Edge-Disjoint Spanners in tori*

Arthur L. Liestman[1], **Thomas C. Shermer**[1], **and Ladislav Stacho**[2]

[1]School of Computing Science,
Simon Fraser University,
Burnaby BC, V5A 1S6 Canada,
{art,shermer}@sfu.ca

[2]Department of Mathematics,
Simon Fraser University,
Burnaby BC, V5A 1S6 Canada,
lstacho@sfu.ca

Abstract. A spanning subgraph $S = (V, E')$ of a connected graph $G = (V, E)$ is an $(x + c)$-spanner if for any pair of vertices u and v, $d_S(u, v) \leq d_G(u, v) + c$ where d_G and d_S are the usual distance functions in G and S, respectively. The parameter c is called the delay of the spanner. We study edge-disjoint spanners in graphs in multidimensional tori. We show that each two-dimensional torus has a set of two edge-disjoint spanners of delay approximately the size of the smaller dimension. Moreover, we show that this delay is close to the best possible. In three-dimensional tori, we find a set of three edge-disjoint spanners with delay approximately the sum of the sizes of the two smaller dimensions when all dimensions are of even size. Surprisingly, we also find a set of two edge-disjoint spanners in three-dimensional tori of constant delay. In d-dimensional tori, we show that for any $k \leq d/5$, there is a set of k edge-disjoint spanners with delay depending only on k and the size of the smaller k dimensions.

1 Introduction

A spanner of a graph is a spanning subgraph in which the distance between any pair of vertices approximates the distance in the original graph. Although spanners were introduced by Peleg and Ullman [19] for simulation of synchronous distributed systems, they are an interesting graph-theoretical structure with application to many problems in interconnection networks [2, 3, 17, 18]. The use of spanners as a network topology (as a substitute for an expensive original topology) was suggested by Richards and Liestman [20] and further studied in a series of papers by Liestman and Shermer [12–16] and Heydemann, Peters, and Sotteau [6]. Algorithms for constructing spanners have also been studied by others [1, 5, 7, 8].

In parallel computing, multiple computers are interconnected by a network of communication links. One problem encountered in parallel computing is to share these resources among several users concurrently. One way to approach this problem is to multitask on the computers but to dedicate each link to an individual user. In graph-theoretic terms, this corresponds to partitioning the edges into a set of edge-disjoint spanners. With various coauthors, we have studied edge-disjoint spanners in complete graphs, complete digraphs, complete bipartite graphs, hypercubes, and in Cartesian products of other graphs [10, 9, 4]. In this

* The research of the authors was supported by NSERC Canada grant.

paper, we continue this line of study, investigating edge-disjoint spanners in tori (Cartesian products of cycles).

A network is represented by a connected simple graph $G = (V, E)$. We use standard graph theoretic terminology as needed. See West [21] for example. We use $d_G(u, v)$ to denote the distance from vertex u to vertex v in graph G. Any spanning subgraph of G is called a *spanner* of G. A spanner S of G is called an $f(x)$-*spanner* if for any pair of vertices u and v, $d_S(u, v) \leq f(d_G(u, v))$. We call $d_S(u, v) - d_G(u, v)$ the *delay* between vertices u and v in S. For an $f(x)$-spanner S, we refer to $f(x) - x$ as the *delay* of the spanner. Note that $f(x) - x$ is an upper bound (but not necessarily a tight bound) on the maximum delay in S between any pair of vertices at distance x in G.

We use $H \times G$ to denote the Cartesian product of base graphs H and G. The vertex set $V(H \times G)$ is $V(H) \times V(G) = \{[u, v] : u \in V(H) \, \text{and} \, v \in V(G)\}$. The edge set $E(H \times G)$ contains all pairs $([u, v], [u', v'])$ such that either (1) $u = u'$ and $(v, v') \in E(G)$, or (2) $v = v'$ and $(u, u') \in E(H)$.

Our goal is to investigate small-delay spanners of tori. We are particularly interested in those spanners with constant delay, i.e. $(x + c)$-spanners for constant c. More precisely, given a constant c and a torus G, we are interested in the maximum number of edge-disjoint $(x + c)$-spanners that can be found in G. We let $EDS(G, c)$ denote this number.

In the next section, we focus on tori in two and three dimensions. In the final section, we consider higher dimensional tori.

2 Two and Three-Dimensional Tori

Let $T_{n_1,n_2,\ldots,n_d}$ denote the $n_1 \times n_2 \times \cdots \times n_d$ torus, i.e., the Cartesian product of d cycles $C_{n_1}, C_{n_2}, \ldots, C_{n_d}$. We follow the convention that $n_1 \geq n_2 \geq \cdots \geq n_d$, with $n_d \geq 2$ and $d \geq 2$. The vertices of the cycle C_{n_j} are labeled from 0 to $n_j - 1$ and the vertices of $T_{n_1,n_2,\ldots,n_d}$ are denoted $[i_1, i_2, \ldots, i_d]$, where i_j is the label of the corresponding vertex of C_{n_j}. For convenience, we use $[i_1, i_2, \ldots, i_d]$ to mean $[i_1 (\bmod\, n_1), i_2 (\bmod\, n_2), \ldots, i_d (\bmod\, n_d)]$. The edges of $T_{n_1,n_2,\ldots,n_d}$ are $([i_1, i_2, \ldots, i_d], [i_1{+}1, i_2, \ldots, i_d]), ([i_1, i_2, \ldots, i_d], [i_1, i_2{+}1, \ldots, i_d]), \ldots, ([i_1, i_2, \ldots, i_d], [i_1, i_2, \ldots, i_d{+}1])$ for all vertices $[i_1, i_2, \ldots, i_d]$.

In this section, we prove bounds on the number of edge-disjoint spanners that can be found in two and three-dimensional tori. We begin with the following observation which can be established using a simple counting argument.

Observation 1 *Any torus $T_{n_1,n_2,\ldots,n_d}$ has at most d edge-disjoint spanners; i.e. $EDS(n_1, n_2, \ldots, n_d, c) \leq d$ for all c.*

2.1 Two spanners in the two-dimensional torus

In this sub-section, we investigate edge disjoint spanners in the two-dimensional torus. By the above proposition, there can be at most two.

Theorem 1 *$EDS\left(T_{n_1,n_2}, 2\left\lceil \frac{n_2}{2} \right\rceil\right) = 2$, when n_1 is even.*

Proof. The upper bound follows from the above proposition. It remains to show that $EDS\left(T_{n_1,n_2}, 2\left\lceil\frac{n_2}{2}\right\rceil\right) \geq 2$. We construct one spanner S_1 as follows.

Figure 1 shows examples of this construction. We draw T_{n_1,n_2} and spanners of it in the plane in the standard fashion: each vertex $[i,j]$ of T_{n_1,n_2} is placed at Cartesian coordinates (i,j). Each edge except for those between columns n_2-1 and 0 and those between rows n_1-1 and 0 is drawn as the appropriate straight line segment. Each remaining edge (called a wrap-around edge) is represented by two half edges as shown in Figure 1 for spanner of $T_{10,7}$ and $T_{9,7}$.

