(N)$ was determined to within a factor of 2 for all N and all d :

$$\rho_d(N) \leq \Delta_d^{(\text{HRN})}(N), \Delta_d^{(\text{MR})}(N) \leq 2\rho_d(N). \quad (1.3)$$

Independently of [12], we improved the bounds in both (1.2) and (1.3), deriving the asymptotically tight upper and lower bounds for both MRs and HRNs that we report in Section 3.

One finds in [31] a “state of the art” survey on diameter structure trade-offs, including the few that are known for augmented networks other than paths and rings.

2 EXPOSING THE STRUCTURE OF AUGMENTED RINGS

This section is devoted to establishing the qualitative results described in Section 1.2.1. Informally, in Section 2.1, we show that the four ring augmentations of interest are “essentially” equivalent in communication power. Then, in Section 2.2, we show that MRs (hence, also noncrossing chordal rings and express rings) are “almost” grid graphs [32].

2.1 Topological Equivalences

In the following results, “linear time” is measured relative to the number of edges in the subject network.

We begin by noting the obvious topological equivalence of chordal rings and express rings. The only issue of interest here is the ease of orienting the chords of a chordal ring in a way that minimizes the cutwidth of the resulting express ring.

Theorem 2.1 [14]. *In linear time, one can transform any planar drawing of a chordal ring into a planar drawing of an isomorphic express ring of minimum cutwidth and vice versa.*

Proof Sketch. 1) To transform any express ring to an equivalent chordal ring, simply ignore the (clockwise, counterclockwise) senses of the arcs and route them inside the ring. 2) One finds in [14] an efficient algorithm for transforming a chordal ring to an equivalent express ring having minimal cutwidth. (A slightly simpler, but less efficient, algorithm appears in [27].) \square

We turn next to the more challenging equivalence between express rings and MRs.

Theorem 2.2. *In linear time, one can transform any planar cutwidth- c drawing of an express ring into a drawing of an isomorphic depth- c MR and vice versa.*

Proof. We prove the claimed equivalence via a pair of explicit transformations which facilitate keeping track of the cutwidth-depth correspondence.

Transformation ER \rightarrow MR. Let \mathcal{R} be any express ring that is presented via a cutwidth- c planar drawing. We transform \mathcal{R} into an isomorphic depth- c MR by “turning it inside out” while decomposing it into c shells: Shell 1 will be a subring of \mathcal{R} while each other shell will be a set of subsidiary subrings. The subrings within each shell $\#k$ will turn out to be the level- k subrings of \mathcal{R} 's isomorphic MR. The decomposition proceeds as follows:

Shell 1 of \mathcal{R} is the subring formed by the hops of \mathcal{R} that are “exposed” in the sense of not being contained within any other hop of \mathcal{R} . The level-1 node-set of \mathcal{R} comprises the nodes of the shell-1 ring.

In Fig. 2, Shell 1 comprises the three-node ring ($0 \leftrightarrow 7 \leftrightarrow 14$).

Shell 2 of \mathcal{R} is obtained by focusing, in turn, on each hop of Shell 1. Say that (u, v) is such a hop, but that nodes u and v are not adjacent in the ring underlying \mathcal{R} . Then, say there is a path in \mathcal{R} that connects nodes u and v and uses hops that become “exposed” when the hops of Shell 1 are removed. The intermediate nodes of that path become one of the level-2 node-sets $V_{2,j}$ of \mathcal{R} and that path, plus hop (u, v) , becomes one of the subrings of Shell 2.

In Fig. 2, shell 2 comprises the three subrings

$$(0 \leftrightarrow 4 \leftrightarrow 6 \leftrightarrow 7 \leftrightarrow 0), (7 \leftrightarrow 11 \leftrightarrow 13 \leftrightarrow 14 \leftrightarrow 7), \\ (14 \leftrightarrow 18 \leftrightarrow 20 \leftrightarrow 0 \leftrightarrow 14).$$

In general, Shell $i + 1$ of \mathcal{R} is obtained by focusing, in turn, on each hop of Shell i that contains at least one node from node-set $V_i \stackrel{\text{def}}{=} \bigcup_j V_{i,j}$. Say that (u, v) is such a hop, but that nodes u and v are not adjacent in the ring underlying \mathcal{R} . Then, say there is a path in \mathcal{R} that connects nodes u and v and uses hops that become “exposed” when the hops of all shells $k \leq i$ are removed.




Fig. 5. Folding \mathcal{H} 's subsidiary subrings recursively within its root subring.

The intermediate nodes of that path become one of the level- $(i + 1)$ node-sets $V_{i+1,j}$ of \mathcal{R} and that path, plus hop (u, v) , becomes one of the subrings of Shell $i + 1$.

In Fig. 2, Shell 3 comprises the six subrings

- $(0 \leftrightarrow 1 \leftrightarrow 4 \leftrightarrow 0)$, $(4 \leftrightarrow 5 \leftrightarrow 6 \leftrightarrow 4)$, $(7 \leftrightarrow 8 \leftrightarrow 11 \leftrightarrow 7)$,
- $(11 \leftrightarrow 12 \leftrightarrow 13 \leftrightarrow 11)$, $(14 \leftrightarrow 15 \leftrightarrow 18 \leftrightarrow 14)$,
- $(18 \leftrightarrow 19 \leftrightarrow 20 \leftrightarrow 18)$;

Shell 4 comprises the three subrings

- $(1 \leftrightarrow 3 \leftrightarrow 4 \leftrightarrow 1)$, $(8 \leftrightarrow 10 \leftrightarrow 11 \leftrightarrow 8)$,
- $(15 \leftrightarrow 17 \leftrightarrow 18 \leftrightarrow 15)$;

and Shell 5 comprises the three subrings

- $(1 \leftrightarrow 2 \leftrightarrow 3 \leftrightarrow 1)$, $(8 \leftrightarrow 9 \leftrightarrow 10 \leftrightarrow 8)$,
- $(15 \leftrightarrow 16 \leftrightarrow 17 \leftrightarrow 15)$.

