Optimal Turn Prohibition for Deadlock Prevention in Networks with Regular Topologies

Mark Karpovsky, *Fellow, IEEE,* Lev Levitin, *Life Fellow, IEEE,* and Mehmet Mustafa*, *Member, IEEE*

***Abstract*—In this paper we consider the problem of constructing minimal cycle-breaking connectivity preserving sets of turns for graphs that model communication networks, as a method to prevent deadlocks. Cycle-breaking provides for deadlock-free wormhole routing constrained by turns prohibited at some nodes. We present lower and upper bounds for minimal cardinalities of cycle-breaking connectivity preserving sets for several classes of graphs such as homogeneous meshes, p-ary n-cubes, cube-connected cycles, hexagonal and honeycomb meshes and tori, etc.**

***Index Terms*—deadlock, livelock, turn prohibition, wormhole routing.**

I. Introduction

BECAUSE of its simplicity, low channel setup times, and its high performance in delivering messages, wormhole routing has been widely investigated [1]–[5], and recently is being revisited for Networks-on-Chips technologies [6], [7]. Wormhole routing and its variants, [8] virtual cut through and pipelined circuit switching, PCS, have been used in regular topologies from chip-scale networks [6], [7], to rack-packed Blue Gene [9], to irregular topologies formed by interconnecting low-cost workstations in an ad hoc manner, forming what is referred to as Network of Workstations or NOWs [10]–[12]. Messages, also known as 'worms', are made up of flits that are transmitted atomicly, one flit at a time, from node to node in the network. The header flit, containing the destination address is immediately followed by the payload or data flits [5]. One aspect that makes wormhole routing and routers attractive is that each channel requires buffers that are only a few flits deep [13], [14]. In wormhole routed networks, messages traverse the network in a pipelined fashion, such that parts of the message occupy different network resources, while the header flit requests yet other resources. When there is no contention, such as in lightly loaded networks, the latency of message delivery varies very slowly with the distance [5]. When a message is blocked, the header and the rest of the message wait until the blockage is removed. In wormhole routed networks messages could hold potentially large number of network resources while attempting to reserve others. In congested networks with high injected traffic, improperly designed routing protocols can lead to a network state, in which no progress can be made in delivering, not only of the current messages but all subsequent messages in the network. This network state, in which worms are in a cyclic dependency of each other's held-up resources, is known as deadlock.

M. Karpovsky, (617) 353-9592, L. Levitin, (617) 353-4607 and M. Mustafa, (781) 647-1359 are with the Department of Electrical and Computer Engineering, Boston University, 8, St. Mary's Street, Boston, MA 02215 USA, e-mail: {markkar, levitin, mmustafa}@bu.edu.

The work of the first author is sponsored by the NSF grant CNS 1012910.

(*) Communicating Author.

Manuscript submitted August 24, 2013.

When a message from a source is intended to be sent to a single destination, the delivery mechanism of such a message is known as unicast. Cycles in channel dependency graphs (CDG), have been identified as the root cause of deadlocks in wormhole networks for unicast messaging. In [15], [16], it has been shown that necessary and sufficient condition for eliminating deadlocks is the elimination of cycles in the corresponding CDG. This condition is equivalent to elimination of all "cycles of edges" (as defined below) in the original undirected network graph. This is the approach used in the present paper.

Considerable body of work has been dedicated to designing wormhole routing algorithms that prevent deadlocks from occurring [1], [13], [15], [17], [5], [18], [19], [20]–[23]. In these proactive deadlock prevention schemes, either virtual channels were added [13], [22], or some resources were prevented from being used.

To provide deadlock-free adaptive routing, Glass and Ni [18], presented a method that requires neither additional physical nor virtual channels. The turn model is based on analyzing the directions in which packets can turn in regular networks and the cycles that the turns can form. Prohibiting just sufficient number of turns to break all of the cycles, produces a routing algorithm that is deadlock free and livelock free.

The motivation for seeking the minimal fraction of prohibited turns is originally due to Glass and Ni [18]. They have found that reduction in the number of prohibited turns results in a decrease of average path length and the average message delivery time, thereby increasing the throughput. This conclusion was confirmed by other authors [24], [25] for irregular topologies as well. Experimental data show that there is a considerable gain of approximately 7-8% in the maximum sustainable throughput in the network, for each percentage point reduction in the fraction of prohibited turns.

The simplest deadlock prevention approach utilizes spanning tree based routing for message delivery. Since messages propagate along the tree edges deadlocks are prevented from occurring. However, in this approach a large number of network communication channels are not being used as they are not a part of the spanning tree. This is not only inefficient and ineffective use of the available resources but can also lead to hot spots in the network close to the root of the spanning tree.

The Up/Down algorithm, first introduced in Autonet [26] routing algorithm improves the shortcoming of the spanning tree approach by using the cross links, non-tree links, under certain conditions. Nevertheless the Up/Down approach still suffers from the other shortcomings of the spanning tree based approach.

Sancho and Robles [27] explored all spanning trees with every node acting as the root node and then selected the best tree. The selection of the best tree and root combination is accomplished by two heuristic rules with a run-time complexity of $O(N^3)$, where N is the number of switches in the network.

In [28] a version of the turn prohibition algorithm was used that enabled generalizing the application of Network Calculus to arbitrary topologies.

Virtual channels have been introduced and considered as a tool to avoid deadlocks in a number of papers [4], [29], [30], often in combination with the dimension-ordered routing (DOR) technique [31]. However, as pointed out in [29], the use of virtual channels may have a negative effect on the message latency.

A variant of turn prohibition algorithm, called Tree-Based Turn-Prohibition, TBTP, where it [32] has been shown to have polynomial-time complexity and to be backward compatibile with the IEEE 802.1d standard. Authors claim that the throughput has been increased by a factor of up to 2.48.

A distributed version of the TBTP algorithm is reported in [33]. With an upper bound of $1/2$ for the fraction of prohibited turns, the shortcoming of the TBTP approach is that it could potentially restrict the use of a large number of turns.

A hybrid methodology using both proactive and reactive approaches was proposed in [34] , in which, routing restrictions are adjusted dynamically based on network congestion.

Another class of deadlock-preventing algorithms, the so-called, tree-free cycle-breaking algorithms, was developed in [19]–[21], [24], [25], [35], [36]. These algorithms (TP and SCB) have been proved to create a minimal (irreducible) set of prohibited turns the size of which never exceeds $1/3$ of the total number of turns in any graph. They have been shown to outperform the tree-based algorithms with respect to three basic characteristics: fraction of prohibited turns, distance dilation, latency and the saturation load. For some broad classes of network topologies, those algorithms provide an optimum solution of the turn prohibition problem [36]. The computational complexity of the tree-free algorithms is $O(N^2\Delta)$, where Δ is the maximum node degree (number of neighbors) in the graph. The algorithms are topology agnostic. However, the application of those general algorithms may be still unnecessarily complex in the case of graphs with certain regularities in their structure.

This paper deals with certain classes of networks with regular topologies. Here, we do not use general algorithms developed previously for arbitrary topologies, e.g., [1]–[5], [13], [15], [17], [19]–[22], [26], [27], [32], [33], [37], [38], [24], [25], [35], [36]. Instead, we present optimal or asymptotically optimal solutions of the turn prohibition problem for general classes of special topologies. These solutions are obtained by application of simple rules, run-time complexity of which does not exceed $O(N)$ (i.e., linear in the number of nodes N), and, in many cases, is $O(1)$ (i.e., constant). The memory requirements for computing the solutions do not exceed $O(\log N)$. The proposed turn prohibition rules can be easily implemented for execution in a distributed way.

It should be pointed out that turn prohibition algorithms are, in fact, pre-routing procedures; they do no prescribe any specific routing policy, but just restrict the set of permitted turns in routing tables. Therefore, they are compatible with any routing algorithm, in particular, with the fully adaptive minimal routing (of course, paths that include prohibited turns are excluded from consideration).

A few particular regular topologies have been considered in several papers [18], [39]–[44], [13]. This paper presents methods applicable to a number of classes of popular regular graphs, such as homogeneous meshes, p-ary n-cubes, cube connected cycles, hexagonal and honeycomb meshes and tori.

The dimension-ordered routing (DOR) [31] has been popular for meshes. However, as shown in Section III, the use of DOR algorithm results in prohibition of much larger fraction of turns in the network than the approach developed in the present paper. For multi-dimensional meshes, the fraction of turns prohibited by DOR tends to $1/2$. Our methods guarantee that the fraction of prohibited turns never exceeds $1/4$. We note also that for the DOR approach in meshes some of the messages will not be delivered even with just a single link failure. For the routing techniques based on turn prohibition approach described in this paper, in the case of n-dimensional meshes all messages will be delivered as long as the number of faulty links does not exceed $n-1$. Thus the proposed techniques provide for a higher reliability than DOR.

Section II includes definitions, notations and lower bounds on the number and the fraction of prohibited turns. Then we introduce and analyze embedded graphs and homogeneous meshes in Section III followed by analysis of a number of well known regular topologies in Section IV. Finally in Section V we discuss dilation as a result of turn prohibitions and present our conclusions in Section VI.

II. Definitions, Notations, and Lower Bounds

Similar to Duato, Glass & Ni and others [4], [14], [15], [18], throughout this paper, we use the general abstract model of a communication network as an undirected graph, in which, every node incorporates a local processor and a router. Nodes are interconnected via full duplex and symmetric communication channels. We note that in our model, the graph representing the network is not a CDG but the undirected graph representing the topology of the network. The transfer of flits which takes place over the communication links, are under tight wormhole-handshaking protocol between the nodes. When a header flit arrives at a router, either from the local processor or from an adjacent node, the router would determine if the outgoing communication channel necessary to forward the flit is busy or not. If the channel is busy, the flit waits until the channel is freed up. If and when the communication channel is freed up, the message with the waiting flit takes ownership of the channel and the flit is transferred to the adjacent node. In this paper we are dealing only with the wormhole routed networks.

We do not consider resource dependent or message dependent deadlocks [45] which may appear associated with IP blocks in NOCs. We do not consider additional features, such as in-network synchronization, priority discipline, credit logic, arbiters, virtual channels, and other more complex methods of fighting deadlocks, as in [46], [47]. Strategies that include additional flit buffers or additional control circuitry that are used to implement virtual channels or escape ports or gateways to avoid deadlocks are viewed as costly approaches.