(See Figure 1 a), for an example.) For each i, we include the edge $\left(\left[i, \left\lfloor\frac{n_2}{2}\right\rfloor\right], \left[i+1, \left\lfloor\frac{n_2}{2}\right\rfloor\right]\right)$; these edges form a cycle. For each even i, we add a path composed of edges $([i,j],[i,j+1])$ for $0 \leq j \leq \left\lfloor\frac{n_2}{2}\right\rfloor - 1$, and the pendant edges $([i,j],[i+1,j])$ where $1 \leq j \leq \left\lfloor\frac{n_2}{2}\right\rfloor - 1$. For each odd i, we add a path composed of edges $([i,j],[i,j+1])$ for $\left\lfloor\frac{n_2}{2}\right\rfloor \leq j \leq n_2-1$, and the pendant edges $([i,j],[i+1,j])$ where $\left\lfloor\frac{n_2}{2}\right\rfloor + 1 \leq j \leq n_2-1$.

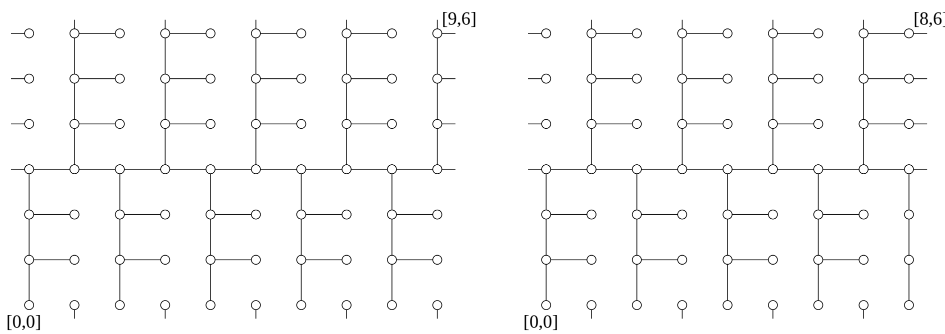

Fig. 1. a) The spanner S_1 in $T_{10,7}$. b) The spanner S_1 in $T_{9,7}$.

We now compute the delay of S_1. Consider a pair of vertices $u = [i_1, j_1]$ and $v = [i_2, j_2]$. In the torus, the distance between these vertices is $m_1 + m_2$ where m_1 is the distance in the first dimension, which is $\min\{(i_1 - i_2)(\bmod n_1), (i_2 - i_1)(\bmod n_1)\}$, m_2 is the distance in the second dimension, which is $\min\{(j_1 - j_2)(\bmod n_2), (j_2 - j_1)(\bmod n_2)\}$.

In S_1, a shortest path from u to v can be obtained as follows. Let u' be the vertex on the cycle of S_1 closest to u in S_1, and v' be the vertex on the cycle closest to v in S_1. Note that u' may be u and/or v' may be v.

A shortest path from u to v in S_1 consists of the path from u to u', followed by a shortest path from u' to v' along the cycle, followed by the shortest path from v' to v. The length of the path from u to u' is at most $\left\lceil\frac{n_2}{2}\right\rceil$, as is the distance from v' to v. The length of the path from u' to v' is at most $m_1 + 1$. This gives a total path length of at most $2\left\lceil\frac{n_2}{2}\right\rceil + m_1 + 1$ from u to v in S_1. Since the distance from u to v in the torus is at least m_1, this is a delay of at most $2\left\lceil\frac{n_2}{2}\right\rceil + 1$. Since the delay in the torus cannot be odd, we can subtract 1 from the previous bound, giving us delay $2\left\lceil\frac{n_2}{2}\right\rceil$.

The spanner S_2 is the complement of S_1 in T_{n_1,n_2}. Note that $S_1$and S_2 have same structure: The function mapping vertex $[i,j]$ to vertex $[i+1, \lfloor \frac{n_2}{2} \rfloor - j]$ is both an automorphism of T_{n_1,n_2} and an isomorphism from S_1 to S_2. Thus, the delay of S_2 is equal to the delay of S_1. We conclude that $EDS\left(T_{n_1,n_2}, 2\left\lceil \frac{n_2}{2} \right\rceil\right) = 2$ if n_1 is even.

We now consider the case when n_1 is odd.

Theorem 2 *$EDS\left(T_{n_1,n_2}, 2\left\lceil \frac{n_2}{2} \right\rceil + 2\right) = 2$, when n_1 is odd.*

Proof. Again, the upper bound follows from Proposition 1. We construct the spanner S_1 as follows. (See Figure 1 b), for an example.) First, we include all edges $\left(\left[i, \lfloor \frac{n_2}{2} \rfloor\right], \left[i+1, \lfloor \frac{n_2}{2} \rfloor\right]\right)$ forming a cycle. For each even i, we add a path composed of edges $([i,j],[i,j+1])$ for $0 \le j \le \lfloor \frac{n_2}{2} \rfloor - 1$. For each odd i, we add a path composed of edges $([i,j],[i,j+1])$ for $\lfloor \frac{n_2}{2} \rfloor \le j \le n_2 - 1$. For each even $i \le n_1 - 3$, we add the edges $([i,j],[i+1,j])$ where $1 \le j \le \lfloor \frac{n_2}{2} \rfloor - 1$. These edges are pendant to the path constructed above for the same i. For each odd $i \le n_1 - 2$, we add the edges $([i,j],[i+1,j])$ where $\lfloor \frac{n_2}{2} \rfloor + 1 \le j \le n_2 - 1$. Again, these edges are pendant to the path constructed above for the same i. Finally, we add the edges $([n_2-1,j],[0,j])$ where $\lfloor \frac{n_2}{2} \rfloor + 1 \le j \le n_2 - 1$. This construction is similar to that of S_1 in the proof of Theorem 1 except for minor adjustments to accommodate an imbalance of the number of dimension one (horizontal) paths on either side of the cycle. These adjustments appear on the vertical wrap-around edges in the figure.

We now compute the delay for S_1. Consider a pair of vertices $u = [i_1, j_1]$ and $v = [i_2, j_2]$. In the torus, the distance between these vertices is $m_1 + m_2$ where $m_1 = \min\{(i_1 - i_2)(\text{mod}\, n_1), (i_2 - i_1)(\text{mod}\, n_1)\}$ and $m_2 = \min\{(j_1 - j_2)(\text{mod}\, n_2), (j_2 - j_1)(\text{mod}\, n_2)\}$.

In S_1, a shortest path from u to v can be obtained as follows. Let u' be the vertex on the cycle of S_1 closest to u, and v' be the vertex on the cycle closest to v. Note that u' may be u and/or v' may be v.

A shortest path from u to v in S_1 consists of the path from u to u', followed by a shortest path from u' to v' along the cycle, followed by the shortest path from v' to v. The length of the path from u to u' is at most $\lceil \frac{n_2}{2} \rceil + 1$, as is the distance from v' to v. If both of these distances are equal to $\lceil \frac{n_2}{2} \rceil + 1$, then $j_1 = j_2 = 0$ and consequently, $d_{S_1}(u', v') = m_1 = 0$. Otherwise, $d_{S_1}(u', v') \le m_1 + 2$. In both cases, this gives a total path length of at most $2\left\lceil \frac{n_2}{2} \right\rceil + m_1 + 3$ from u to v. Since the distance from u to v in the torus is at least m_1, this is a delay of at most $2\left\lceil \frac{n_2}{2} \right\rceil + 3$. Since the delay in the torus cannot be odd, the upper bound is $2\left\lceil \frac{n_2}{2} \right\rceil + 2$.