One can directly read off from these subrings the sets V_1, \dots, V_5 enumerated in Section 1.1.3.

We have, thus, constructed an MR whose nodes and edges are, respectively, the nodes and hops of the express ring \mathcal{R} . In particular, our running example has constructed the MR of Fig. 3 from the express ring of Fig. 2.

Transformation MR \rightarrow ER. The easiest way to perform this transformation proceeds in two steps. First, trace the obvious Hamiltonian cycle in a given MR, making all unused edges into chords of the ring obtained from the trace. Then, apply the algorithm of [14] to transform the resulting chordal ring into an express ring. The problem with this approach is that one must proceed carefully if one wants a depth- d MR to produce a noncrossing cutwidth- d express ring. We choose, therefore, to use an iterative transformation that allows us to perspicuously keep track of the parameter d in both the MR and the express ring.

We begin with a depth- d MR \mathcal{M} and transform it into an isomorphic express ring by successively ‘‘collapsing’’ its subsidiary rings. We begin with the level-1 ring of \mathcal{M} . Make all of its edges level-1 tentative edges of the express ring $\mathcal{R}_{\mathcal{M}}$ that we are constructing. In general, we look at each tentative edge in turn. When we are looking at a level- i tentative edge (u, v) , if there is no level- $(i + 1)$ ring that contains edge (u, v) , then we make this a final edge of $\mathcal{R}_{\mathcal{M}}$; if there is a level- $(i + 1)$ ring,

$$(u \leftrightarrow w_1 \leftrightarrow w_2 \leftrightarrow \dots \leftrightarrow w_k \leftrightarrow v \leftrightarrow u),$$

that contains edge (u, v) , then we make this edge a final arc between nodes u and v , and we incorporate, under this new arc, the path $(w_1 \leftrightarrow w_2 \leftrightarrow \dots \leftrightarrow w_k)$, with each edge of the path being a level- $(i + 1)$ tentative edge of $\mathcal{R}_{\mathcal{M}}$. When no tentative edges remain—which must happen since each tentative edge is scanned just once before being converted to a final entity (edge or arc)—we shall have produced an express ring $\mathcal{R}_{\mathcal{M}}$ whose nodes and hops are, respectively, the nodes and edges of \mathcal{M} . This recipe constructs the express ring in Fig. 2 from the MR in Fig. 3. \square

Our final structural result exposes the relationship between HRNs and MRs.

Theorem 2.3. *In linear time, one can add edges to a depth- d HRN that transform it into a depth- $(2d - 1)$ MR. Hence, the former network is a spanning subgraph of the latter.*

Proof. Let us be given a depth- d HRN \mathcal{H} . We proceed in a manner reminiscent of transformation MR \rightarrow ER of Theorem 2.2 to transform \mathcal{H} into a noncrossing cutwidth- $(2d - 1)$ express ring $\mathcal{E}(\mathcal{H})$. We then invoke transformation ER \rightarrow MR of Theorem 2.2 to transform $\mathcal{E}(\mathcal{H})$ into the desired depth- $(2d - 1)$ MR. Fig. 5 may make the following procedure more intuitive.

We begin with the level-1 cycle

$$(v_{1,0} \leftrightarrow v_{1,1} \leftrightarrow \dots \leftrightarrow v_{1,k_1-1})$$

of \mathcal{H} . For each $j \in \{0, 1, \dots, k_1 - 1\}$, we ‘‘fold’’ the sub-HRN $\mathcal{H}_{1,j}$ rooted at node $v_{1,j}$ under the cycle-edge between $v_{1,j}$ and $v_{1,j+1 \bmod k_1}$ as follows: We place the nodes $v_{1,j}, v_{1,j,1}, v_{1,j,2}, \dots, v_{1,j,k_1,j-1}$ of the level-1 cycle of $\mathcal{H}_{1,j}$ (which, of course, is a level-2 cycle of \mathcal{H}) in that order between nodes $v_{1,j}$ and $v_{1,j+1 \bmod k_1}$ and add edges that connect:

- node $v_{1,j}$ to node $v_{1,j,1}$,
- each node $v_{1,j,k}$ to its clockwise neighbor $v_{1,j,k+1}$, and
- node $v_{1,j,k_1,j-1}$ to node $v_{1,j+1 \bmod k_1}$.

Of course, all of these edges, save the very last, are edges of \mathcal{H} . We then recursively fold the sub-sub-HRNs rooted at these new nodes under the appropriate edges of this cycle.




Fig. 6. the embedding of \mathcal{M} 's level-1 subring.

The indicated folding process clearly produces a noncrossing express ring $\mathcal{E}(\mathcal{H})$ which contains \mathcal{H} as a spanning subgraph. The initial ring (level 1) of \mathcal{H} contributes 1 to the cutwidth of $\mathcal{E}(\mathcal{H})$. Each successive level of folding lays down a subring of \mathcal{H} in a way that contributes 2 to the cutwidth of $\mathcal{E}(\mathcal{H})$. In the end, therefore, $\mathcal{E}(\mathcal{H})$ has cutwidth $2d - 1$.

A direct application of transformation $\text{ER} \rightarrow \text{MR}$ of Theorem 2.2 now produces the desired depth- $(2d - 1)$ MR. \square

The transformation of Theorem 2.3 produces the MR of Fig. 3 from the HRN of Fig. 4 via the express ring of Fig. 2.

2.2 Augmented Rings as Grid-Graphs

We show now that the networks of interest are "almost" grid graphs in the sense of being embeddable into high-dimensional grids with small dilation. We use the standard notion of graph embedding [6]: One embeds a guest graph \mathcal{G} into a host graph \mathcal{H} via a one-to-one mapping of nodes of \mathcal{G} to nodes of \mathcal{H} together with a compatible routing of edges of \mathcal{G} along paths in \mathcal{H} . The *dilation* of the embedding is the length of the longest path of \mathcal{H} used to route an edge of \mathcal{G} . We strengthen the applications of our embeddings by insisting that some node of the guest augmented ring reside at a corner of the host grid.