Let us consider an undirected connected graph $G(V, E)$, with $N = |V|$ vertices (nodes), denoted by $a, b, \ldots$, and $M = |E|$ edges, denoted by (a, b), etc, to represent a communication network. Here, each node $a \in V$ represents a router and a processor, and each edge $(a, b) \in E$ represents a bidirectional communication link between nodes a and b. A **turn** in G is a triplet of nodes (a, b, c) if (a, b) and (b, c) are edges in G and $a \neq c$. In an undirected graph turns (a, b, c) and (c, b, a) are considered to be the same turn. If the degree of node j is d_j, the total number of turns $T(G)$ in G is given by $T(G) = \sum_{j=1}^{N} \binom{d_j}{2}$. A **path** $P = (v_0, v_1, \ldots, v_{L-1}, v_L)$ of length L, $L \geq 1$ from node a to node b in G is a sequence of nodes $v_i \in V$ such that, $v_0 = a$ and $v_L = b$, and every two consecutive nodes are connected by an edge. Subsequences of the form (v_i, v_k, v_i) are not permitted in a path. Nodes and edges in the path are not necessarily all different. A turn (a, b, c) **belongs** to path $P = (v_0, v_1, \ldots, v_L)$ if $(a, b, c) = (v_i, v_{i+1}, v_{i+2})$, $i = 0, \ldots, L-2$. A set of turns $W(G)$ is called the **set of prohibited turns** and any path that includes turns from $W(G)$ cannot be used for communication (such a path is called **prohibited**). This set is called **connectivity preserving**, if for any $a, b \in V$ there exists a path in G that is not prohibited. Path $P = (v_0, v_1, v_2, \ldots, v_k, v_0, v_1)$ in G is called a **cycle**. If no proper subset of nodes of cycle P forms a cycle, we call P a **simple cycle**. Set $W(G)$ of prohibited turns in G is called **cycle-breaking** if every cycle in G includes at least one turn from $W(G)$. The minimum cardinality of connectivity preserving set $W(G)$ for a given graph G is denoted by $Z(G)$ and the minimum fraction of prohibited turns is denoted by $z(G) = Z(G)/T(G)$. Since prohibition of turns imposes routing constraints, by preventing certain communication paths from being used during the routing of messages in the network, it must be done in a way that minimizes the fraction of link pairs (i.e. turns) that are prevented from being used.

In Fig. 1 these concepts are illustrated using a simple example of a graph with nine nodes and 13 turns. The turn at node a for example is (c, a, d). If a turn is prohibited we denote such a turn graphically as an arc drawn centered at the node and ending on the two edges of the turn. For example, the turn (a, d, h) is shown to be prohibited. A path from node a to node f is $P = (a, b, e, f)$, and an alternative but non-minimal path is $P = (a, c, h, d, e, f)$. A simple cycle in the graph would be $C_1 = (a, d, h, c, a, d)$. Note that in our definition of a cycle, one edge, here the edge (a, d) is repeated at most once in the same direction as the path is traversed. Because of the prohibited turn at node d, if a message is to be routed from node a to node h it would have to be sent out via node c. Since the cycle C_1 includes the turn (a, d, h) it is a cycle-breaking turn. Note that if we prohibit an additional turn at node f, namely the turn (b, f, k), we would have a set of prohibited turns $W_1(G) = \{(a, d, h), (b, f, k)\}$. This set would break the two simple cycles but it does not break all cycles. A set of prohibited turns $W_2(G) = \{(a, d, h), (d, e, f), (b, f, k)\}$ is a cycle breaking set but it is not a connectivity preserving set since it disconnects the graph. Finally the set $W_3(G) = \{(a, d, h), (f, b, g)\}$ is a minimal connectivity preserving cycle-breaking set of turns.

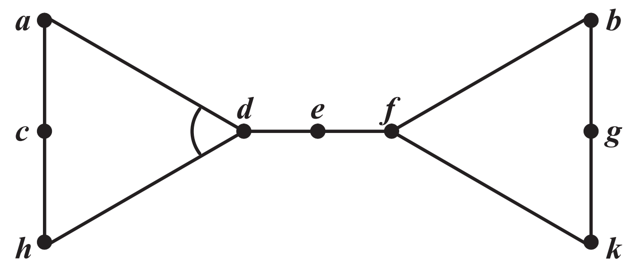

Fig. 1. An example of a simple network with $N = 9$ nodes, $M = 10$ bidirectional communication links and $T = 13$ turns.

Let G be a connected graph with minimum degree δ. Consider a set of R cycles in G such that no more than r cycles are covered by the same turn. Then [25], the number of prohibited turns $Z(G)$ and fraction of prohibited turns $z(G)$ satisfy the following inequalities:

$$Z(G) \geq M - N + 1, \tag{1}$$

$$z(G) \geq \frac{R}{rT(G)}, \tag{2}$$

and

$$Z(G) \geq M - N + \binom{\delta - 1}{2} + 1, \quad \delta > 2. \tag{3}$$

Bound (3) is tight. For example, in the Petersen graph (see Fig. 11 b) with $M = 15$, $N = 10$, and $\delta = 3$, the number of prohibited turns, $Z(G) = 7$.

III. Embedded Graphs and Homogeneous Meshes

Consider a graph $G = (V, E)$ which is embedded in the n-dimensional real space $\mathbb{R}^n$, so that each node $\mathbf{x}$ is a point in $\mathbb{R}^n$.

Definition 1: Given a neighborhood set $D = \{\pm \mathbf{a}_i, i = 1, \ldots, t\}$, where $\mathbf{a}_i$ are vectors in $\mathbb{R}^n$, an embedded graph G is a ***homogeneous mesh***, if each node $\mathbf{x}$ has a degree $d = 2t$, and if $\mathbf{x} \in V$, then its neighbors are nodes $\mathbf{x} \pm \mathbf{a}_i$, $i = 1, 2, \ldots, t$.

For example, for an infinite 2-D Mesh, $D = \{\pm \mathbf{a}_1, \pm \mathbf{a}_2\}$ where $\mathbf{a}_1 = (0, 1)$, $\mathbf{a}_2 = (1, 0)$.

Several important topologies, such as multi-dimensional meshes and tori, can be embedded into n-dimensional real spaces and can be considered as homogeneous meshes.

We call $\mathbf{a} \in D$ **positive**, $\mathbf{a} > 0$, if the first non-zero component of $\mathbf{a}$ is positive, otherwise $\mathbf{a}$ is negative, $\mathbf{a} < 0$. For example, in a two dimensional space,$(0, 1) > 0$, and $(-1, 1) < 0$.

A. Infinite Meshes

Consider the following turn prohibition rule for homogeneous meshes. Turn $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = (\mathbf{x}_2 - \mathbf{x}_1, \mathbf{x}_3 - \mathbf{x}_2)$ is prohibited *iff* $\mathbf{x}_2 - \mathbf{x}_1 < 0$ and $\mathbf{x}_3 - \mathbf{x}_2 < 0$. Let $W(M_D)$ be a set of prohibited turns for a homogeneous mesh M_D.

Theorem 1: As described, the turn prohibition rule has the following properties.

1) For any mesh M_D and any $\mathbf{x}, \mathbf{y} \in V$ there exists a path from $\mathbf{x}$ to $\mathbf{y}$ not containing any turns from $W(M_D)$.
2) For any cycle in M_D there exists a turn which belongs to the cycle and also belongs to $W(M_D)$, the set of prohibited turns.
3) The set of prohibited turns is minimum
4) The minimum fraction of prohibited turns for a homogeneous mesh M_D with size of D equal to d is

$$z(G) = \frac{1}{4}\left(1 - \frac{1}{d-1}\right). \tag{4}$$

Proof:

1) Consider a path $P = (\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k)$ from node $\mathbf{x}_0$ to $\mathbf{x}_k$, where $\mathbf{x}_{i+1} = \mathbf{x}_i + \mathbf{b}_i$; $i = 0, \ldots, k-1$, $\mathbf{b}_i \in D$. The corresponding sequence of edges is $S = (\mathbf{b}_0, \mathbf{b}_1, \ldots, \mathbf{b}_{k-1})$. Note that path P is prohibited *iff* there exists a pair of consecutive edges $(\mathbf{b}_{i-1}, \mathbf{b}_i)$ in S such that $\mathbf{b}_{i-1} > 0$ and $\mathbf{b}_i < 0$. It follows from Definition 1 that if S forms a path from $\mathbf{x}_0$ to $\mathbf{x}_k$, then any permutation of $\mathbf{b}_0, \mathbf{b}_1, \ldots, \mathbf{b}_{k-1}$ also corresponds to a path from $\mathbf{x_0}$ to $\mathbf{x_k}$, since the mesh is homogeneous and $\mathbf{x}_k = \mathbf{x}_0 + \sum_{i=0}^{k-1} \mathbf{b}_i$. Then there exists a permutation $S' = (\mathbf{b'}_0, \mathbf{b'}_1, \ldots, \mathbf{b'}_{k-1})$ of S in which all negative vectors (if any) appear before all positive ones (if any). The corresponding path $P' = (\mathbf{x}_0, \mathbf{x'}_1 = \mathbf{x}_0 + \mathbf{b'}_0, \ldots, \mathbf{x}_k = \mathbf{x'}_{k-1} + \mathbf{b'}_{k-1})$ has no prohibited turns and thus, nodes $\mathbf{x}_0$ and $\mathbf{x}_k$ are connected.
2) Consider a cycle $C = (\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k, \mathbf{x}_0, \mathbf{x}_1)$ and the corresponding cycle of edges $S = (\mathbf{b}_0, \mathbf{b}_1, \ldots, \mathbf{b}_k, \mathbf{b}_0)$, where $\mathbf{b}_i = \mathbf{x}_{i+1} - \mathbf{x}_i$, $i = 0, 1, \ldots, k-1$; $\mathbf{b}_k = \mathbf{x}_0 - \mathbf{x}_k$. Note that $\sum_{i=0}^{k} \mathbf{b}_i = 0$. Therefore, among vectors $\mathbf{b}_0, \mathbf{b}_1, \ldots, \mathbf{b}_k$ both positive and negative ones must exist. Since sequence S starts and ends with the same vector (either positive or negative), it must include at least one pair $\mathbf{b}_{i-1}, \mathbf{b}_i$, where $\mathbf{b}_{i-1}$ is positive and $\mathbf{b}_i$ is negative. Thus, the corresponding cycle is prohibited.
3) Let us consider cycles of length four, $C = (\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_0, \mathbf{x}_1)$, where $\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{b}_0$, $\mathbf{x}_2 = \mathbf{x}_1 + \mathbf{b}_1$, $\mathbf{x}_3 = \mathbf{x}_2 - \mathbf{b}_0 = \mathbf{x}_3 + \mathbf{b}_1$. All sets of turns corresponding to different choices of nodes $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ are disjoint. Hence, in order to break all cycles, it is necessary to prohibit at least one turn in each of such cycles. Indeed, according to our prohibition rule, in the sequence of edges $(\mathbf{b}_0, \mathbf{b}_1, -\mathbf{b}_0, -\mathbf{b}_1, \mathbf{b}_0)$ exactly one turn is prohibited (e.g., if $\mathbf{b}_0, \mathbf{b}_1 > 0$, then turn $(\mathbf{b}_1, \mathbf{b}_0)$ is prohibited). Thus, the set of prohibited turns is the smallest possible.
4) Obviously, in the set $D = \{\pm \mathbf{a}_i, i = 1, \ldots, t\}$ exactly $t = \frac{d}{2}$ vectors are positive, and the other half are negative. Therefore