The spanner S_2 is the complement of S_1 in T_{n_1,n_2}. Note that now the spanners S_1 and S_2 need not be self-complementary factors of T_{n_1,n_2}. However, the delay of S_2 is at most $2\left\lceil \frac{n_2}{2} \right\rceil + 2$, by a similar argument. We conclude that $EDS\left(T_{n_1,n_2}, 2\left\lceil \frac{n_2}{2} \right\rceil + 2\right) = 2$ if n_1 is odd.

In summary, we can find two edge-disjoint spanners in T_{n_1,n_2} with delay between n_2 and $n_2 + 3$, depending on the parity of n_1 and n_2. In the following Theorem, we show that n_2 is the best possible delay, matching Theorem 1 when n_1 and n_2 are both even.

Theorem 3 *If $n_2 \geq 3$, then $EDS(T_{n_1,n_2}, c) = 1$ for any delay $c < n_2$.*

Proof. We use the term *row* l to denote the vertices $[i, l]$ where $0 \leq i \leq n_1 - 1$. Similarly, the term *column* l is used to denote the vertices $[l, j]$ where $0 \leq j \leq n_2 - 1$. The *dual* $T^*_{n_1,n_2}$ of the torus T_{n_1,n_2} is the usual dual graph when T_{n_1,n_2} is embedded in a geometric torus: $T^*_{n_1,n_2}$ has a vertex corresponding to each face of T_{n_1,n_2}, and an edge between any two faces that share an edge. Note that $T^*_{n_1,n_2}$ is isomorphic to T_{n_1,n_2}. Since each edge of T_{n_1,n_2} separates two faces, each edge has a unique *dual edge* of $T^*_{n_1,n_2}$ which joins the vertices corresponding to those two faces.

We sometimes include the dual in our standard drawing of the torus, placing each vertex of the dual in the center of the corresponding face of the primal. This is shown in Figure 2 b) for $T_{3,4}$ where the dual edges are represented by dashed lines.

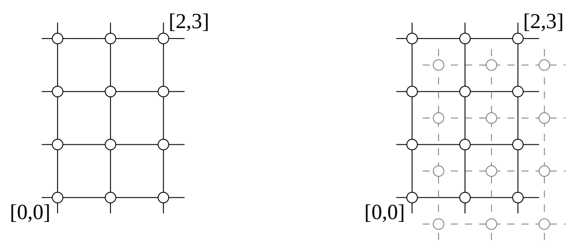

Fig. 2. Drawings of a) $T_{3,4}$, and b) $T^*_{3,4}$ in $T_{3,4}$.

For any subgraph S of T_{n_1,n_2}, we define the *dual* S^* of S as the subgraph of $T^*_{n_1,n_2}$ having all of the vertices of $T^*_{n_1,n_2}$ and exactly those edges that are dual to edges of S. As is easily seen in Figure 2, the only edge of the dual that intersects a given edge is the dual of that particular edge.

By way of contradiction, let S_1 and S_2 be two edge-disjoint spanners of T_{n_1,n_2} each with delay less than n_2. Without loss of generality assume that each edge of T_{n_1,n_2} is in either S_1 or S_2, and that S_2 has at least as many edges as S_1. Since T_{n_1,n_2} has $2n_1n_2$ edges, S_2 must have at least n_1n_2 edges. Therefore, S_2^* also has at least n_1n_2 edges and must contain a cycle C^*. For each $e^* = (a, b)$ of C^* let $e = (u, v)$ be the edge of S_2 dual to e^*. Let P_e be a shortest path from u to v in S_1, and P'_e be the cycle $P_e + e$.

Consider the situation when there are indices i and j such that C^* contains no edge dual to any edge between vertices of columns j and $j + 1$ and no edge dual to any edge between vertices of rows i and $i + 1$. By renaming the rows and columns, we may without loss of generality assume that $i = n_1 - 1$ and $j = n_2 - 1$. With this assumption, C^* uses none of the half edges of our standard drawing of T_{n_1,n_2} and its dual. Therefore, in this drawing, C^* is a closed curve with exactly one of u and v inside and the other outside. Since P_e is a path from u to v, it must cross the drawing of C^* by the Jordan Curve Theorem. However, this is a contradiction because no edge of S_1 can cross the dual of any edge of S_2.

Therefore, it must be the case that C^* either crosses every column (that is, for each $0 \le j \le n_2 - 1$, C^* contains an edge dual to an edge joining vertices of column j) or crosses every row.

If there are indices i and j such that P'_e contains no edge between columns j and $j+1$ and no edge between rows i and $i+1$, then we may arrive at a contradiction via the Jordan Curve Theorem, as above (with P'_e playing the role of C^* and vice versa). We can conclude that P'_e either has an edge between column j and column $j+1$ for each $0 \le j \le n_2 - 1$, or it has an edge between row i and row $i+1$ for each $0 \le i \le n_1 - 1$.

Suppose that C^* crosses every row and let $e^* = (a, b)$ be an edge of C^* that crosses the edge $e = (u, v)$ in a row i. If P'_e contains an edge between rows i and $i+1$ for each $0 \le i \le n_1 - 1$, then P_e contains all of these edges plus at least one edge between adjacent columns because u and v are in different columns. Thus P_e contains at least $n_1 + 1$ edges giving delay of at least $n_1 \ge n_2$ between u and v in S_1, a contradiction.

Therefore, P'_e contains an edge between columns j and $j+1$ for every $0 \le j \le n_2 - 1$. If P'_e also contains an edge between some adjacent rows, then it must contain at least two such edges, and then the length of P_e would be at least $n_2 - 1 + 2 = n_2 + 1$, giving delay at least n_2 between u and v in S_1, a contradiction. Hence, the path P_e consists of all of the edges in row i, except edge e.

Since C^* crosses each row, spanner S_1 lacks at least one edge $e_i = ([i, m_i], [i, m_i + 1])$ from each row i. Furthermore, the above argument implies that e_i is the only edge between vertices of row i that is not in S_1. Thus, S_1 contains $n_1(n_2 - 1)$ edges in dimension two. As S_1 must be connected and cannot have more than half the edges of T_{n_1,n_2} by assumption, then S_1 must contain either $n_1 - 1$ or n_1 edges in dimension one.

Suppose there is some row i such that S_1 contains no edge between a vertex of row i and a vertex of row $i+1$. Let l be a column and consider a shortest path in S_1 between the vertices $[i, l]$ and $[i+1, l]$. This path must contain an edge between every pair of consecutive rows except rows i and $i+1$. This accounts for $n_1 - 1$ of the at most n_1 edges of S_1 between different rows. So, between rows $i-1$ and i, there are at most two edges of S_1. Since $n_2 \ge 3$, we may choose a column l such that there is no edge between vertices $[i-1, l]$ and $[i, l]$ in S_1. Any shortest path between $[i, l]$ and $[i+1, l]$ in S_1 contains at least two edges between vertices in the same row. This means that such a path has length at least $n_1 + 1$, giving a delay of at least n_1, a contradiction. Therefore S_1 contains exactly $n_1 n_2$ edges including precisely one edge f_i between each pair of consecutive rows i and $i+1$.