Theorem 2.4. *Every depth- d MR \mathcal{M} can be embedded into the positive orthant of the d -lattice with dilation 3 in such a way that some node of \mathcal{M} is mapped to the origin of the lattice.*

Proof. We embed the given \mathcal{M} into the positive orthant of the d -lattice via the following strategy which avoids interference among \mathcal{M} 's subrings:

1. We number the dimensions of the d -lattice, starting from 1, from left to right and allocate dimension $i \in \{1, 2, \dots, d\}$ of the lattice to the level- i subrings of \mathcal{M} .
2. We embed \mathcal{M} 's (unique) level-1 subring in the familiar interleaved ring-into-path pattern along dimension 1 so that, up to parity, one "half" of the ring-nodes occupy even-numbered lattice nodes while the other "half" occupy the interleaved odd-numbered lattice nodes.
3. We embed \mathcal{M} 's level- i subrings in an interleaved pattern so that, for $i > 1$, (at most) the two level- i rings incident to each level- $(i - 1)$ node of \mathcal{M} use disjoint, alternating sets of lattice nodes along dimension i .

We now supply details of the node-mapping component of our embedding, allowing the simple edge-routing to be specified implicitly.

Embedding \mathcal{M} 's level-1 subring. The familiar dilation-2 embedding of a ring in a path proceeds as

follows: Letting the ring have nodes (in clockwise order) v_0, v_1, \dots, v_{n-1} , we take an n -step "walk" along dimension 1 of the lattice, starting at the origin $\langle 0, \vec{0} \rangle$. During the i th step of the walk, where $0 \leq i < n$, we visit node $\langle i, \vec{0} \rangle$ of the lattice. When i is even, we deposit node $v_{i/2}$ of the ring at this node. When i is odd, we deposit node $v_{n-\lceil i/2 \rceil}$ of the ring at this node (cf. Fig. 6).

Embedding \mathcal{M} 's level- k subrings, $k \geq 2$. Let us focus on an arbitrary level $k \geq 2$. Each level- k subring \mathcal{R} of \mathcal{M} contains two adjacent nodes of one of \mathcal{M} 's level- $(k - 1)$ subrings, which we call the *anchors* of \mathcal{R} . By induction, we assume that the anchors of each level- k subring lie on a line "parallel" to one of the axes of the d -lattice; this is certainly true of the level-1 ring. If a level- k subring \mathcal{R} has n nodes, then we prepare to embed it as follows:

Path reservation. Letting $m \stackrel{\text{def}}{=} \lfloor (n - 2)/2 \rfloor$, we reserve a path of length $3m$ emanating from each anchor node in the positive direction of and "parallel" to the dimension- k axis of the d -lattice. We view this path as consisting of m three-node *boxes*, as depicted in Fig. 7.

Subring labeling. We return now to the structure of \mathcal{M} . We label all level- k subrings with the "colors" $\{0, 1, 2\}$ in such a way that no two like-colored level- k subrings share an anchor node. This is possible since each node of each level- $(k - 1)$ subring is incident to at most two level- k subrings.

Note. We can "almost" color the subrings with two colors. We need the third color only when the number of subrings is odd and each subring is adjacent to two others (so that the subrings form a complete "daisy"). However, we actually make good use of the third color, even when it is not needed for subring discrimination.

We next label each nonanchor node of each level- k subring as either "outgoing" or "incoming" in the following way: We traverse all of the level- k subrings that emanate from a given level- $(k - 1)$ subring \mathcal{R}' by traversing \mathcal{R}' in a clockwise sense and traversing each of its level- k subrings as it is encountered for the first time. During this traversal of an n -node level- k subring, we label the first $\lceil (n - 2)/2 \rceil$ nodes "outgoing" and the remaining $\lfloor (n - 2)/2 \rfloor$ nodes "incoming." Our embedding will place at most one "incoming" node in each of the boxes we have reserved and up to two "outgoing" nodes in the manner we describe now.

Subring embedding. We are finally ready to embed the level- k subrings of \mathcal{M} . Let us focus on a subring that was assigned color $\kappa \in \{0, 1, 2\}$. If this ring has an even length, then we embed its nodes in the boxes emanating from its anchor nodes using the node labeled κ in each box (see Fig. 7). Our coloring regimen assures that "adjacent" subrings will not conflict with one another. If the subring has odd length n , then we modify the preceding embedding by placing the "outgoing" nodes $\lceil (n - 2)/2 \rceil - 2$ and $\lceil (n - 2)/2 \rceil - 1$ (counting outward from the anchor node) in the same box. One of these nodes will be placed in the node labeled κ in the box. The other will be placed in the node that is not used by the "incoming" node of the other subring that shares this anchor node. While the dilation of the embedding is not affected by the relative positions of "outgoing" nodes $\lceil (n - 2)/2 \rceil - 2$ and $\lceil (n - 2)/2 \rceil - 1$ in their shared box,




Fig. 7. A length- m path emanating from an anchor node.




Fig. 8. A portion of the interleaving of subrings A and B in the embedding.

the traversal of the embedded MR may be facilitated if node $\lceil (n - 2)/2 \rceil - 2$ is placed closer to the anchor node.

Validation and analysis. It should be clear from our construction that the described “embedding” (we have yet to show that it is, indeed, a valid embedding) has dilation 3. What we must verify is that the described node-assignment is one-to-one, i.e., that we have never assigned two nodes of \mathcal{M} to the same lattice point. To the end of this verification, let us consider the potential sources of collisions.

Adjacent level- k subrings. If two level- k subrings are incident on the same anchor node, then they share the path of boxes emanating from that anchor. However, our node-assignment reserves the path for “outgoing” nodes of one of the subrings and for “incoming” nodes of the other. Since 1) the two subrings are labeled with distinct colors, 2) every lattice-node occupied by an “incoming” node has the same color as the “incoming” node’s subring, and 3) every lattice-node occupied by an “outgoing” node has a color distinct from the color assigned to “incoming” nodes, it is clear that these two level- k subrings never collide in the embedding.