$$z(G) = \frac{\binom{d/2}{2}}{\binom{d}{2}} = \frac{1}{4}\left(1 - \frac{1}{d-1}\right).$$

■

Remarkably, the result (4) does not depend on the choice of the coordinate system and on the particular topology of the mesh. For example, Fig. 2 shows two different non-isomorphic topologies which have the same node degree d and, thus, the same $z(G)$.

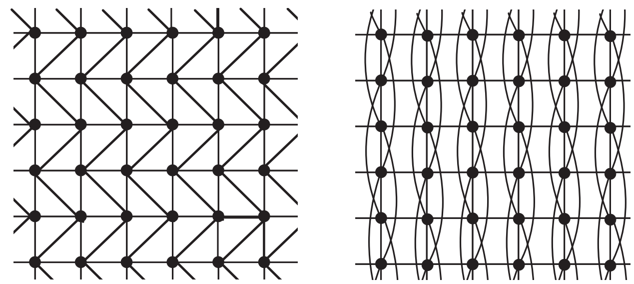

Fig. 2. Different non-isomorphic topologies with the same degree $d = 6$ have the same $z(G)$.

It is interesting to compare (4) with the fraction of prohibited turns when one uses the popular DOR algorithm [31].

For the case of an n-dimensional mesh ($d = n$) the fraction of prohibited turns given by (4) is $\frac{n-1}{2(2n-1)}$. The DOR algorithm prohibits a portion of the turns equal to $\frac{n-1}{2n-1}$, i.e., twice as large as our approach. We note also that for the DOR approach in meshes some of the messages will not be delivered even with just a single link failure. For the routing techniques based on turn prohibition approach described in this paper, in the case of n-dimensional meshes all messages will be delivered as long as the number of faulty links does not exceed $n-1$. Thus the proposed techniques provide for a higher reliability than DOR.

A more general situation can be described as follows. Consider an embedded graph $G = (V, E)$ that consists of m different types of nodes, $V = \bigcup_{k=1}^{m} V_k$ such that all nodes of type k have the same degree d_k, and if $\mathbf{x} \in V_k$, then its neighbors are $\mathbf{x} + \mathbf{a}_{ki}$, $i = 1, 2, \ldots, d_k$. Let $d_k = d_k^{(+)} + d_k^{(-)}$, where $d_k^{(+)}$ and $d_k^{(-)}$ are the numbers of positive and negative vectors, respectively, in the set $A_k = \{\mathbf{a}_{ki}\}$. We call such embedded graphs *multicomponent meshes*.

Suppose we prohibit all turns $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$, such that $\mathbf{x}_1 - \mathbf{x}_2 < 0$ and $\mathbf{x}_3 - \mathbf{x}_2 < 0$, or, alternatively, such that $\mathbf{x}_1 - \mathbf{x}_2 > 0$ and $\mathbf{x}_3 - \mathbf{x}_2 > 0$. Let us call such turns "negative" or, respectively, "positive". Assuming that the connectivity is preserved and following the same reasoning, as in the proof of Theorem 1, we obtain Corollary 1.

Corollary 1: Prohibition of all negative or of all positive turns in graph G described above breaks all the cycles in G. The fraction of prohibited turns $z(G)$ obeys an upper bound

$$z(G) \le \frac{\min\left\{\sum_{k=1}^{m} \rho_k \binom{d_k^{(-)}}{2}, \sum_{k=1}^{m} \rho_k \binom{d_k^{(+)}}{2}\right\}}{\sum_{k=1}^{m} \rho_k \binom{d_k}{2}}, \tag{5}$$

where ρ_k is the density of nodes of type k.

Here, as usual (c.f. [48]), the density ρ_k of a subset V_k of nodes in an infinite embedded graph $G(V, E)$ is defined as follows. Consider a ball $B(R)$ of radius R in $\mathbb{R}^n$. Then

$$\rho_k = \limsup_{R\to\infty} \frac{|V_k \bigcap B(R)|}{|V \bigcap B(R)|}.$$

Note that if $\mathbf{y} = \mathbf{x}+\mathbf{a}$, where $\mathbf{a} > 0$, then $\mathbf{x} = \mathbf{y}+\mathbf{b}$, where $\mathbf{b} = -\mathbf{a} < 0$. Therefore, $\sum_{k=1}^{m} \rho_k d_k^{(+)} = \sum_{k=1}^{m} \rho_k d_k^{(-)}$. However, for some structures prohibition of positive vs. negative turns can give rather different results, as shown by Example 1.

Example 1: The embedded graph in Fig. 3 has three different types of nodes with degrees 2, 3, and 5, each with a density of $\rho = 1/3$. As shown in the enlarged view, all positive turns prohibited at the node of degree 5, and all negative turns prohibited at nodes of degree 2 and degree 3. Prohibition of negative and positive turns yields different fractions of prohibited turns equal to $3/7$ and $1/7$, respectively.

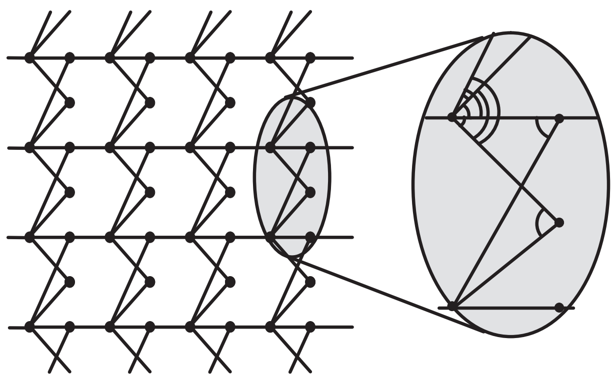

Fig. 3. A multicomponent mesh with three different types of nodes of degrees 2, 3, and 5. In the enlarged view we show all positive turns prohibited at the node of degree 5, and all negative turns prohibited at nodes of degree 2 and degree 3.

Example 2: The embedded graph that we call the "Brick Mesh" shown in Fig. 4. There are six types of nodes in this mesh; type 1, type 2, type 3, and type 4 nodes are of degree 3, and type 5, and type 6 nodes are of degree 4, as shown in the insert. The densities of type 1 and type 6 are each $3/14$ and the density of each of the others is $1/7$. If we consider the prohibition of the negative turns as shown in the enlarged view in Fig. 4, we determine that the fraction of prohibited turns is $z = 23/84$. The prohibition of positive turns gives a different result: $z = 3/14$.

Another interesting topology is the honeycomb mesh (see Section IV, Fig. 8b).

In general, the bounds in (5) depend on the choice of the coordinate system, in particular, on the order of the coordinates.

Note also that the prohibition rule given above for a multicomponent mesh does not guarantee, in general, the preservation of connectivity. However, it can be shown that for a two-component mesh ($m = 2$) connectivity is always preserved, provided that $d_k^{(+)} > 0$ and $d_k^{(-)} > 0$ for $k = 1, 2$. For example, for the honeycomb mesh (Fig. 8b), $m = 2$, $d_1^{(+)} = 2$, $d_1^{(-)} = 1$, $d_2^{(+)} = 1$, $d_2^{(-)} = 2$, $\rho_1 = \rho_2 = 1/2$ and $z(G) = 1/6$ (see Section IV).

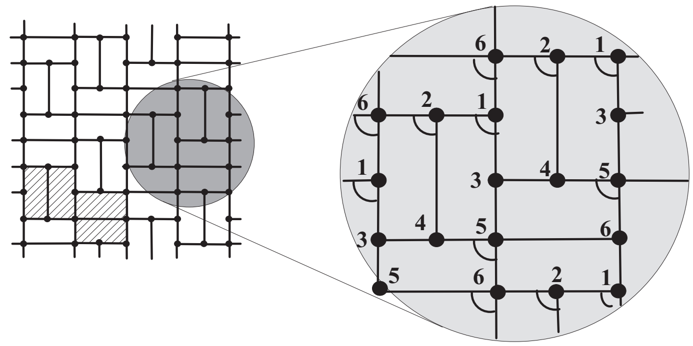

Fig. 4. A multicomponent brick mesh in which six different node types are identified in the enlarged view by the numbers adjacent to the nodes.

B. Finite and Wraparound Meshes

Homogeneous meshes considered so far in this section are of infinite extent with infinite number of nodes. We will define now finite D-Meshes $M_D(p_1, \ldots, p_n)$ and finite wraparound D-meshes $M_D^W(p_1, \ldots, p_n)$.