Suppose that for some i, the edges f_i and f_{i+1} are in the same column. Without loss of generality they are in column 0 with common vertex $[i+1, 0]$. The edge e_{i+1} missing from row $i+1$ in S_1 must be either $([i+1, 0], [i+1, 1])$ or $([i+1, n_2 - 1], [i+1, 0])$ as otherwise the vertex $[i+1, 0]$ would have degree 0 in S_2. Without loss of generality, $e_{i+1} = ([i+1, n_2 - 1], [i+1, 0])$. Any path in S_1 from any vertex on row $i+1$ to any vertex not on row $i+1$ must include $[i+1, 0]$. The path in S_1 from $[i+1, n_2 - 1]$ to $[i, n_2 - 1]$ therefore includes $n_2 - 1$ edges of row $i+1$ to get to $[i+1, 0]$ and at least 2 edges to get from $[i+1, 0]$ to $[i, n_2 - 1]$ yielding a delay of n_2, a contradiction. thus there is no i so that f_i and f_{i+1} are in the same column.

Consider any pair of rows i and $i+1$ and the edge f_i between them in S_1. Without loss of generality, let $f_i = ([i,0],[i+1,0])$. Note that the shortest path in S_1 from any vertex $[i,l]$ to any vertex $[i+1,l]$ must include f_i (and no other f_j) as otherwise their delay would be at least n_1. Consider m_i and m_{i+1} (the columns to the left of missing edges e_i and e_{i+1}, respectively). If $m_i > m_{i+1}$, then the delay from $[i,m_i]$ to $[i+1,m_i]$ is n_2. Similarly, if $m_{i+1} > m_i$, then the delay from $[i+1,m_{i+1}]$ to $[i,m_{i+1}]$ is n_2. Thus, $m_i = m_{i+1}$. It follows that $m_0 = m_1 = \cdots = m_{n_1-1}$.

Note that if n_2 is even, then either $a = [i,m_i]$ or $b = [i,m_i+1]$ is at distance at least $n_2/2$ from $[i,0]$. Thus, the distance from either a or b to its neighbor (in T_{n_1,n_2}) in row $i+1$ is at least n_2+1 in S_1, a contradiction. If n_2 is odd, then $m_i = (n_2-1)/2$ or a similar contradiction arises. In other words, the distances from $[i,m_i]$ and $[i,m_i+1]$ to the edge f_i must be exactly $(n_2-1)/2$. Since $m_i = m_{i+1}$, applying the argument to the pair of row $i+1$ and $i+2$, we conclude that f_i and f_{i+1} are in the same column, a contradiction.

Since in any case we have shown that our original assumption that S_1 and S_2 are edge-disjoint spanners of T_{n_1,n_2} both with delay less than n_2 is false. To complete the proof we have to consider the case when C^* crosses every column (instead of every row). The arguments which lead to a contradiction are similar to those in the previous case.

In summary, we have nearly determined the least delay that allows the construction of two edge-disjoint spanners in T_{n_1,n_2}. We believe that the delays of Theorems 1 and 2 are, for the most part, the best possible. However, the situation is at least slightly more complicated as Theorem 2 can be improved when $n_1 = n_2$. Figure 3 shows an example of a pair of edge-disjoint spanners with delay 7 in $T_{7,7}$. This construction is generalizable to two spanners of delay n in $T_{n,n}$.

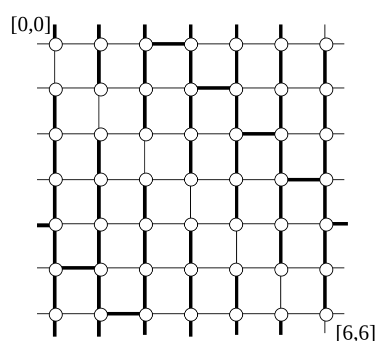

Fig. 3. The two spanners in $T_{7,7}$ are depicted in bold and regular lines, respectively.

2.2 Three spanners in three-dimensional tori

In this and the next subsection, we investigate edge-disjoint spanners in the three-dimensional torus. There can be at most three such spanners; in this section we restrict our attention to when there are exactly three.

Theorem 4 *$EDS(T_{n_1,n_2,n_3}, n_2+n_3+6) = 3$, whenever n_1, n_2, and n_3 are all even and $n_3 \geq 4$.*

Proof. The upper bound comes from Proposition 1. We can construct three spanners of delay $n_2 + n_3 + 6$ in the three-dimensional torus. For any fixed i, we refer to the set of vertices $\{[i, x, y] : 0 \leq x \leq n_2 - 1, 0 \leq y \leq n_3 - 1\}$ and the edges connecting them as "layer i": Each layer is a two-dimensional torus. Layers are either even or odd depending on the parity of i.

Informally, the main idea of the construction is to use the two spanner pattern of Theorem 1 in each layer. In even layers, we use this pattern for spanners S_1 and S_2 and in odd layers for S_1 and S_3. To connect spanner S_2 to the vertices in an odd layer i, we use all of the edges between layers i and $i + 1$. Similarly, to connect spanner S_3 to the vertices in an even layer i, we use all of the edges between layers i and $i+1$. As stated, each of these "spanners" is disconnected. We reassign edges to create a dimension-one cycle for each spanner, interconnecting its different layers. This necessitates some further local adjustments.

To be more specific, the dimension-one cycle in any spanner S_i is called the *hub cycle* of S_i. The hub cycle of S_1 consists of all vertices $\{h_1^i = [i, 0, 0] : 0 \leq i \leq n_1 - 1\}$ and the edges of dimension one connecting them. The hub cycle of S_2 consists of all vertices $\{h_2^i = [i, 0, \frac{n_3}{2}] : 0 \leq i \leq n_1 - 1\}$ and the edges of dimension one connecting them. The hub cycle of S_3 consists of all vertices $\{h_3^i = [i, \frac{n_2}{2}, 0] : 0 \leq i \leq n_1 - 1\}$ and the edges of dimension one connecting them.

In each even layer i, there are two edges of S_3 and the other edges are split equally between S_1 and S_2 using a modification of Figure 1 a). First, the pattern of Figure 1 a) is placed so that the highway of S_1 is in dimension two and passes through the hub vertex h_1^i as shown in Figure 4 a).

Of necessity, this means that the highway of S_2 is also in dimension two and passes through the hub vertex h_2^i. Next, to connect the hub vertex h_1^i to spanner S_3 we let $g^i = [i, 0, n_3 - 1]$ and reassign the edge (h_1^i, g^i) from S_1 to S_3. Refer to Figure 4 b) for the result of this and the subsequent adjustments. The reassignment of (h_1^i, g^i) disconnects S_1 within this layer. Let $r^i = [i, \frac{n_2}{2} - 1, 0]$. To reconnect S_1, we break the highway cycle of S_1 in this layer by reassigning (r^i, h_3^i) to S_2, for each j, $1 \leq j \leq \frac{n_2}{2} - 2$, changing the assignment of edges $([i, j, 0], [i, j, n_3 - 1])$ and $([i, j, 0], [i, j, 1])$ from S_1 to S_2 or vice versa, and reassigning the edge $(r^i, [i, \frac{n_2}{2} - 1, 1])$ from S_2 to S_1. Similarly, to connect h_2^i to spanner S_3, we let $p^i = [i, 0, \frac{n_3}{2} - 1]$ and reassign the edge (h_2^i, p^i) from S_2 to S_3, disconnecting S_2 within this layer. Letting $h^i = [i, \frac{n_2}{2}, \frac{n_3}{2}]$ and $w^i = [i, \frac{n_2}{2} - 1, \frac{n_3}{2}]$, we reassign (w^i, h^i) to S_1, for each j, $1 \leq j \leq \frac{n_2}{2} - 2$, change the assignment of edges $([i, j, \frac{n_3}{2}], [i, j, \frac{n_3}{2} - 1])$ and $([i, j, \frac{n_3}{2}], [i, j, \frac{n_3}{2} + 1])$ from S_1 to S_2 or vice versa, and reassigning the edge $(w^i, [i, \frac{n_2}{2} - 1, \frac{n_3}{2} + 1])$ from S_1 to S_2. This completes the construction in even layers and a sample layer is depicted in Figure 4 b).