Interleaved level-2 subrings. The only other potential collisions arise from level-2 subrings that are embedded in an interleaved fashion that arises from the interleaving of the nodes of the level-1 subring. The situation is depicted schematically in Fig. 8. Now, the only potential for collision in this picture arises when $m_A = m_B$, in which case, subrings A and B want to “turn” (from “outgoing” to “incoming” or vice versa) at the same place along the path and, hence, are going to cross one another. We claim that this crossing can cause no collisions. To wit, if subrings A and B share an anchor, then they are adjacent and, hence, cannot collide by our earlier reasoning about colors and labels. Alternatively, if the subrings do not share an anchor, then subring B will end up on a lattice-path that is distinct from subring A ’s two lattice-paths once it “turns.” Once again, the “turn” cannot cause collisions.

We conclude that the described embedding is valid and, hence, verifies the claim of the theorem. \square

The $ER \rightarrow MR$ transformation of Theorem 2.2 allows us to adapt the embedding algorithm of Theorem 2.4 to cutwidth- d noncrossing express rings, yielding the following:

Corollary 2.1. *Every cutwidth- c noncrossing express ring \mathcal{E} can be embedded into the positive orthant of the c -dimensional lattice with dilation 3 in such a way that some node of \mathcal{E} is mapped to the origin of the lattice.*

We close this section with the HRN analogue of Theorem 2.4.

Theorem 2.5. *Every depth- d HRN \mathcal{H} can be embedded into the positive orthant of the d -lattice with dilation 2 in such a way that some node of \mathcal{H} is mapped to the origin of the lattice.*

Proof Sketch. Our embedding strategy follows that of Theorem 2.4 with the following simplification: Since every subring of an HRN has only one anchor node, each such level- k subring can be laid out in the k th dimension of the d -lattice with dilation 2, as in Fig. 6, in the same way that an MRs level-1 subring is laid out in the proof of Theorem 2.4. \square

3 DIAMETER-STRUCTURE TRADE-OFFS FOR AUGMENTED RINGS

This section is devoted to establishing the quantitative results described in Section 1.2.2. In Section 3.1, we bound the decrease of $\Delta_d^{(HRN)}(N)$ and $\Delta_d^{(MR)}(N)$ with increasing depth d , thereby simultaneously bounding the diameter of a noncrossing express ring with increasing cutwidth. In Section 3.2, we show that our diameter-cutwidth trade-off for express rings weakens by at most a factor of two if we allow hops to cross and we present a simple example to show that this factor is best possible.

3.1 A Diameter-Depth Trade-Off for HRNs and MRs

In this section, we study diameter-depth trade-offs for HRNs and MRs. Clearly, for $d = 1$, an HRN or an MR is just a ring so that

$$\Delta_d^{(\text{HRN})}(N) = \Delta_d^{(\text{MR})}(N) = \lfloor N/2 \rfloor.$$

Since we cannot derive precise expressions for the optimal diameters of HRNs and MRs when $d > 1$, we settle for asymptotically tight upper and lower bounds for general $\Delta_d^{(\text{HRN})}(N)$ and $\Delta_d^{(\text{MR})}(N)$.

We establish the overall strategy of our bounding technique by deriving the simpler bounds one finds in [12] before we turn to our tighter bounds in Theorem 3.1.

Lemma 3.1 [12]. *For any N and d ,*

$$\rho_d(N) \leq \Delta_d^{(\text{HRN})}(N), \Delta_d^{(\text{MR})}(N) \leq 2\rho_d(N).$$

Proof. We derive the bounds for HRNs, the argument for MRs being similar.

The upper bound. We construct an N -node depth- d HRN \mathcal{H} by packing the d -lattice with N nodes as densely as possible, i.e., in a manner that minimizes the maximum distance from the origin. By definition, the ℓ_1 -distance of any of the packed nodes from the origin is $\leq \rho_d(N)$. Now, we add edges to the packed nodes which turn them into the HRN \mathcal{H} . We add an edge between any two nodes u and v such that u resides at a lattice-point of the form $\langle \vec{x}, c, \vec{0} \rangle$ while v resides at one of u 's "neighbors" of the form $\langle \vec{x}, c \pm 1, \vec{0} \rangle$. Adding these edges creates a tree that is embedded into the d -lattice with unit dilation. The diameter of the tree is $\leq 2\rho_d(N)$ since any node can reach the origin via a path of length $\leq \rho_d(N)$. Finally, we add nontree edges to close the tree paths into rings, thereby creating the HRN \mathcal{H} . Since the nontree edges can only shorten distances within \mathcal{H} , we conclude that \mathcal{H} 's diameter is $\leq 2\rho_d(N)$.

The lower bound. Let \mathcal{H} be any N -node depth- d HRN. Consider the breadth-first (hence, shortest-path) tree $\mathcal{T}[u]$ rooted at an arbitrary node u of \mathcal{H} . Note first that $\mathcal{T}[u]$ can be laid out in a path with cutwidth d since it is a subgraph of \mathcal{H} . Next, invoke the result from [12] that shows how to embed $\mathcal{T}[u]$ into the d -lattice with unit dilation in such a way that node u is mapped to the origin of the lattice. Under this embedding, some node of $\mathcal{T}[u]$ must be embedded at a distance $\geq \rho_d(N)$ from the origin. It follows that $\rho_d(N)$ is a lower bound on the diameter of \mathcal{H} . \square

We now craft a more sophisticated lower-bound argument which closes the factor-of-2 gap between the upper and lower bounds of Lemma 3.1. The key to our argument is the following estimate for $\rho_d(N)$ whose proof is presented in Appendix A:

Lemma 3.2. *For any fixed d , $\rho_d(N) = \frac{1}{2}N^{1/d}(d!)^{1/d}(1 \pm o(1))$.*

We turn now to our diameter bounds.