Let $D = \{\pm\mathbf{a}_1, \pm\mathbf{a}_2, \ldots, \pm\mathbf{a}_t\}$, $\mathbf{a}_i \in \mathbb{R}^n$, $i = 1, 2, \ldots, t$ and $d = 2t$ be the degree of every node. Then $n \le t$, (otherwise the mesh can be embedded in a space of a smaller dimensionality), and there are n linearly independent vectors in D. Henceforth we will assume that there exists a basis $B = \{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$, $B \subseteq D$ such that any point in the mesh can be represented as a linear combination of vectors from B with integer coefficients. Denote $\mathbf{C} = \mathbf{A}^{-1}$ where $\mathbf{A}$ is the matrix with columns $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$. Then any node $\mathbf{x}$ in the mesh can be represented in basis B as $\widetilde{\mathbf{x}} = \mathbf{C}\,\mathbf{x} = (\tilde{x}^{(1)}, \tilde{x}^{(2)}, \ldots, \tilde{x}^{(n)})$, where all $\tilde{x}^{(i)}$ are integers, $i = 1, 2, \ldots, n$.

Let $p_1, p_2, \ldots, p_n$ be positive integers, $p_i \ge 2$, $i = 1, 2, \ldots, n$.

Definition 2: A graph $G(V, E)$ is a finite D-mesh $M_D(p_1, p_2, \ldots, p_n)$ if $V = \{\mathbf{x} \mid \widetilde{x}^{(i)} \in \{0, 1, \ldots, p_i - 1\},\ i = 1, \ldots, n\}$. Then $(\mathbf{x}, \mathbf{y}) \in E$ if $\mathbf{C}\ (\mathbf{x} - \mathbf{y}) \in D_{\mathbf{C}}$ or $\mathbf{C}\ (\mathbf{y} - \mathbf{x}) \in D_{\mathbf{C}}$, where $D_{\mathbf{C}} = \{\pm\mathbf{C}\mathbf{a}_i \mid i = 1, \ldots, t\}$ $= \{\pm(1, 0, 0, \ldots, 0), \pm(0, 1, 0, \ldots, 0), \ldots, \pm(0, 0, 0, \ldots, 1), \pm\mathbf{C}\mathbf{a}_{n+1}, \ldots, \pm\mathbf{C}\mathbf{a}_t\}$.

Example 3: Let $n = 2$ and $D = \{\pm\mathbf{a}_1, \pm\mathbf{a}_2, \pm\mathbf{a}_3\} = \left\{\pm\left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm(1, 0)\right\}$. Note that $\mathbf{a}_3 = \mathbf{a}_1 - \mathbf{a}_2$, and

$$\mathbf{A} = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ \frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} \end{bmatrix},$$

$$\mathbf{C} = \begin{bmatrix} 1 & \frac{\sqrt{3}}{3} \\ -1 & \frac{\sqrt{3}}{3} \end{bmatrix},$$

and $D_{\mathbf{C}} = \{\pm(1, 0),\ \pm(0, 1),\ \pm(1, -1)\}$. The finite mesh $M_D(5, 3)$ is shown in Fig. 5.

Next we define finite wraparound meshes $M_D^W(p_1, p_2, \ldots, p_n)$. Let p_i be positive integers larger than 2. We will also assume that for the set $D = \{\pm\mathbf{a}_1, \pm\mathbf{a}_2, \ldots, \pm\mathbf{a}_t\}$, $(\mathbf{a}_i \in \mathbb{R}^n, n \le t)$, vectors $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n\}$ are linearly independent and each $\mathbf{a}_{n+j} = \sum_{i=1}^{n} u^{(i)} \mathbf{a}_j$ $(j = 1, \ldots, t-n)$, where $c^{(i)}$ are integers, such that $|u^{(i)}| \le p_i - 1$. Let $\mathbf{U}_1 = \left(u_1^{(1)}, u_1^{(2)}, \ldots, u_1^{(n)}\right)$ and $\mathbf{U}_2 = \left(u_2^{(1)}, u_2^{(2)}, \ldots, u_2^{(n)}\right)$

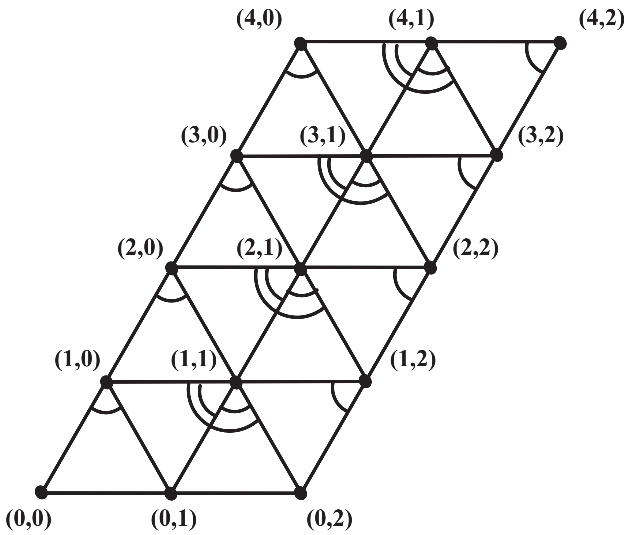

Fig. 5. A finite D-Mesh $M_D(5,3)$ with $D = \left\{\pm\left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm(1,0)\right\}$ and $D_{\mathbf{C}} = \{\pm(1,0), \pm(0,1), \pm(1,-1)\}$.

be vectors with $u_1^{(i)}, u_2^{(i)} \in \{0,1,\ldots,p_i-1\}$. Denote $\mathbf{U}_3 = \mathbf{U}_1 \oplus \mathbf{U}_2$, if $u_3^{(i)} = u_1^{(i)} + u_2^{(i)} \bmod p_i$, $i = 1,2,\ldots,n$.

Definition 3: A graph $G(V,E)$ is a wraparound D-Mesh $M_D^W(p_1,p_2,\ldots,p_n)$ if $V = \{\mathbf{x} \mid \widetilde{x}^{(i)} \in \{0,1,\ldots,p_i-1\},\ i = 1,2,\ldots,n\}$ and the edge $(\mathbf{x},\mathbf{y}) \in E$ if there exists a vector $\mathbf{h}$ such that $\widetilde{\mathbf{x}} \oplus \widetilde{\mathbf{h}} = \widetilde{\mathbf{y}}$, and $\widetilde{\mathbf{h}} = \widetilde{\mathbf{b}}$ for some $\widetilde{\mathbf{b}} \in D$. (Here, $\widetilde{\mathbf{x}} = \mathbf{C}\,\mathbf{x}$, $\widetilde{\mathbf{h}} = \mathbf{C}\,\mathbf{h}$, $\widetilde{\mathbf{y}} = \mathbf{C}\,\mathbf{y}$, and $\widetilde{\mathbf{b}} = \mathbf{C}\,\mathbf{b}$.)

Example 4: Let $n = 2$ and $D = \{\pm\mathbf{a}_1, \pm\mathbf{a}_2, \pm\mathbf{a}_3\} = \left\{\pm\left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm(1,0)\right\}$. As in Example 3, select $\mathbf{a}_1 = (\frac{1}{2}, \frac{\sqrt{3}}{2})$ and $\mathbf{a}_2 = (-\frac{1}{2}, \frac{\sqrt{3}}{2})$. Then

$$\mathbf{C} = \begin{bmatrix} 1 & \frac{\sqrt{3}}{3} \\ -1 & \frac{\sqrt{3}}{3} \end{bmatrix},$$

$\mathbf{a}_3 = \mathbf{a}_1 - \mathbf{a}_2$, and $\mathbf{C}\,\mathbf{a}_3 = \widetilde{\mathbf{a}}_3 = (1,-1)$. With this neighborhood definition, the wraparound mesh $M_D^W(5,5)$ is shown in Fig. 6. This wraparound mesh has five wraparound cycles $(\widetilde{\mathbf{x}}, \widetilde{\mathbf{x}} \oplus (0,1), \widetilde{\mathbf{x}} \oplus 2\cdot(0,1), \widetilde{\mathbf{x}} \oplus 3\cdot(0,1), \widetilde{\mathbf{x}})$ of length 4, where $\oplus$ stands for addition of vectors such that first components are added modulo 5 and the second components are added modulo 4, four wraparound cycles $(\widetilde{\mathbf{x}}, \widetilde{\mathbf{x}} \oplus (1,0), \widetilde{\mathbf{x}} \oplus 2\cdot(1,0), \widetilde{\mathbf{x}} \oplus 3\cdot(1,0), \widetilde{\mathbf{x}} \oplus 4\cdot(1,0), \widetilde{\mathbf{x}})$ of length 5, and one wraparound cycle $(\widetilde{\mathbf{x}}, \widetilde{\mathbf{x}} \oplus (-1,1), \widetilde{\mathbf{x}} \oplus 2\cdot(-1,1), \widetilde{\mathbf{x}} \oplus 3\cdot(-1,1), \ldots, \widetilde{\mathbf{x}} \oplus 19\cdot(-1,1), \widetilde{\mathbf{x}})$ of length 20. In the figure, a path from node $\widetilde{\mathbf{x}} = (3,2)$ to node $\widetilde{\mathbf{y}} = (0,3)$, $P = ((3,2),(2,2),(1,2),(0,2),(0,3))$ is shown using thick lines.

To construct sets of prohibited turns for $M_D(p_1,p_2,\ldots,p_n)$ or $M_D^W(p_1,p_2,\ldots,p_n)$ we will introduce a total ordering of nodes in these meshes.

Definition 4: If $\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}} \in V$ where V is the set of nodes in $M_D(p_1,p_2,\ldots,p_n)$ or $M_D^W(p_1,p_2,\ldots,p_n)$, we will say that $\widetilde{\mathbf{x}} > \widetilde{\mathbf{y}}$ if $\widetilde{\mathbf{x}}^{(i)} > \widetilde{\mathbf{y}}^{(i)}$ where i is the smallest integer such that $\widetilde{\mathbf{x}}^{(i)} \neq \widetilde{\mathbf{y}}^{(i)}$ ($\widetilde{\mathbf{x}} = \mathbf{C}\,\mathbf{x}$, $\widetilde{\mathbf{y}} = \mathbf{C}\mathbf{y}$).