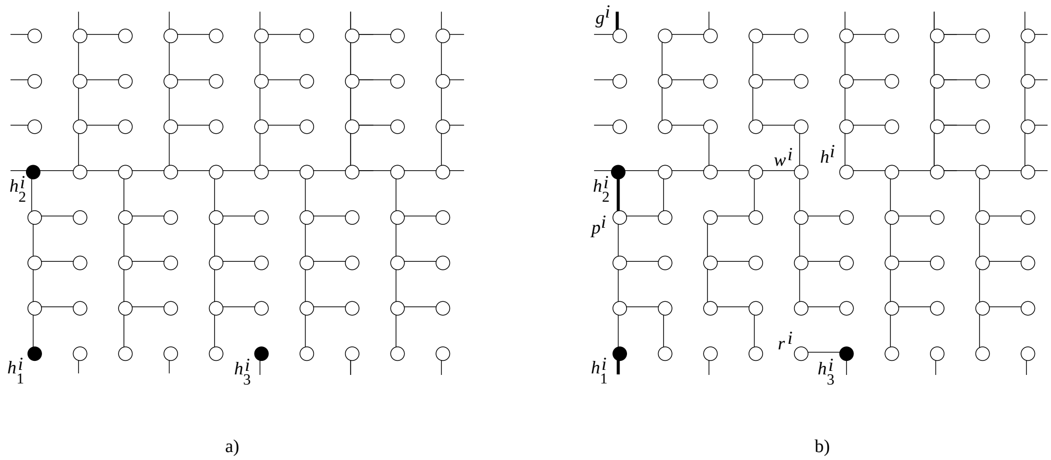

Fig. 4. a) Starting configuration in an even layer. b) The resulting configuration in an even layer. The bold lines depict edges of S_3, regular lines depict edges of S_2. The spanner S_1 contains all remaining edges in this layer (they are not depicted in the figure).

The construction of odd layers is analogous. In each odd layer i, there are two edges of S_2 and the other edges are split equally between S_1 and S_3 using a modification of Figure 1 a). First, the pattern of Figure 1 a) is placed so that the highway of S_1 is in dimension three and passes through the hub vertex h_1^ias shown in Figure 5 a). Of necessity, this means that the highway of S_3 is also in dimension three and passes through the hub vertex h_3^i.

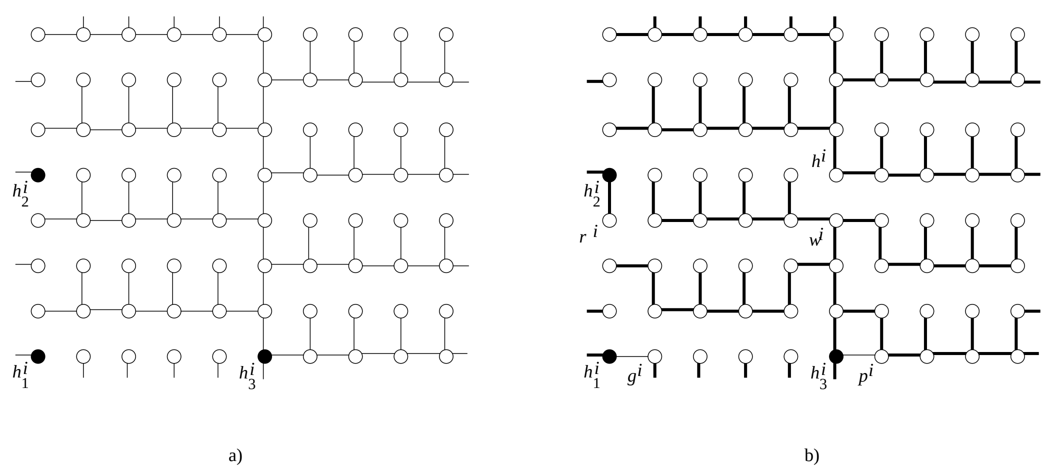

Fig. 5. a) Starting configuration in an odd layer. b) The resulting configuration in an odd layer. The bold lines depict edges of S_3, regular lines depict edges of S_2. The spanner S_1 contains all remaining edges in this layer (they are not depicted in the figure).

Next, to connect the hub vertex h_1^i to spanner S_2 we let $g^i = [i, 1, 0]$ and reassign the edge (h_1^i, g^i) from S_3 to S_2. Refer to Figure 5 b) for the result of this and the subsequent adjustments. The reassignment of (h_1^i, g^i) disconnects S_3

within this layer. Let $r^i = [i, 0, \frac{n_3}{2} - 1]$. To reconnect S_3, we break the highway cycle of S_1 in this layer by reassigning (r^i, h_2^i) to S_3, for each j, $1 \leq j \leq \frac{n_3}{2} - 2$, changing the assignment of edges $([i, 0, j], [i, n_2 - 1, j])$ and $([i, 0, j], [i, 1, j])$ from S_1 to S_3 or vice versa, and reassigning the edge $(r^i, [i, 1, \frac{n_3}{2} - 1])$ from S_3 to S_1. Similarly, to connect h_3^i to spanner S_2, we let $p^i = [i, \frac{n_2}{2}+1, 0]$ and reassign the edge (h_3^i, p^i) from S_1 to S_2, disconnecting S_1 within this layer. Letting $h^i = [i, \frac{n_2}{2}, \frac{n_3}{2}]$ and $w^i = [i, \frac{n_2}{2}, \frac{n_3}{2} - 1]$, we reassign (w^i, h^i) to S_1, for each j, $1 \leq j \leq \frac{n_3}{2} - 2$, change the assignment of edges $([i, \frac{n_2}{2}, j], [i, \frac{n_2}{2} - 1, j])$ and $([i, \frac{n_2}{2}, j], [i, \frac{n_2}{2} + 1, j])$ from S_1 to S_3 or vice versa, and reassign the edge $(w^i, [i, \frac{n_2}{2}+1, \frac{n_3}{2}-1])$ from S_1 to S_3. This completes the construction in odd layers and a sample layer is depicted in Figure 5 b).