Theorem 3.1. *For any fixed dimensionality d ,*

$$2^{1-2/d}\rho_d(N)(1 - o(1)) \leq \Delta_d^{(\text{HRN})}(N), \Delta_d^{(\text{MR})}(N) \leq 2\rho_d(N).$$

Proof. Retaining the upper bound of Lemma 3.1, we now prove our tighter lower bound for HRNs, leaving the easy adaptation of the bounding argument to MRs to the reader.

Let \mathcal{H} be any N -node depth- d HRN and let $V_1 = \{v_1, v_2, \dots, v_k\}$ be the node-set of \mathcal{H} 's level-1 subring. Each node $v_i \in V_1$ belongs to a sub-HRN \mathcal{H}_i of \mathcal{H} of depth $\leq d-1$. Note that all of \mathcal{H}_i 's other nodes (if there are any) come from higher-level subrings of \mathcal{H} . We consider two cases.

Some \mathcal{H}_i is "big." Say first that some

$$\mathcal{H}_i \in \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k\}$$

has at least $N^{1-1/(2d)}$ nodes. By Lemma 3.1, the diameter of \mathcal{H}_i is at least $\rho_{d-1}(N^{1-1/(2d)})$; by Lemma 3.2, therefore,

$$\begin{aligned} \frac{\rho_{d-1}(N^{1-1/(2d)})}{\rho_d(N)} &= \frac{N^{(2d-1)/(2d(d-1))}((d-1)!)^{1/(d-1)}(1 \pm o(1))}{N^{1/d}(d!)^{1/d}(1 \pm o(1))} \\ &= \Omega(N^{1/(2d(d-1))}) \\ &\geq 2(1 - o(1)). \end{aligned}$$

It follows that the diameter of \mathcal{H} , which is clearly no smaller than that of \mathcal{H}_i , is at least $2\rho_d(N)(1 - o(1))$.

All \mathcal{H}_i are "small." Say next that each

$$\mathcal{H}_i \in \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k\}$$

has fewer than $N^{1-1/(2d)}$ nodes. By the pigeon-hole principle, there must exist an $l \in \{1, 2, \dots, k\}$ such that the cumulative number of nodes in $\mathcal{H}_l, \mathcal{H}_{l+1}, \dots, \mathcal{H}_{l+k/2}$ is at least $N/2$. Further, since each \mathcal{H}_i is "small," there exist subgraphs

$$L \stackrel{\text{def}}{=} \mathcal{H}_l \cup \mathcal{H}_{l+1} \cup \dots \cup \mathcal{H}_j$$

and

$$R \stackrel{\text{def}}{=} \mathcal{H}_{j+1} \cup \mathcal{H}_{j+2} \cup \dots \cup \mathcal{H}_{l+k/2}$$

of \mathcal{H} , each of which has $\geq \frac{1}{4}N(1 - o(1))$ nodes. Importantly, for any pair of nodes, u from L and w from R , the shortest path between u and w in \mathcal{H} crosses edge (v_j, v_{j+1}) of \mathcal{H} 's level-1 ring. Consider now the breadth-first subtree $\mathcal{T}[v_j]$ of \mathcal{H} that is rooted at node v_j . Since $\mathcal{T}[v_j]$ can be laid out in a path with cutwidth d , it can be embedded into the d -lattice with unit dilation [12]. Since L has $\geq \frac{1}{4}N(1 - o(1))$ nodes, there exists a node u of L whose shortest path to v_j has length $\geq \rho_d(\frac{1}{4}N(1 - o(1)))$. Analogously, there exists a node w of R whose shortest path to v_{j+1} has length $\geq \rho_d(\frac{1}{4}N(1 - o(1)))$. Since the shortest path from u to w in \mathcal{H} must pass through edge (v_j, v_{j+1}) , the distance between u and w —hence, the diameter of \mathcal{H} —can be no smaller than $2\rho_d(\frac{1}{4}N(1 - o(1))) + 1$. We use Lemma 3.2 to bound this quantity relative to $\rho_d(N)$ as follows:

$$\begin{aligned} \frac{\rho_d((1/4)N(1 - o(1)))}{\rho_d(N)} &= \frac{((1/4)N(1 - o(1)))^{1/d}(d!)^{1/d}(1 \pm o(1))}{N^{1/d}(d!)^{1/d}(1 \pm o(1))} \\ &\geq 2^{-2d}(1 - o(1)). \end{aligned} \tag{3.1}$$

It follows from (3.1) that the diameter of \mathcal{H} is no smaller than $2^{1-1/(2d)}\rho_d(N)(1 - o(1))$, as was claimed. \square

3.2 The Impact of Noncrossing Hops

The chordal rings and express rings that we have studied here are unconventional due to our insistence that chords




Fig. 9. Resolving a Case 1 crossing.

and hops, respectively, not cross. We show now that this structural simplification cannot cost us more than a factor of two in diameter. This factor cannot be decreased in general. To wit, any *unit-diameter* N -node express (resp., chordal) ring is isomorphic to an N -node clique which is not outerplanar for any $N \geq 4$ [20].

Theorem 3.2. Any N -node cutwidth- c express ring \mathcal{R} , in which arcs are allowed to cross, can be transformed to a crossing-free cutwidth- c N -node express ring \mathcal{R}' such that

$$\text{Diameter}(\mathcal{R}') \leq 2 \cdot \text{Diameter}(\mathcal{R}).$$

Proof. We produce ring \mathcal{R}' from ring \mathcal{R} in two stages. First, we replace \mathcal{R} by a cutwidth- c N -node express ring \mathcal{R}'' whose hops form a tree (whose edges may cross). Specifically, we obtain \mathcal{R}'' by choosing an arbitrary node v of \mathcal{R} and growing a breadth-first spanning tree of hops outward from v . Each path within \mathcal{R}'' constitutes a walk up the tree to its root v , then, down the tree to the desired destination. Since \mathcal{R}'' is a *shortest-path* spanning tree of \mathcal{R} , it follows that $\text{Diameter}(\mathcal{R}'') \leq 2 \cdot \text{Diameter}(\mathcal{R})$. Next, we perform a series of transformations on \mathcal{R}'' to




Fig. 10. Resolving a Case 3 crossing.