Theorem 2: For a finite mesh $M_D(p_1,p_2,\ldots,p_n)$ or a wraparound mesh $M_D^W(p_1,p_2,\ldots,p_n)$, let the set of prohibited turns $F = \{(\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}}, \widetilde{\mathbf{z}}) \mid \widetilde{\mathbf{x}}, \widetilde{\mathbf{y}}, \widetilde{\mathbf{z}} \in V$ and $\widetilde{\mathbf{y}} > \widetilde{\mathbf{x}}$, $\widetilde{\mathbf{y}} > \widetilde{\mathbf{z}}\}$. Then

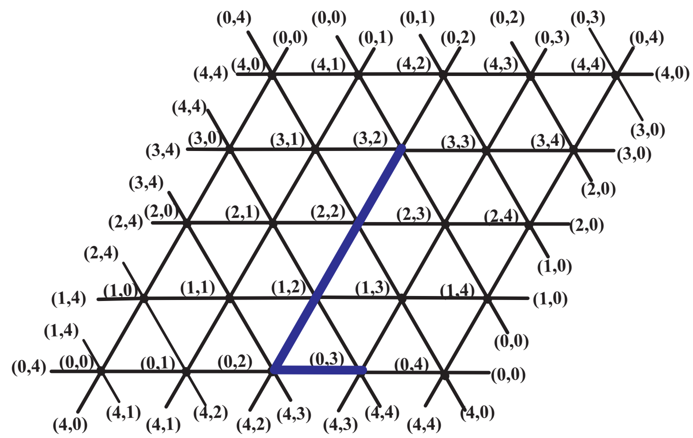

Fig. 6. A wraparound D-Mesh $M_D^W(5,5)$ with $D = \{\pm\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm(1,0)\}$, $D_{\mathbf{C}} = \{\pm(1,0), \pm(0,1), \pm(1,-1)\}$

1) For any $\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}} \in V$ there exists a path from $\widetilde{\mathbf{x}}$ to $\widetilde{\mathbf{y}}$ containing no turns from F.
2) For any cycle there exists a turn in the cycle that belongs to F.
3) The set F is asymptotically optimal if $p_i \to \infty$ ($i = 1,\ldots,n$), and the minimum fraction z of prohibited turns for $M_D(p_1,p_2,\ldots,p_n)$ or $M_D^W(p_1,p_2,\ldots,p_n)$ with $|D| = d$ is, asymptotically,

$$\lim_{\substack{p_i\to\infty\\ i=1,\ldots,n}} z = \frac{1}{4}\left(1 - \frac{1}{d-1}\right).$$

Proof:

1) First we will prove that if $\widetilde{\mathbf{x}} = \mathbf{C}\,\mathbf{x} = \left(\widetilde{\mathbf{x}}^{(1)}, \widetilde{\mathbf{x}}^{(2)}, \ldots, \widetilde{\mathbf{x}}^{(n)}\right)$ and $\widetilde{\mathbf{y}} = \mathbf{C}\mathbf{y} = \left(\widetilde{\mathbf{y}}^{(1)}, \widetilde{\mathbf{y}}^{(2)}, \ldots, \widetilde{\mathbf{y}}^{(n)}\right)$, there exists a path from $\widetilde{\mathbf{x}}$ to $\widetilde{\mathbf{y}}$ in $M_D(p_1,p_2,\ldots,p_n)$ or in $M_D^W(p_1,p_2,\ldots,p_n)$ containing no turns from F. Let $S_+(\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}}) = \{i \mid \widetilde{\mathbf{x}} \geq \widetilde{\mathbf{y}}\}$ and $S_-(\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}}) = \{i \mid \widetilde{\mathbf{x}} < \widetilde{\mathbf{y}}\}$. Consider now a node $\widetilde{\mathbf{z}}$ such that $\widetilde{\mathbf{z}}^{(i)} = \min\left(\widetilde{\mathbf{x}}^{(i)}, \widetilde{\mathbf{y}}^{(i)}\right)$. Obviously, there exists a path from $\widetilde{\mathbf{x}}$ to $\widetilde{\mathbf{z}}$, such that any next node in the path is smaller than the previous one. Similarly, there exists a path from $\widetilde{\mathbf{z}}$ to $\widetilde{\mathbf{y}}$ such that any next node is larger than the previous one. Now take the concatenation of these two paths. The turn at node $\widetilde{\mathbf{z}}$ is permitted, since $\widetilde{\mathbf{z}}$ is smaller than the two neighboring nodes in the path. Thus, there exists a permitted path from $\widetilde{\mathbf{x}}$ to $\widetilde{\mathbf{y}}$.
2) In every cycle $(\widetilde{\mathbf{x}}_{\mathbf{1}}, \widetilde{\mathbf{x}}_{\mathbf{2}}, \ldots, \widetilde{\mathbf{x}}_{(\ell-\mathbf{1})}, \widetilde{\mathbf{x}}_\ell)$ where $\widetilde{\mathbf{x}}_{(\ell-\mathbf{1})} = \widetilde{\mathbf{x}}_{\mathbf{1}}$ and $\widetilde{\mathbf{x}}_\ell = \widetilde{\mathbf{x}}_{\mathbf{2}}$ there exists $i \in \{1,2,\ldots,\ell\}$ such that $\widetilde{\mathbf{x}}_{\mathbf{i}} > \widetilde{\mathbf{x}}_{(\mathbf{i-1})}$, $\widetilde{\mathbf{x}}_{\mathbf{i}} > \widetilde{\mathbf{x}}_{(\mathbf{i+1})}$, and turn $(\widetilde{\mathbf{x}}_{(\mathbf{i-1})}, \widetilde{\mathbf{x}}_{\mathbf{i}}, \widetilde{\mathbf{x}}_{(\mathbf{i+1})}) \in F$.
3) We will say that the node $\widetilde{\mathbf{x}} \in V$ is internal in $M_D(p_1,p_2,\ldots,p_n)$ or in $M_D^W(p_1,p_2,\ldots,p_n)$ if $0 < \widetilde{\mathbf{x}}^{(i)} < p_i - 1$ for all $i = 1,\ldots,n$. If $\mathbf{x}$ is an internal node, then in each pair of its neighbors, $\mathbf{x} \pm \mathbf{a}_i$ ($i = 1,\ldots,t$) one neighbor is larger than $\mathbf{x}$ and the other is smaller than $\mathbf{x}$. Thus for any internal node $\mathbf{x}$ exactly t neighbors are larger than $\mathbf{x}$, and exactly t neighbors are smaller

than $\mathbf{x}$. Hence, for every internal node $\mathbf{x}$ there are $\binom{t}{2}$ turns $(\mathbf{y}, \mathbf{x}, \mathbf{z})$ which belong to F. Thus,

$$\lim_{\substack{p_i \to \infty \\ i=1,\ldots,n}} z \leq \frac{\binom{t}{2}}{\binom{2t}{2}} = \frac{1}{4}\left(1 - \frac{1}{d-1}\right).$$

On the other hand, similar to the proof of Theorem 1, for any internal node $\mathbf{x}$ there are $\binom{t}{2}$ cycles in $M_D(p_1, p_2, \ldots, p_n)$ or in $M_D^W(p_1, p_2, \ldots, p_n)$ which contain 4 nodes each that do not have common turns. In the union of these sets for all internal nodes any two cycles do not have common turns. Since ate most $\binom{d}{2}$ turns are prohibited at any non-internal node, the contribution of the non-internal nodes to Z does not exceed their fraction among all nodes, and, therefore, is infinitesimal when $p_i \to \infty \quad (i = 1, \ldots, n)$. Thus, it follows that

$$\lim_{\substack{p_i \to \infty \\ i=1,\ldots,n}} z \geq \frac{1}{4}\left(1 - \frac{1}{d-1}\right).$$

■

The set $F = W(M_D(5,3))$ of prohibited turns for the $M_D(5,3)$ with $D = \left\{ \pm\left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right), \pm(1,0) \right\}$ is shown in Fig. 5.

IV. Special Topologies

A. Finite Meshes and Tori

Meshes and tori have been the most widely used communication network topologies for multiprocessors [5], [49]. Recently, ”TOFU”, a 6-dimensional mesh and torus topologies have been used to provide the extremely high performance and fault tolerant interconnection network, achieving 10 petaflops [50]. In this section, we first consider square meshes, with each inner node connected with 2n nodes, where n is the dimension of a mesh. Meshes of this type were investigated in [18], where only 90-degree turns were taken into account. It was shown, that $1/4$ of all such turns has to be prohibited. With a more general turn model, our results are in agreement with authors' conclusion in [18].

Theorem 3: For n-dimensional p-ary mesh, M_p^n

$$z(M_p^n) = \frac{(n-1)(p-1)^2}{2p(p-2) + 4(n-1)(p-1)^2}, \tag{6}$$

and for n-dimensional p-ary tori, T_p^n, with $p > 2$,

$$\frac{(n-1)p + 2}{2(2n-1)p} < z(T_p^n) < \frac{(n-1)(p^2+2) + 2p}{2(2n-1)p^2}. \tag{7}$$

Proof: To prove the lower bound for meshes we consider the system of all cycles of length 4. There are $R = \binom{n}{2}(p-1)^2 p^{n-2}$ turn-disjoint cycles of this type and the total number of turns in M_p^n is equal to

$$T(M_p^n) = n(p-2)p^{n-1} + 4\tbinom{n}{2}(p-1)^2 p^{n-2}. \tag{8}$$

The lower bound for $Z(M_p^n)$ follows now by observing that at least as many turns must be prohibited as there are turn-disjoint cycles.

To prove the upper bound of Theorem 3 for p-ary meshes, we prohibit all turns $(\mathbf{a}, \mathbf{b}, \mathbf{c})$, where $l(\mathbf{a}) < l(\mathbf{b})$, $l(\mathbf{b}) > l(\mathbf{c})$, and $l(\mathbf{a}), l(\mathbf{b}), l(\mathbf{c})$ are distances in terms of number of hops from node $(0, 0, \ldots, 0)$ to $\mathbf{a}, \mathbf{b}$ and $\mathbf{c}$. The number of prohibited turns is equal to

$$Z(M_p^n) = \binom{n}{2}(p-1)^2 p^{n-2}. \tag{9}$$

Then (6) follows from (8) and (9).