To route between an arbitrary pair of vertices u and v in S_i, we first find u' and v', the closest hub vertices in S_i to u and v, respectively. We can then construct a path from u to u' to v' to v. The distance from u to u' is at most $\frac{n_2+n_3}{2}+1$ if u and u' are in the same layer. If u and u' are in adjacent layers, then there is a path of at most two edges from u in S_i to a vertex u'' in the same layer as u'. Thus the distance from u to u' is at most $\frac{n_2+n_3}{2}+3$. The same bounds apply for the distance from v to v'. If u and u' are in the same layer and v and v' are in the same layer, then the distance from u' to v' is the distance in dimension one from u to v, and therefore the delay from u to v is at most $2(\frac{n_2+n_3}{2} + 1) = n_2 + n_3 + 2$. If both u and v are in different layers than u' and v', respectively, then again the distance from u' to v' is the distance in dimension one from u to v, and therefore the delay from u to v is at most $2\left(\frac{n_2+n_3}{2} + 3\right) = n_2 + n_3 + 6$. Otherwise, the distance from u' to v' may be one larger than the distance in dimension one from u to v, giving a delay of at most $\left(\frac{n_2+n_3}{2} + 1\right) + 1 + \left(\frac{n_2+n_3}{2} + 3\right) = n_2 + n_3 + 5$. Thus, in any case, the delay is at most $n_2 + n_3 + 6$. This is best possible, as vertices h_2^0 and h_2^2 in spanner S_3 have this delay.

2.3 Two spanners in three-dimensional tori

In this subsection, we investigate the conditions under which there are two edge-disjoint spanners in the three-dimensional torus. We first show that when each of the dimensions is sufficiently large then the delay of the two spanners must be at least 4. We then show that there are two spanners of delay 8 for a class of arbitrarily large tori.

Let S_1 and S_2 be two edge-disjoint spanners of a three-dimensional torus. We call a vertex v *pierced in dimension* ρ in the spanner S_1 (similarly in S_2) if both edges in dimension ρ incident on v are included in S_1 (in S_2). We call a vertex v *ordinary in dimension* ρ in S_1 (in S_2) if either v is pierced in dimension ρ in S_1 (in S_2) or it is adjacent by an edge in dimension different from ρ in S_1 (in S_2) to some vertex that is pierced in dimension ρ in S_1 (in S_2). Finally, we call a vertex *ordinary in* S_1 (in S_2) if it is ordinary in each of the three dimensions in S_1 (in S_2). By *subgrid* of a torus we mean a subgraph that is a grid. To obtain our first result, we require the following two lemmas.

Lemma 1 *Let $j \geq 2$ and $n_3 \geq 122j^2 + 5$. If S_1 and S_2 are edge-disjoint spanners of $T = T_{n_1,n_2,n_3}$ of delay 2, then T contains a $j \times j \times 1$ subgrid containing only vertices ordinary in both S_1 and S_2.*

Proof. Let k be an integer such that $k \leq (n_3 - 5)/2$ and consider a $k \times k \times k$ subgrid B of T. Let B' be the $(k+2) \times (k+2) \times (k+2)$ subgrid centered on and surrounding B. Since $n_3 \geq 2k + 5$ and both S_1 and S_2 have delay 2, every shortest path in either S_1 or S_2 between a pair of vertices in B stays within B'. We proceed to find an upper bound on the number of non-ordinary vertices (in either spanner) in B.

We use the term *line* to denote a longest contiguous path in a single dimension in B. Each line extends from one side of the "box" B to the opposite side. There are k^2 distinct lines in each dimension, or $3k^2$ distinct lines in toto. Consider a single line L of consecutive vertices $l_1, l_2, \ldots, l_k$ and, without loss of generality, suppose that L is in dimension one. Let l_0 and l_{k+1} be the vertices of B' adjacent in dimension one to the vertices l_1 and l_k, respectively. We use L^* to denote the path of consecutive vertices $l_0, l_1, \ldots, l_k, l_{k+1}$.

Consider the subgraph of S_1 induced by the vertices of L^*. Let p be the number of connected components in this subgraph. Since L^* is a path, each of these components is also a path on consecutive vertices $l_{u_i}, l_{u_i+1}, \ldots, l_{v_i-1}, l_{v_i}$. We assume that these paths are labeled so that $v_i = u_{i+1} - 1$ for $1 \leq i \leq p-1$. Also note that it is possible for a path to have length zero, i.e., $u_i = v_i$. We claim that there are at most 10 vertices of L that are not ordinary in dimension one in S_1: possibly l_{v_1}, l_{u_p}, at most four other l_{u_i}'s and at most four other l_{v_i}'s. First, note that the internal vertices of each of these paths are pierced and therefore ordinary in dimension one in S_1. Thus, we are only concerned with whether the l_{u_i}'s and l_{v_i}'s are ordinary in S_1. Note that $l_{u_1} = l_0$ and $l_{v_p} = l_{k+1}$ are not part of L. Suppose by way of contradiction that there are more than four l_{u_i}'s, with $1 \leq i \leq p$, that are not ordinary in dimension one in S_1. The shortest paths from each of these to l_0 start with an edge that is not in dimension one. There are four "directions" possible for these first edges: increasing or decreasing, in dimension two or 3. By the pigeon hole principle, there is a pair of these vertices l_{u_x}, l_{u_y}, where $x < y$, which use the same direction. Let l'_{u_x} and l'_{u_y} be the neighbors of l_{u_x} and l_{u_y}, respectively, in this direction. The shortest path P from l_{u_y} to l_0 in this direction proceeds from l_{u_y} to l'_{u_y}, follows edges in dimension one to some vertex l'_b, follows one edge in dimension other than 1 to $l_b \in L^*$, and zero or more edges in dimension one to l_0. Note that $b < u_2$, as l_b is in the same component as l_0. Therefore, the subpath of P from l'_{u_y} to l'_b passes through l'_{u_x} and therefore, l'_{u_x} is pierced in dimension one in S_1. Thus, l_{u_x} is ordinary in dimension one in S_1, a contradiction. See Figure 6.

Hence there are at most four l_{u_i}'s, with $1 \leq i \leq p$, that are not ordinary in dimension one in S_1. A similar argument about shortest paths from the l_{v_i}'s to l_{k+1} instead of from l_{u_i}'s to l_0 completes the claim.

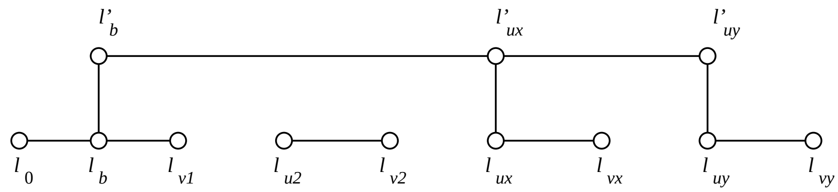

Fig. 6. Components of a line in dimension one in S_1. The vertex l_{v_y} must be joined by a path which has at most two edges in dimension other than one to the vertex l_0. This will force the vertex l_{u_x} to be pierced in dimension one.

Similarly, there are at most 10 vertices of L that are not ordinary in dimension one in S_2. Thus, there are at most 20 vertices of L that are not ordinary in dimension one in either S_1 or S_2.

Recall that there are $3k^2$ distinct lines through the subgrid B. The claim, therefore, implies that there are at most $60k^2$ vertices that are not ordinary in some dimension in either S_1 or S_2; the remaining $k^3 - 60k^2$ vertices are ordinary (in all three dimensions in both spanners). Consider the $k(k-j+1)^2$ distinct $j \times j \times 1$ subgrids contained in B. Each of the at most $60k^2$ non-ordinary vertices is contained in at most j^2 of these subgrids. Thus, there are at least $k(k-j+1)^2 - 60k^2j^2$ subgrids containing only ordinary vertices. We require $k(k-j+1)^2 - 60k^2j^2 \geq 1$. By choosing $k = 61j^2$ we satisfy this condition and since $n_3 \geq 122j^2 + 5$, the condition $k \leq (n_3 - 5)/2$ is also satisfied.