Fig. 11. Resolving a Case 4 crossing.

eliminate all crossings of hops—without increasing either diameter or cutwidth. We eliminate crossings starting with the edges/hops that emanate from the root of the tree and proceed steadily down toward the leaves. Since each transformation replaces (at least) one offending hop with a shorter hop, no “uncrossing” step can create a new crossing. This fact guarantees that our algorithm terminates. Our “uncrossing” algorithm is the analog for bidirectional trees of the similarly motivated algorithm in [18]—which has many fewer cases since its trees each have a nonroot node lying to the right of its parent.

We describe our algorithm by specifying how it “uncrosses” two arbitrary crossing hops (a, c) and (b, d) while never increasing the *level* of any node, i.e., the node’s distance from the root. With no loss in generality (since the hops cross), assume that nodes $a, b, c,$ and d appear, in that order, in a clockwise traversal of the ring. Our primary breakdown into cases considers how hops (a, c) and (b, d) are oriented relative to the root of the spanning tree.

Case 1. $\text{level}(a) < \text{level}(c); \text{level}(b) < \text{level}(d)$; see Fig. 9a. We branch on the relative levels of nodes a and b .

- 1.1. $\text{level}(a) < \text{level}(b)$. Eliminate hop (b, d) ; add hop (c, d) ; see Fig. 9b.
- 1.2. $\text{level}(a) \geq \text{level}(b)$. Eliminate hop (a, c) ; add hop (b, c) ; see Fig. 9c.

Case 2. $\text{level}(a) > \text{level}(c); \text{level}(b) > \text{level}(d)$. This case is clearly symmetric to Case 1.

Case 3. $\text{level}(a) < \text{level}(c); \text{level}(b) > \text{level}(d)$; see Fig. 10a. We branch on the relative levels of nodes a and d .

- 3.1. $\text{level}(a) < \text{level}(d)$. Eliminate hop (d, b) ; add hop (c, b) ; see Fig. 10b.
- 3.2. $\text{level}(a) > \text{level}(d)$. This case is symmetric to Case 3.1.
- 3.3. $\text{level}(a) = \text{level}(d)$. Eliminate hops (a, c) and (b, d) ; add hops (a, b) and (d, c) ; see Fig. 10c.

Case 4. $\text{level}(a) > \text{level}(c); \text{level}(b) < \text{level}(d)$; see Fig. 11a. We branch on the relative levels of nodes b and c .

- 4.1. $\text{level}(b) \geq \text{level}(c)$. Eliminate hop (b, d) ; add hop (c, d) ; see Fig. 11b.
- 4.2. $\text{level}(b) \leq \text{level}(c)$. This case is symmetric to Case 4.1.

By the time we reach the leaves of the spanning tree, we shall have transformed the express ring \mathcal{R}'' to the

express ring \mathcal{R}' , which has no crossing hops and a diameter no greater than that of \mathcal{R}'' . \square

APPENDIX

ESTIMATING $\rho_d(N)$: THE PROOF OF LEMMA 3.2

We estimate the number of lattice points in a d -dimensional ℓ_1 -sphere $S_{\rho,d}$ of radius ρ . Note, first, that each orthant of the sphere has $\binom{d+\rho}{d}$ lattice points. Since there are 2^d orthants, $S_{\rho,d}$ contains no more than $2^d \binom{d+\rho}{d}$ lattice points. We can improve this estimate by using the inclusion-exclusion principle to subtract out points that appear in the pairwise intersection of the quadrants. We thereby find that $S_{\rho,d}$ contains *at least*

$$2^d \binom{d+\rho}{d} - 2^{d-1} d \binom{d+\rho-1}{d-1}$$

points since there are $d2^{d-1}$ pairs of intersecting quadrants and each intersection contains

$$\binom{d+\rho-1}{d-1}$$

points.

We now use our upper and lower bounds on the number of lattice points in $S_{\rho,d}$ to derive upper and lower bounds on $\rho_d(N)$. By definition of $\rho_d(n)$, for all d and N ,

$$2^d \binom{d+\rho_d(N)-1}{d} - 2^{d-1} d \binom{d+\rho_d(N)-2}{d-1} < N \leq 2^d \binom{d+\rho_d(N)}{d}.$$

Manipulating the lefthand side of the above inequality, we obtain

$$2^d \binom{d+\rho_d(N)-1}{d} \left(1 - \frac{d^2}{2(d+\rho_d(N)-1)}\right) < N \leq 2^d \binom{d+\rho_d(N)}{d}.$$

Bounding the lefthand and righthand sides of the above inequality, we obtain

$$2^d \frac{(\rho_d(N))^d}{d!} e^{-d^2/\rho_d(N)} < N \leq 2^d \frac{(\rho_d(N))^d}{d!} e^{d^2/\rho_d(N)}.$$

It follows that $\rho_d(N) = \frac{1}{2} N^{1/d} (d!)^{1/d} (1 \pm o(1))$. \square

ACKNOWLEDGMENTS

The research of A. Rosenberg was supported in part by US National Science Foundation (NSF) grants CCR-92-21785 and CCR-97-10367 and in part by US-Israel Binational Science Foundation grant 94-00266. The research of R. Sitaraman was supported in part by NSF Grant CCR-94-10077 and by NSF CAREER Award CCR-9703017. A portion of this research was done while F. Chung and A. Rosenberg were visiting Bellcore (now Telcordia Technologies) and a portion was done while A. Rosenberg was visiting the Department of Computer

Science at the Technion (Israel Institute of Technology). A portion of this paper was presented at the Sixth International Conference on Structural Information and Communication Complexity (SIROCCO'99).