The lower bound for tori is obtained by counting all cycles with disjoint sets of turns, namely, all cycles of length 4 and all np^{n-1} one-dimensional cycles.

To obtain the upper bound, consider in any of n dimensions one direction as positive and the other one as the negative. Then at each of p^n nodes prohibit all $\binom{n}{2}$ turns from positive direction to negative. Also, to break all one-dimensional cycles, prohibit turns along each of n coordinates of the form $\{(x_1, \ldots, x_{i-1}, p-2, x_{i+1}, \ldots, x_n), (x_1, \ldots, x_{i-1}, p-1, x_{i+1}, \ldots, x_n), (x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n)\}$ and the values $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n$ are all possible. To break all more complex cycles, it is sufficient to prohibit all turns from the positive direction along one of the coordinates to the positive direction along another coordinate at the point where both coordinates have values equal to $p-1$. There are $2\binom{n}{2}p^{n-2}$ such turns.

■

B. Hexagonal and Honeycomb Meshes

Next, we consider *hexagonal* meshes [40], [41], [49] in which each node has up to 6 neighbors and *honeycomb* meshes [42], [43], [49] where each node has up to 3 neighbors, and their corresponding tori. In a hexagonal mesh of size p denoted by HeM_p, peripheral edges form a regular hexagon where each side has p nodes. A honeycomb mesh of size p, denoted by HoM_p, where each side of the mesh has p hexagonal cells whose centers also form a regular hexagon. The hexagonal and honeycomb tori are degree six and degree three regular topologies, respectively.

In a hexagonal mesh HeM_p, there are $N = 3p^2 - 3p + 1$ nodes with labels $0, 1, \ldots, (N-1)$ with the center node having the label 0 [44]. Adjacent nodes of any given node a are identified to have labels $a \pm 1$, $a \pm (3p-1)$, $a \pm (3p-2)$ where arithmetic operations are $\mod N$. In the corresponding torus, wraparound edges are also identified using the same adjacency rules. Labels of adjacent nodes are shown in Fig. 7(a) for the case of a size $p = 3$ torus.

In a honeycomb torus, nodes that are connected by the wraparound edges are those nodes that are mirror symmetric with respect to the three lines passing through the center and normal to each of three edge orientations [43]. These axes are shown as dashed lines in Fig. 7(b).

Theorem 4: For a hexagonal mesh of size p, HeM_p, with $N = 3p^2 - 3p + 1$ nodes,

$$z(HeM_p) = \frac{9p^2 - 21p + 13}{45p^2 - 99p + 51}, \tag{10}$$

and for a hexagonal torus HeT_p of size p,

$$z(HeT_p) = \frac{9p^2 - 15p + 10}{45p^2 - 45p + 15}. \tag{11}$$

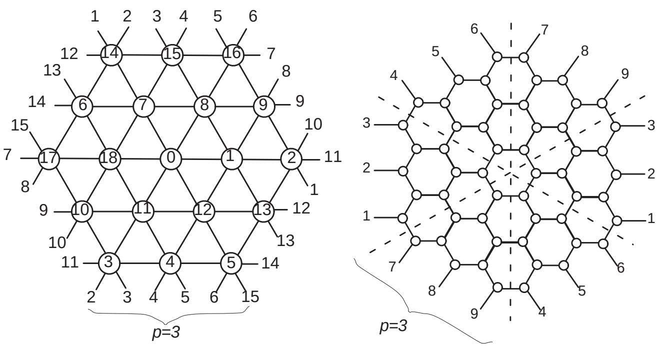

(a) (b)

Fig. 7. Examples hexagonal torus HeT_3 in (a), and honeycomb torus HoT_3 in (b) for $p = 3$, where wraparound links are identified.

Proof: First, note that total number of turns in a HeM_p is equal to: $T(HeM_p) = 15(3p^2 - 9p + 7) + 6(6p - 12) + 18 = 45p^2 - 99p + 51$.

To prove the lower bound, we consider the set of all turn-disjoint $6(p-1)$ triangles, and $3p^2 - 9p + 7$ hexagons and observe that we must prohibit at least as many turns as there are turn-disjoint cycles, e.g., triangles and hexagons.

Upper bound on $Z(HeM_p)$ can be obtained as shown in Fig. 8(a).

For the case of hexagonal tori with $N(HeT_p) = 3p^2 - 3p + 1$ nodes, $m(HeT_p) = 3N(HeT_p)$ edges, and $T(HeT_p) = 15N(HeT_p)$ turns, additional $6(2p-1)$ turns have to be prohibited to prevent all wraparound cycles. Therefore, $6(2p-1)$ cycles must be added to the system of turn-disjoint cycles due to triangles and hexagons. Again, observe that we must prohibit at least as many turns as there are turn-disjoint cycles. To prove the upper bound, we cut the wraparound cycles in the hexagonal torus and prohibit all $6(2p-1)$ turns at the nodes on the border of the resulting mesh. ■

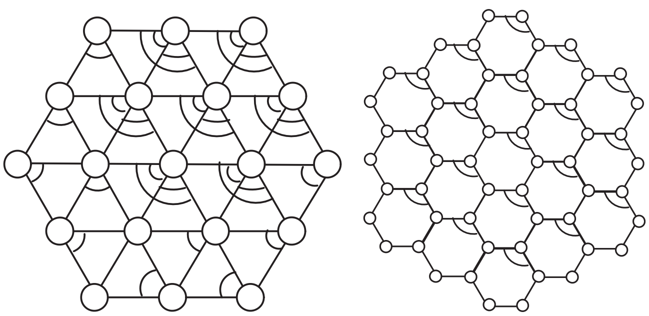

(a) (b)

Fig. 8. Prohibited turns for hexagonal (a) and honeycomb (b) meshes showing the prohibited turns.

C. Locally Complete Tree-Like Topologies

Locally complete tree-like topologies are hybrid topologies incorporating the properties and attributes of its components [49]. Consider a tree $T' = G'(V', E')$ with $M' = |E'|$ undirected edges $\{v_i, v_j\} \in E'$ and $N' = |V'|$ nodes v_i, $i = 0, \ldots, N'-1$. Assume that each node of the tree is now replaced with a complete graph K_n with $n \geq d_i$ nodes where d_i is the degree of node v_i of the original tree T', to obtain the augmented graph $G(V, E)$ which is **locally complete**. The locally complete graph has $N = |V| = N'n$ nodes and $|E| = N' - 1 + N'\binom{n}{2}$ edges. Let us denote the nodes of K_n that replaces node v_i of the original tree by $v_{i,m}$ $(m = 0, 1, \ldots, n-1)$. Embedding of the complete graph K_n is done in such a way that if the v_i is the parent of nodes v_j and v_k, then in the locally complete graph $v_{j,0}$ is connected to node $v_{i,r}$ and $v_{k,0}$ is connected to node $v_{i,s}$, where $r \neq s$ and $r, s \neq 0$ (Fig. 9).

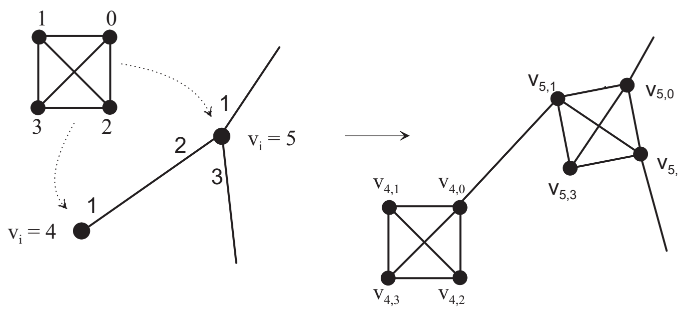

Fig. 9. Embedding a complete graph K_4 at tree nodes $v_i = 4$ and $v_i = 5$. Port numbers at nodes $v_i = 4$, $v_i = 5$, and the node numbers of the complete graph K_4 are displayed.

Theorem 5: For a locally complete tree-like graph obtained as described above, the fraction of prohibited turns is given by

$$z(G) = \frac{1}{3}\frac{N'n(n-2)}{N'n(n-2) + 4(N'-1)}. \tag{12}$$

Proof: Since the minimum degree nodes will always be at the leaf node positions of the original tree, the number of prohibited turns in each embedded K_n is given by $Z(K_n) = \binom{n-1}{2} + Z(K_{n-1})$. Solving this recursion equation, we obtain $Z(K_n) = \binom{n}{3}$. Hence, for the augmented graph G with N' nodes we have

$$Z(G) = N'Z(K_n) = N'\binom{n}{3}.$$

In embedding a K_n at a tree node of degree d_i, only d_i nodes of the K_n will be connected directly to the original tree. This means that embedding a K_n graph at an original tree node, will create nodes of at most degree n in the locally complete graph. Also, note that when a K_n is embedded at a tree node with degree d_i, there will be $n\binom{n-1}{2}$ turns contributed by the K_n and $(n-1)d_i$ turns contributed by the d_i edges of the original tree. With these observations the total number of turns is $T(G) = N'n\binom{n-1}{2} + \sum_{i=1}^{N'}(n-1)d_i = N'n\binom{n-1}{2} + (n-1)\sum_{i=1}^{N'} d_i$ or

$$T(G) = \frac{1}{2}N'n(n-1)(n-2) + 2(n-1)(N'-1).$$

Hence,

$$z(G) = \frac{1}{3}\frac{N'n(n-2)}{N'n(n-2) + 4(N'-1)}.$$

■

For example, for $n = 3$ and $N' \to \infty$, $z(G) = \frac{1}{3}\frac{3}{3+4} = \frac{1}{7}$, and for $n = 4$, and $N' \to \infty$, $z(G) = \frac{1}{3}\frac{8}{12} = \frac{2}{9}$.

D. Cube Connected Cycles

We will consider now a binary n-cube connected cycles, CCC [49], where each node of an n-dimensional binary cube is replaced by a cycle of n nodes of degree 3 (see Fig. 10 for $n = 3$). These interconnection networks are popular, since they combine the properties of small node degree and small diameter of the network graph [51]. First, we will establish upper and lower bounds with Theorem 6 for a slightly larger class of graphs.