Lemma 2 *If there exists a vertex v that is ordinary and pierced in more than one dimension in S_1, then S_2 has delay at least 4. Similarly, if there exists a vertex v that is ordinary and pierced in more than one dimension in S_2, then S_1 has delay at least 4.*

Proof. Suppose by way of contradiction that S_2 has delay at most 2. If v is pierced in three dimensions in S_1, then it has degree 0 in S_2, a contradiction. Otherwise, let w be the vertex adjacent in S_1 to v that is pierced in the dimension that v is not as shown in Figure 7. The vertex w exists since v is ordinary. Now, the delay between v and w in S_2 is at least 4, a contradiction.

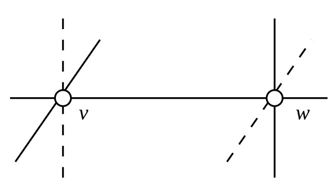

Fig. 7. Vertices v and w in spanners S_1(indicated by solid lines) and S_2(indicated by dashed lines).

Theorem 5 *If $n_3 \geq 3055$, then $EDS(T_{n_1,n_2,n_3}, 2) = 1$.*

Proof. Assume by way of contradiction that for some $n_3 \geq 3055$ the torus T_{n_1,n_2,n_3} has two edge-disjoint spanners S_1 and S_2, both of delay 2. By Lemma 1, there is a $5 \times 5 \times 1$ subgrid B of the torus T such that all vertices of B are ordinary in both spanners. Let u be the vertex at the center of B. Since u is ordinary in S_1, it is either pierced in dimension three (the "thin" dimension) in which case we let $v = u$ or it is adjacent to a vertex v of B that is pierced in dimension three. The vertex v must have degree at least 3 in S_1 as otherwise v is twice pierced in S_2 and using Lemma 2 we obtain a contradiction. Let a and b be the neighbors of v in dimension 3 and let c_1, c_2, c_3, c_4 be the neighbors of v in B, ordered successively around v. Without loss of generality assume that c_1 is the neighbor of v and that it is the neighbor of v in dimension two, see Figure 8 a).

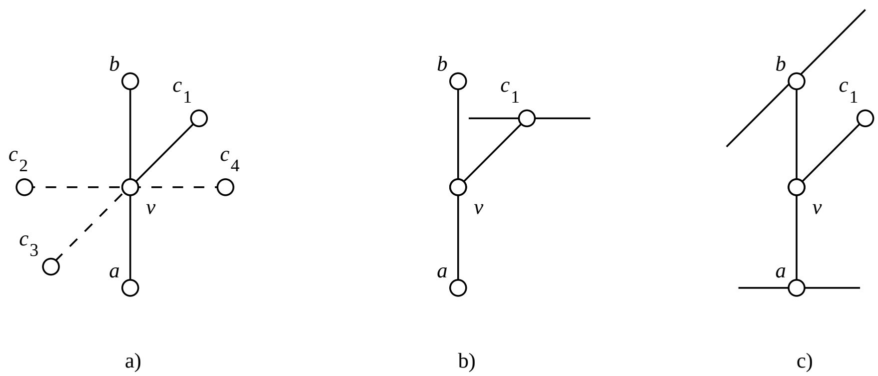

Fig. 8. A close neighborhood of the vertex v in the spanner S_1.

We claim that v has degree exactly 3 in S_1. Certainly, c_3 is not a neighbor of v in S_1 as Lemma 2 would give a contradiction. Again, by Lemma 2, at most one of c_2 and c_4 can be a neighbor of v in S_1. Without loss of generality assume it is c_2. Since v is ordinary and not pierced in S_2, its neighbors c_3 and c_4 together must be pierced in S_2 in all three dimensions. Thus, one of these vertices is pierced in S_2 in two different dimensions. Since both c_3 and c_4 are in B, they are ordinary so we have a contradiction by Lemma 2. Thus, our claim that v has degree exactly 3 in S_1 is established.

Vertex v must be adjacent in S_1 to a vertex that is pierced in dimension one. If this vertex is c_1, then the delay between v and c_1 in S_2 is at least 4, a contradiction; see Figure 8 b). Otherwise, without loss of generality, this vertex is a. Vertex v must also be adjacent in S_1 to a vertex other than c_1 pierced in dimension two, by the definition of ordinary. If this vertex is a, then a has degree at most 1 in S_2 and the delay from a to v in S_2 is at least 4, a contradiction. Thus, b is pierced in dimension two in S_1. Now the delay from a to b in S_1 is at least 4, a contradiction; see Figure 8 c).

Theorem 6 *$EDS(T_{n_1,n_2,n_3}, 8) \geq 2$, whenever any two of n_1, n_2, n_3 are congruent to $0 (\bmod 4)$.*

Proof. To show that $EDS(T_{n_1,n_2,n_3}, 8) \geq 2$, we begin by considering the case when n_1 and n_2 are congruent to 0 mod 4, and n_3 is even. For any $0 \leq i \leq n_3 - 1$, we call the two-dimensional torus on the set of vertices $L_i = \{[x, y, i] \ : \ 0 \leq x \leq n_1 - 1, 0 \leq y \leq n_2 - 1\}$ the *i-layer*. Further, we call a layer *even* or *odd* depending on whether i is even or odd. We use the term *type* 1 vertex to indicate a vertex $[x, y, i]$ such that $x + y \equiv 0 (\mathrm{mod}\, 4)$ or $x + y \equiv 1 (\mathrm{mod}\, 4)$. The remaining vertices are called *type* 2 vertices.

We construct one spanner S_1 as follows. In every even layer, include each edge in dimension one. In every odd layer, include each edge in dimension two. For every type 1 vertex $[x, y, i]$, include the edge $([x, y, i], [x, y, i + 1])$. Figure 9 a) shows a portion of an even layer of S_1, and the same portion of an odd layer (depicted in b)). In the figure, the type 1 vertices are indicated as solid circles.

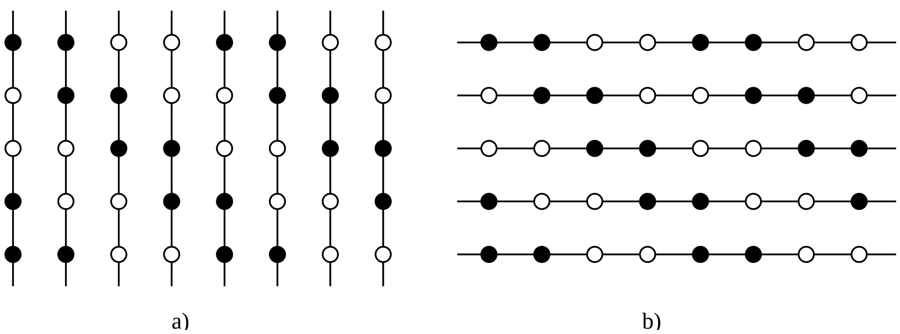

Fig. 9. a) an even layer of S_1 b) an odd layer of S_1.