REFERENCES

- [1] W. Aiello, S.N. Bhatt, F.R.K. Chung, A.L. Rosenberg, and R.K. Sitaraman, "Augmented Ring Networks," *Invited Talk, Int'l Conf. Math. Computer Modeling and Scientific Computing*, (no proceedings) 1997.
- [2] F.S. Annexstein, M. Baumslag, and A.L. Rosenberg, "Group Action Graphs and Parallel Architectures," *SIAM J. Computing*, vol. 19 pp. 544-569, 1990.
- [3] B.W. Arden and H. Lee, "Analysis of Chordal Ring Networks," *IEEE Trans. Computers*, vol. 30 pp. 291-295, 1981.
- [4] J.-C. Bermond, F. Comellas, and D.F. Hsu, "Distributed Loop Computers: A Survey," *J. Parallel and Distributed Computing*, vol. 24, pp. 2-10, 1995.
- [5] J.-C. Bermond, J.M. Fourneau, and A. Jean-Marie, "A Graph Theoretical Approach to the Equivalence of Multistage Interconnection Networks," *Discrete Applied Math.*, vol. 22, pp. 201-214, 1988/89.
- [6] S.N. Bhatt, F.R.K. Chung, J.-W. Hong, F.T. Leighton, B. Obrenic, A.L. Rosenberg, and E.J. Schwabe, "Optimal Emulations by Butterfly-Like Networks," *J. ACM*, vol. 43, pp. 293-330, 1996.
- [7] M. Blum and D. Kozen, "On the Power of the Compass," *Proc. 19th IEEE Symp. Foundations of Computer Science*, pp. 132-142, 1978.
- [8] B. Bollobás and F.R.K. Chung, "The Diameter of a Cycle Plus a Random Matching," *SIAM J. Discrete Math.*, vol. 1, pp. 328-333, 1988.
- [9] H. Burkhardt et al., "Overview of the KSR1 Computer System," Technical Report KSR-TR 9202001, Kendall Square Research, 1992.
- [10] S. Cosares, I. Saniee, and O. Wasem, "Network Planning with the SONET Toolkit," *Bellcore Exchange*, pp. 8-15, Sept./Oct. 1992.
- [11] W.J. Dally, "Express Cubes: Improving the Performance of k -Ary n -Cube Interconnection Networks," *IEEE Trans. Computers*, vol. 40, pp. 1016-1023, 1991.
- [12] Y. Dinitz, M. Feigelson, and S. Zaks, "On Optimal Graphs Embedded into Paths and Rings Using ℓ_1 -Spheres," *Proc. 23rd Int'l Workshop Graph-Theoretic Concepts in Computer Science*, 1997.
- [13] R. Feldmann and W. Unger, "The Cube-Connected Cycles Network Is a Subgraph of the Butterfly Network," *Parallel Processing Letters*, vol. 2, pp. 13-19, 1992.
- [14] A. Frank, T. Nishizeki, N. Saito, H. Suzuki, and É. Tardos, "Algorithms for Routing around a Rectangle," *Discrete Applied Math.*, vol. 40, pp. 363-378, 1992.
- [15] O. Gerstel, I. Cidon, and S. Zaks, "The Layout of Virtual Paths in ATM Networks," *ACM/IEEE Trans. Networking*, vol. 4, pp. 873-884, 1996.
- [16] O. Gerstel and S. Zaks, "The Virtual Path Layout Problem in Fast Networks," *Proc. 13th ACM Symp. Principles of Distributed Computing*, pp. 235-243, 1994.
- [17] O. Gerstel and S. Zaks, "The Virtual Path Layout Problem in ATM Networks," *Proc. First Int'l Colloquium Structure, Information and Comm. Complexity*, pp. 151-166, 1994.
- [18] O. Gerstel, A. Wool, and S. Zaks, "Optimal Layouts on a Chain ATM Network," *Discrete Applied Math.*, vol. 83, no. 1-3, pp. 157-178, 1998.
- [19] V.C. Hamacher and H. Jiang, "Comparison of Mesh and Hierarchical Networks for Multiprocessors," *Proc. Int'l Conf. Parallel Processing*, vol. I pp. 67-71, 1994.
- [20] F. Harary, *Graph Theory*. Reading, Mass.: Addison-Wesley, 1969.
- [21] C.-T. Ho and S.L. Johnson, "Embedding Meshes in Boolean Cubes by Graph Decomposition," *J. Parallel and Distributed Computing*, vol. 8, pp. 325-339, 1990.
- [22] R. Koch, F.T. Leighton, B.M. Maggs, S.B. Rao, A.L. Rosenberg, and E.J. Schwabe, "Work-Preserving Emulations of Fixed-Connection Networks," *J. ACM*, vol. 44, pp. 104-147, 1997.
- [23] E. Kranakis, D. Krizanc, and A. Pelc, "Hop-Congestion Tradeoffs for High-Speed Networks," *Int'l J. Foundations of Computer Science*, vol. 8, pp. 117-126, 1997.
- [24] C.P. Kruskal and M. Snir, "A Unified Theory of Interconnection Network Structure," *Theoretical Computer Science*, vol. 48, pp. 75-94, 1986.