Theorem 6: If graph G is obtained from d-regular graph H ($d_i = d$ for all i, $d > 2$) with $N(H)$ nodes by replacing each node by the cycle of d nodes, then

$$\frac{1}{6} + \frac{2}{3dN(H)} \leq z(G) \leq \frac{1}{6} + \frac{1}{3d}. \qquad (13)$$

Proof: The lower bound can be obtained from $Z(G) \geq M - N + 1$, since for G there are $M(G) = 1.5N(H)d$ edges and $T(G) = 3N(H)d$ turns.

To prove the upper bound, we label all nodes in G as (i, j), where i is the number of the cycle containing the node i in G, and j is the number of a node within each cycle of length d, $i \in \{1, \ldots, N(H)\}$, $j \in \{0, 1, \ldots, (d-1)\}$, as shown in Fig. 10. In each cycle, nodes are labeled subsequently. In cycle i we prohibit the turn $((i, d-1), (i, 1), (i, 2))$. There exist $N(H)$ such turns. Also, for each of $N(H)d/2$ edges between different cycles (edges between cycles in G correspond to edges in H), we prohibit turn (a, b, c), where $a = (i_1, j_1)$, $b = (i_2, j_2)$, $c = (i_2, j_3)$, if $i_1 < i_2$ and $j_3 = j_2 + 1 \mod d$. Then it follows that $z(G) \leq \frac{|W|}{T(G)} = \frac{1}{6} + \frac{1}{3d}$. ■

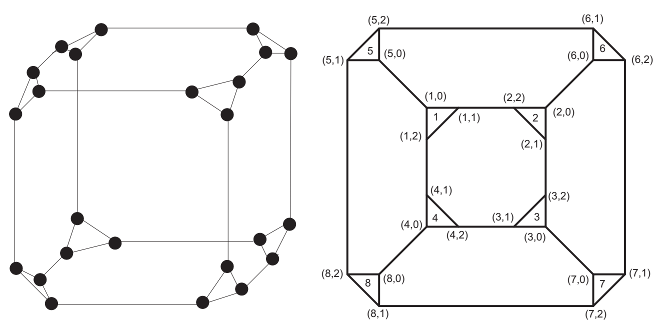

Fig. 10. Binary 3-cube connected cycles and their labels

For example, for the binary n-dimensional cube connected cycle (CCC with $d = n$, $N = 2^n$) if $n \to \infty$, then $z(G) \to \frac{1}{6}$. We note that for binary cube Z_2^n we have $z(Z_2^n) = \frac{1}{4}$.

The following theorem is a generalization of Theorem 6.

Theorem 7: If all nodes of 3-regular graph G with N nodes can be covered by k non-intersecting **simple** cycles, then

$$\frac{1}{6} + \frac{2}{3N} \leq z(G) \leq \frac{1}{6} + \frac{k}{3N}. \qquad (14)$$

Proof: The proof of Theorem 7 is similar to the proof of Theorem 6. ■

To illustrate Theorem 7, let us consider the 4-pancake graph [49]. E.g., in a 4-pancake graph, nodes have labels that include all 4! = 24 orderings of numbers 1, 2, 3, and 4. For the q-pancake, node $(1, 2, \cdots, i-1, i, i+1, \cdots, q)$ is connected to nodes $(i, i-1, \cdots, 2, 1, i+1, \cdots, q)$ for each i, i.e., $1, 2, \cdots, i$ is flipped, like a pancake [49]. In a 4-pancake, nodes that are adjacent to node $(1, 2, 3, 4)$ are $(2, 1, 3, 4)$, $(3, 2, 1, 4)$ and $(4, 3, 2, 1)$. (see Fig. 11(a)). For this graph, $N = 24, k = 4$ and according to Theorem 7 and (3), $z(G) = 2/9$.

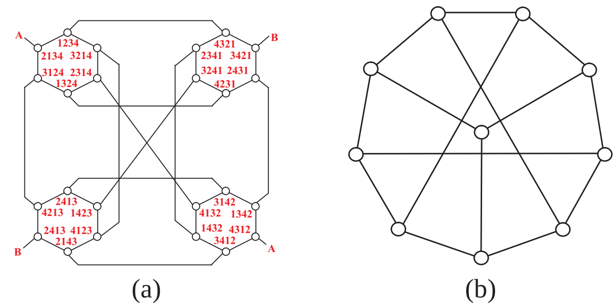

Fig. 11. 4-Pancake graph with node labels (a), and Petersen graph (b).

Another graph, which can be analyzed by Theorem 7, is the Petersen graph [51], (Fig. 11(b)) which has the smallest diameter (equal to 2) among all regular graphs of degree 3. For this graph $N = 10$, $k = 2$ and by Theorem 7 and (3) we obtain $z(G) = 7/30$.

V. Distance Dilation

Consider now the notion of dilation in a network topology due to turn prohibitions. Paths that involve prohibited turns cannot be used for communication. Thus, one side effect of turn prohibitions is that prohibiting certain paths from being used for message routing may increase distances between some nodes. The net result of this is that the average distance of the network graph will be increased. To facilitate the investigation of this phenomenon, the notion of distance dilation is introduced.

Definition 5: The dilation in a graph, is the ratio of the average distance after turn prohibition to the average distance without any turn prohibition.

When the dilation is 1 it implies that the turn prohibitions have not caused any lengthening of the average distance. For example, for complete graphs the fraction of prohibited turns achieves the upper bound, but the dilation is 1. Similarly for homogeneous and D-meshes, for hexagonal meshes, p-ary n-dimensional meshes and, for locally complete tree-like topologies no dilation is introduced by turn prohibitions. In Fig. 12 the distance dilations in p-ary n-dimensional tori are shown. For these calculations, we used the formulation in Section IV to identify the turns to prohibit and then determined the average distance using the shortest distances between all source-destination pairs that do not include any of the prohibited turns. Note that popular West-First, North-Last, and DOR algorithms are not suitable for tori (since they are not deadlock free) and therefore are not included in our calculations. In contrast with our algorithm, those routing methods are not applicable to more general meshes (such as hexagonal and honeycomb meshes) and other regular topologies. We make following observations on dilation using our approach.

For the $p = 3$ and $p = 4$, no dilation is experienced by the tori. For $p = 5$ and $p = 6$ the largest dilation occur for the one-dimensional cases. As expected, the dilation is larger for larger p for the one-dimensional cases with a maximum of 7.5% and diminishing for larger dimensions.

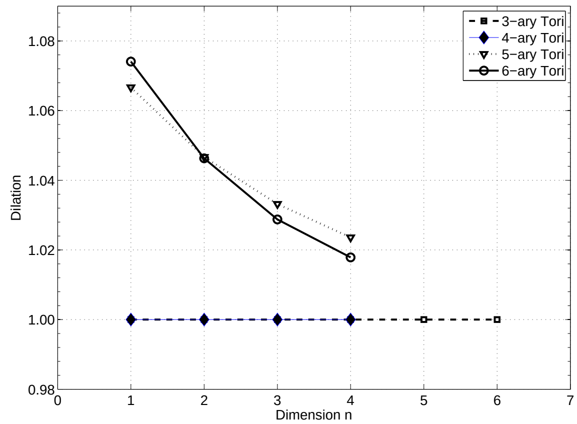

Fig. 12. Dilation in p-ary n-dimensional tori due to turn prohibition, for $p = 3, \ldots, 6$.

VI. Conclusions

In this paper we considered the problem of constructing minimum cycle-breaking sets of turns for graphs that model communication networks. This problem is important for deadlock-free and livelock-free message routing in computer communication networks. In contrast to popular DOR algorithm, the proposed turn prohibition techniques are compatible with any adaptive routing approach. We present a series of new algorithms that are used to obtain optimal or close to optimal sets of prohibited turns to prevent deadlock formation during routing. The results on the fraction of prohibited turns are summarized in Table I. We present the results of our calculations for dilations as a result of prohibitions in p-ary n-dimensional tori. We show that meshes do not suffer from any dilation and the worst case dilation for tori is less than 7.5%.

TABLE I
Lower and upper bounds on fractions of prohibited turns, $z(G)$, in minimal cycle breaking sets for several regular and semiregular topologies

Topology	Lower Bound) on $z(G)$	Upper Bound on $z(G)$
Homogenous meshes Theorem 1	$\frac{1}{4}\left(1 - \frac{1}{d-1}\right)$	
Complete graph K_n, $n > 2$, [35], [36]	$1/3$	
n-dimensional p-ary mesh M_p^n, (6)	$\frac{(n-1)(p-1)^2}{2p(p-2)+4(n-1)(p-1)^2}$	
n-dimensional p-ary torus Z_p^n $(p > 2)$, (7)	$\frac{(n-1)p+2}{2(2n-1)p}$	$\frac{(n-1)p^2+2(n+p-1)}{2(2n-1)p^2}$
Hypercube Z_2^n , (IV-D)	$1/4$	
Hexagonal mesh, size p (10)	$\frac{9p^2-21p+13}{45p^2-99p+51}$	
Hexagonal torus, size p (11)	$\frac{9p^2-15p+10}{45p^2-45p+15}$	
Honeycomb mesh, size p (IV-B)	$\frac{3p^2-3p+1}{18p^2-12p}$	
Honeycomb torus, size p (IV-B)	$\frac{1}{6} + \frac{1}{18p^2}$	
Cube-connected cycle $N = 2^n$ *nodes* (13)	$\frac{1}{6} + \frac{1}{3N}$	$\frac{1}{6} + \frac{1}{3d}$
4-pancake graph (IV-D)	$2/9$	
Petersen graph (IV-D)	$7/30$	
Locally Complete graphs Each one of the N' nodes of a tree is replaced by K_n, (12)	$\frac{1}{3}\frac{N'n(n-2)}{N'n(n-2)+4(N'-1)}$	

References

[1] W. J. Dally and C. L. Seitz, “Deadlock-free message routing in multiprocessor interconnection networks,” *IEEE Transactions on Computers*, vol. 36, pp. 547–553, 1987.

[2] R. Boppana and S. Chalasani, “A comparison of adaptive wormhole routing algorithms,” *Computer Architecture News*, vol. 21, no. 2, pp. 351–360, 1993.

[3] S. Chalasani and R. V. Boppana, “Fault-tolerant wormhole routing algorithms in mesh networks,” *IEEE Transactions on Computers*, vol. 44, no. 7, pp. 848–864, 1995.