- [25] D.-M. Kwai and B. Parhami, "Periodically Regular Chordal Rings: Generality, Scalability, and VLSI Layout," *Proc. Eighth IEEE Symp. Parallel and Distributed Processing*, pp. 148-151, 1996.
- [26] C. Lam, H. Jiang, and V.C. Hamacher, "Design and Analysis of Hierarchical Ring Networks for Shared-Memory Multiprocessors," *Proc. Int'l Conf. Parallel Processing*, pp. 46-50, 1995.
- [27] A. Litman and A.L. Rosenberg, "Balancing Communication in Ring-Structured Networks," Technical Report 93-80, Univ. of Massachusetts, 1993.
- [28] A. Michail, "Optimal Broadcast and Summation on Hierarchical Ring Architectures," *Parallel Processing Letters*, vol. 8, 1998.
- [29] E.J. Schwabe, "On the Computational Equivalence of Hypercube-Derived Networks," *Proc. Second ACM Symp. Parallel Algorithms and Architectures*, pp. 388-397, 1990.
- [30] Z.G. Vranesic, M. Stumm, D.M. Lewis, and R. White, "Hector: A Hierarchically Structured Shared-Memory Multiprocessor," *Computer*, vol. 24, no. 1, pp. 72-79, Jan. 1991.
- [31] S. Zaks, "Path Layout in ATM Networks—A Survey," *Proc. DIMACS Workshop Networks in Distributed Computing*, 1997.
- [32] C. Zamfirescu and T. Zamfirescu, "Hamiltonian Properties of Grid Graphs," *SIAM J. Discrete Math.*, vol. 5, pp. 564-570, 1992.



William Aiello received the BS degree in physics from Stanford University in 1982. After working in the Condensed Matter Group at the IBM Almaden Research Center, he entered the applied mathematics program at the Massachusetts Institute of Technology where he received the PhD degree in 1988. His thesis included research in cryptography, complexity theory, circuit lower bounds, and their relationships. In 1989, he was a postdoctoral fellow at both the MIT Laboratory for

Computer Science and the Computer Science Department of the Royal Institute of Technology, Sweden. From 1990 to 1998, he worked for Bellcore, first in the Parallel and Distributed Algorithms Department, concentrating on research in load balancing, fault tolerance, and backplane routing, and, later, in the Mathematics and Cryptography Department, working on zero-knowledge protocols and efficient cryptographic primitives, such as high-speed software encryption. In 1999, he joined AT&T Labs where he has worked in cryptography and network security and networking algorithms, including dynamic routing and congestion control. He has been active in developing the security standards for IP and voice services over hybrid IP/cable networks. For the past year, he has been the division manager of Secure Systems Research. He is a member of the IEEE and the IEEE Computer Society.



Sandeep Bhatt received his graduate and undergraduate degrees from the Massachusetts Institute of Technology. He is a principal research scientist at Akamai Technologies, Cambridge, Massachusetts. At Akamai, he leads the Performance Assurance Group which is responsible for end-to-end performance measurement and monitoring and analysis of network services. Prior to Akamai, he was on the research staff at Telcordia Technologies (formerly Bellcore) and on the

faculty at Yale University. His research interests include network performance and fault monitoring, the design of self-configuring systems for network security management, models and algorithms for distributed computing and communications, distributed simulations of network signaling and routing protocols, and N -body systems. His theoretical research has focused on the theory of graph embeddings with applications to VLSI circuit layout and parallel computing.



Fan R.K. Chung received the BS degree in mathematics from National Taiwan University in 1970 and the PhD degree in mathematics from the University of Pennsylvania in 1974, after which she joined the technical staff of AT&T Bell Laboratories. From 1983 to 1991, she headed the Mathematics, Information Sciences and Operations Research Division at Bellcore. In 1991, she became a Bellcore Fellow. In 1993, she was the Class of 1965 Professor of Mathematics at the the University of Pennsylvania. Since 1998, she has been a professor of mathematics and professor of computer science and engineering at the University of California, San Diego. Her research interests are primarily in graph theory, combinatorics, and algorithmic design, in particular in spectral graph theory, extremal graphs, graph labeling, graph decompositions, random graphs, graph algorithms, parallel structures, and various applications of graph theory in communication networks, software reliability, chemistry, engineering, and various areas of mathematics. She was awarded the Allendoerfer Award by Mathematical Association of America in 1990. Since 1998, she has been a member of the American Academy of Arts and Sciences.



Arnold L. Rosenberg received the AB degree in mathematics from Harvard College in 1962 and the AM and PhD degrees in applied mathematics from Harvard University in 1963 and 1966, respectively. He is a Distinguished University Professor of Computer Science at the University of Massachusetts at Amherst where he codirects the Theoretical Aspects of Parallel and Distributed Systems (TAPADS) Laboratory. Prior to joining the University of Massachusetts,

he was a professor of computer science at Duke University from 1981 to 1986 and a research staff member at the IBM T.J. Watson Research Center from 1965 to 1981. He has held visiting positions at Yale University and the University of Toronto. He was a Lady Davis Visiting Professor at the Technion (Israel Institute of Technology) and a Fulbright Research Scholar at the University of Paris-South. His research focuses on theoretical aspects of parallel architectures and communication networks with emphasis on developing algorithmic techniques for designing better networks and architectures and using them more efficiently. He is the author of more than 130 technical papers on these and other topics in theoretical computer science and discrete mathematics. He is the editor-in-chief of *Theory of Computing Systems* (formerly, *Mathematical Systems Theory*) and serves on other editorial boards. Information on his publications and other activities can be found at <http://www.cs.umass.edu/~rsnrbg/>. He is a fellow of the ACM, a fellow of the IEEE, a Golden Core member of the IEEE Computer Society, and a member of SIAM.



Ramesh K. Sitaraman received the BTech degree in electrical engineering from the Indian Institute of Technology, Madras. He obtained the PhD degree in computer science from Princeton University in 1993. He is an associate professor of computer science at the University of Massachusetts at Amherst where he is a member of the Theoretical Computer Science Group and codirects the Theoretical Aspects of Parallel and Distributed Systems (TAPADS) Laboratory. He

is currently at Akamai Technologies Inc., a company devoted to efficient and reliable delivery of content on the internet. His research focuses on fundamental issues in the design and use of communication networks and parallel and distributed systems. His specific interests include communication protocols, fault tolerance, and algorithms for delivering web content and streaming media in the internet. He is a recipient of a US National Science Foundation CAREER Award and a Lilly Fellowship. He serves on the editorial board of journals such as the *IEEE Transactions on Parallel and Distributed Systems*. He is a member of Sigma Xi, the ACM, the IEEE, and the IEEE Computer Society.