[4] J. Duato, S. Yalamancili, and L. Ni, *Interconnection Newtworks and Engineering Approach*. IEEE Computer Society Press, 1997.

[5] L. M. Ni and P. K. McKinley, “A survey of wormhole routing techniques in direct networks,” *Computer*, vol. 26, pp. 62–76, 1993.

[6] A. V. de Mello, L. C. Ost, F. G. Moraes, and N. L. V. Calazans, “Evaluation of routing algorithms on mesh based NoCs,” FACULDADE DE INFORMATICA PUCRS-Brazil, Tech. Rep. 040, May 2004.

[7] J. Hu and R. Marculescu, “DyAD - smart routing for networks-on-chip,” *Proceedings of the 41st annual conference on Design Automation*, vol. 00, pp. 260–263, 2004.

[8] P. Gaughan and S. Yalamanchili, “A family of fault tolerant routing protocols for direct multiprocessor networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 6, no. 5, pp. 482–497, 1995.

[9] D. Klepacki, “Blue gene,” http://lobster.bu.edu/SCV/Archive/IBM/BGL-BU.pdf, 2003.

[10] R. Libeskind-Hadas, D. Mazzoni, and R. Rajagopalan, “Tree-base multicast routing in the mesh with no virtual channels,” *Proceedings of the First Merged International Parallel Processing Symposium and Symposium on Parallel and Distributed Processing*, pp. 244–249, 1998.

[11] F. Silla, J. Duato, A. Sivasubramaniam, and C. R. Das, “Virtual channel multiplexing in networks of workstations with irregular topology,” *Proceedings of the International Conference on High Performance Computing*, pp. 147–154, 1998.

[12] F. Silla and J. Duato, “High-performance routing in networks of workstations with irregular topology,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 11, no. 7, pp. 699–719, 2000.

[13] W. J. Dally and C. L. Seitz, “The torus routing chip,” *Journal of Distributed Computing*, vol. 1, no. 3, pp. 187–196, 1986.

[14] C. Glass and L. Ni, “The turn model for adaptive routing,” *Proceedings*

of the 19th Annual International Symposium on Computer Architecture, pp. 278–286, 1992.

[15] J. Duato, “A new theory of deadlock-free adaptive routing in wormhole networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 4, pp. 1320–1331, 1993.

[16] ——, “A necessary and sufficient condition for deadlock-free adaptive routing in wormhole networks,” *Proceedings of International Conference on Parallel Processing*, pp. 142–149, 1994.

[17] W. Dally and H. Aoki, “Deadlock-free adaptive routing in multiprocessor networks using virtual channels,” *IEEE Transactions on Parallel and Distributed Systems*, pp. 466–475, 1997.

[18] C. Glass and L. Ni, “The turn model for adaptive routing,” *Journal of ACM*, vol. 5, pp. 874–902, 1994.

[19] L. Zakrevski, S. Jaiswal, L. Levitin, and M. Karpovsky, “A new method for deadlock elimination in computer networks with irregular toplologies,” *Proceedings of the International Association of Science and Technology for Development (IASTED) Conference, Parallel and Distributed Computing Systems (PDCS-99)*, vol. 1, pp. 396–402, 1999.

[20] L. Zakrevski, S. Jaiswal, and M. Karpovsky, “Unicast message routing in communication networks with irregular topologies,” *Proceeding of Computer Aided Design (CAD-99)*, 1999.

[21] L. Zakrevski, M. Mustafa, and M. Karpovsky, “Turn prohibition based routing in irregular computer networks,” *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pp. 175–179, 2000.

[22] O. Lysne, T. Skeie, S.-A. Reinemo, and I. Theiss, “Layered routing in irregular networks,” *IEEE Transactions on Parallel and Distributed Systems*, pp. 51–56, 2006.

[23] J. C. Sancho, A. Robles, P. L?pez, J. Flich, and J. Duato, “Routing in infiniband¡sup¿tm¡/sup¿ torus network topologies,” *Parallel Processing, International Conference on*, vol. 0, p. 509, 2003.

[24] M. Mustafa, M. Karpovsky, and L. Levitin, “Cycle breaking in wormhole routed computer communication networks,” in *Opnetwork 2005*, Opnet Technologies. Opnet Technologies, August 2005.

[25] L. B. Levitin, M. G. Karpovsky, M. Mustafa, and L. Zakrevski, “A new algorithm for finding cycle-breaking sets of turns in a graph,” *JGAA*, vol. 10, no. 2, pp. 387–420, 2006.

[26] M. D. Schroeder, A. D. Birrell, M. Burrows, H. Murray, R. M. Needham, T. L. Rodeheffer, E. H. Satterthwaite, and C. P. Thacker, “Autonet: A high-speed self configuring local area network using point-to-point links,” Digital Equipment Corporation, SRC, Tech. Rep. SRC Research Report 59, April 1990.

[27] J. Sancho and A. Robles, “Improving the Up*/Down* routing scheme for networks of workstations,” *European Conference on Parallel Computing (Euro-Par 2000)*, 2000.

[28] D. Starobinski, M. Karpovsky, and L. Zakrevski, “Application of network calculus to general topologies using turn prohibition,” *IEEE/ACM Transactions on Networking*, vol. 11, no. 3, pp. 411–421, 2003.

[29] J. Duato, “Deadlock-free adaptive routing algorithms for multicomputers: Evaluation of a new algorithm,” in *Proc. Third IEEE Intl Symp. Parallel Distributed Processing.*, Dec. 1991.

[30] G. D. Pifarré, L. Gravano, G. Denicolay, and J. L. C. Sanz, “Adaptive deadlock- and livelock-free routing in the hypercube network,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 5, no. 11, pp. 1121–1139, 1994.

[31] G. Min, M. Ould-Khaoua, D. Kouvatsos, and I. Awan, “A queuing model of dimension-ordered routing under self-similar traffic loads,” in *Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International*, 26-30 2004, p. 251.

[32] F. D. Pellegrini, D. Starobinski, M. Karpovsky, and L. Levitin, “Scalable cycle-breaking algorithms for gigabit ethernet backbones,” *Proceedings of IEEE Infocom 2004*, 2004.

[33] F. D. Pellegrini, D. Starobinski, M. G. Karpovsky, and L. B. Levitin, “Scalable, distributed cycle-breaking algorithms for gigabit ethernet backbones,” *Journal of Optical Networks*, vol. 5, no. 2, pp. 122–144, February 2006.

[34] A. Shevtekar and L. Zakrevski, “Hybrid turn-prohibition routing algorithm for the networks of workstations,” *Parallel and Distributed Processing Techniques and Applications*, 2004.

[35] L. Levitin, M. Karpovsky, and M. Mustafa, “Deadlock prevention by turn prohibitions in interconnection networks,” in *2009 IEEE International Symposium on Parallel & Distributed Processing*. Rome, Italy: IEEE Computer Society, May, 2009, pp. 1–7.

[36] ——, “Minimal sets of turns for breaking cycles in graphs modeling networks,” *IEEE Transactions in Parallel and Distributed Systems*, vol. 21, no. 9, pp. 1342–1353, September 2010.

[37] T. Skeie, O. Lysne, and I. Theiss, “Layered shortest path (LASH) routing in irregular system area networks,” in *International Parallel and Distributed Processing Symposium: IPDPS 2002 Workshops*, 2002, pp. 0162–0170.

[38] J. C. Sancho, A. Robles, and J. Duato, “An effective methodology to improve the performance of the Up*/Down* routing algorithm,” *IEEE Transactions of Parallel and Distributed Systems*, vol. 15, no. 8, pp. 740–754, August 2004.

[39] R. Horst, “ServerNet(TM) deadlock avoidance and fractahedral topologies,” *Proceedings of IEEE International Parallel Processing Symposium*, pp. 274–280, 1996.

[40] C. Decayeux and D. Seme, “3d hexagonal network: Modeling, topological properties, addressing scheme, and optimal routing algorithm,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 16, pp. 875–884, 9 2005.

[41] F. G. Nocetti, I. Stojmenovic, and J. Zhang, “Addressing and routing in hexagonal networks with applications for tracking mobile users and connection rerouting in cellular networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 13, no. 9, pp. 963–971, 2002.

[42] B. Parhami and D.-M. Kwai, “A unified formulation of honeycomb and diamond networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 12, no. 1, pp. 74–80, 2001.

[43] I. Stojmenovic, “Honeycomb networks: Topological properties and communication algorithms,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 8, no. 10, pp. 1036–1042, 1997.

[44] J. W. Dolter, P. Ramanathan, and K. G. Shin, “Performance analysis of virtual cut-through switching in HARTS: A hexagonal mesh multicomputer,” *IEEE Transactions on Computers*, vol. 40, no. 6, pp. 669–680, 1991.

[45] A. Hansson, K. Goossens, and A. Radulescu, “Avoiding message-dependent deadlock in network-based systems on chip,” *VLSI Design, vol. 2007,Article ID 95850, 10 pages, doi:10.1155/2007/95859*, 1998.

[46] S. Chatterjee, M. Kishinevsky, and m. Y. Ogras, “xmas: Quick formal modeling of communication fabrics to enable verification.” *IEEE Design & Test of Computers*, vol. 29, no. 3, pp. 80–88, 2012.

[47] A. Gotmanov, S. Chatterjee, and M. Kishinevsky, “Verifying deadlock-freedom of communication fabrics,” in *Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation*, ser. VMCAI’11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 214–231. [Online]. Available: http://dl.acm.org/citation.cfm?id=1946284.1946300

[48] I. Honkala, M. G. Karpovsky, and L. B. Levitin, “On robust and dynamic identifying codes,” *IEEE Trans. Info. Theory*, pp. 599–613, Feb 2006.

[49] B. Parhami, *Introduction to Parallel Processing: Algorithms and Architectures*. Plenum Press, 1998.

[50] Y. Ajima, S. Sumimoto, and T. Shimizu, “Tofu: A 6d Mesh/Torus interconnect for exascale computers,” *Computer*, vol. 42, no. 11, pp. 36–40, November 2009.

[51] F. Harary, *Graph Theory*, ser. The Advanced Book Program. Perseus Books, 1998.