We let S_2 be the complement of S_1 in the torus. Note that S_2 has an isomorphic structure to S_1. Furthermore, given any two type 1 vertices in the same layer, there is an automorphism of the torus that maps one of these vertices to the other while preserving the partition of the edges into S_1 and S_2. This is true even if the two type 1 vertices are in different layers of the same parity. This also holds for type 2 vertices. Thus, in order to determine the delay of the spanner, it suffices to consider the delay from four fixed vertices: a type 1 vertex in an even layer, a type 2 vertex in an even layer, a type 1 vertex in an odd layer, and a type 2 vertex in an odd layer. Below, we analyze the delay for sources in the even layer 0. For sources in odd layers, we can use the same descriptions as for sources in even layers except that we exchange the roles of dimensions one and two.

Consider paths in S_1 from a type 1 vertex, without loss of generality $[0, 0, 0]$. We first show that there is a path in S_1 from $[0, 0, 0]$ to any other vertex $[0, y, 0]$ with delay at most 8. If $y \equiv 0 (\mathrm{mod}\, 4)$ or $1 (\mathrm{mod}\, 4)$, there is a path from $[0, 0, 0]$ to $[0, y, 0]$ of delay 2 including the vertices $[0, 0, 1]$ and $[0, y, 1]$. If $y \equiv 2 (\mathrm{mod}\, 4)$, there is a path from $[0, 0, 0]$ to $[0, y, 0]$ of delay 8. This path starts at $[0, 0, 0]$ and proceeds directly to $[0, 0, 1]$ to $[0, 1, 1]$ to $[0, 1, 0]$ to $[-1, 1, 0]$ to $[-1, 1, 1]$. From $[-1, 1, 1]$ it proceeds directly dimension two edges to $[-1, y, 1]$, and then directly to $[-1, y, 0]$ and $[0, y, 0]$. If $y \equiv 3 (\mathrm{mod}\, 4)$, there is a path from $[0, 0, 0]$ to $[0, y, 0]$ of delay 4. This path starts at $[0,0,0]$ and proceeds directly to $[1,0,0]$ and $[1,0,1]$. This is followed by a path of dimension two edges to $[1,y,1]$ and then directly to $[1,y,0]$ and $[0,y,0]$.

There is a path in S_1 from $[0,0,0]$ to any other vertex $[x,y,0]$ with delay at most 8. This path consists of the delay at most 8 path from $[0,0,0]$ to $[0,y,0]$ followed by edges in dimension one to vertex $[x,y,0]$. There is also a path in S_1 from $[0,0,0]$ to any other vertex $[x,y,i]$ with delay at most 8. This path consists of edges in dimension three from $[0,0,0]$ to $[0,0,i]$ followed by the path of delay at most 8 to the vertex $[x,y,i]$.

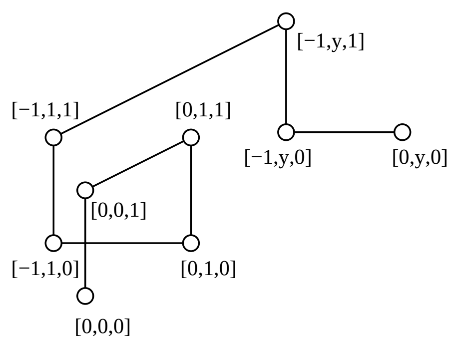

Fig. 10. A path with delay at most delay 8 from $[0,0,0]$ to $[0,y,0]$ where $y \equiv 2(\bmod 4)$.

Now, consider paths in S_1 from a type 2 vertex, without loss of generality $[1,1,0]$. We first show that there is a path in S_1 from $[1,1,0]$ to any other vertex $[1,y,0]$ with delay at most 8. If $y \equiv 1(\bmod 4)$ or $y \equiv 2(\bmod 4)$, there is a path from $[1,1,0]$ to $[1,y,0]$ of delay 6. This path starts at $[1,1,0]$ and proceeds directly to $[2,1,0]$, to $[3,1,0]$ to $[3,1,1]$. From $[3,1,1]$ it proceeds along dimension two edges to $[3,y,1]$ and then directly to $[3,y,0]$ to $[2,y,0]$ to $[1,y,0]$. If $y \equiv 3(\bmod 4)$ or $y \equiv 0(\bmod 4)$, then $[1,y,0]$ is a type 1 vertex and there is a path of delay 8 in S_1 between $[1,1,0]$ and $[1,y,0]$.

There is a path in S_1 from $[1,1,0]$ to any other vertex $[x,y,0]$ with delay 8. This path consists of the path of delay at most 8 form $[1,1,0]$ to $[1,y,0]$ followed by edges in dimension one to the vertex $[x,y,0]$.

Note that for all x,y, the path from the type 2 vertex $[1,1,0]$ to $[x,y,0]$, as constructed above, includes a type 1 vertex. We can now construct a path of delay at most 8 from $[1,1,0]$ to any other $[x,y,i]$ where $i \equiv 0(\bmod 2)$. To do this, we use the path P_0 from $[1,1,0]$ to $[x,y,0]$, which includes a type 1 vertex w. Our new path consists of the part of P_0 from $[1,1,0]$ to w followed by $\min\{i, n_3 - i\}$ edges in dimension three, followed by the remainder of P_0 translated i units (positively or negatively as required) in dimension three.

It remains to construct paths from $[1,1,0]$ to any $[x,y,i]$ where $i \equiv 1(\bmod 2)$. Such a path will consist of a path from $[1,1,0]$ to $[x,1,i]$ followed by edges in dimension two to $[x,y,i]$. The delay of such path is simply the delay of the subpath from $[1,1,0]$ to $[x,1,i]$. To construct this subpath, we consider three different cases.

If $x \equiv 0 (\bmod 4)$ or $x \equiv 3 (\bmod 4)$, then there is a delay 0 subpath consisting of edges in dimension one from $[1, 1, 0]$ to $[x, 1, 0]$ followed by dimension three edges to $[x, 1, i]$. If $x \equiv 1 (\bmod 4)$, there is a path of delay 8 that starts at $[1, 1, 0]$ proceeds directly to $[0, 1, 0]$ to $[0, 1, 1]$ to $[0, 0, 1]$ to $[0, 0, 0]$, followed by dimension one edges to $[x, 0, 0]$, followed by dimension three edges to $[x, 0, i]$, and then directly to $[x, 1, i]$. If $x \equiv 2 (\bmod 4)$, there is a path of delay at most 8 starting at $[1, 1, 0]$ and proceeding directly to $[2, 1, 0]$ to $[3, 1, 0]$ to $[3, 1, 1]$ to $[3, 2, 1]$ to $[3, 2, 0]$, followed by edges of dimension one to $[x, 2, 0]$, followed by edges of dimension three to $[x, 2, i]$ and proceeding directly to $[x, 1, i]$.

Thus, in any case, the subpath has delay at most 8 and we have completed the analysis for type 2 vertices.

We have now shown that S_1 has delay at most 8. Because S_1 and $S_2$are isomorphic, the analysis for S_1 can easily be modified to show that S_2 also has delay at most 8. Thus, we have shown that $EDS(T_{n_1,n_2,n_3}, 8) \geq 2$ when n_1 and n_2 are congruent to $0 (\bmod 4)$ and n_3 is even.

We can modify our construction for the case when n_1 and n_2 are congruent to $0 (\bmod 4)$ and n_3 is odd. The construction is textually the same, resulting in alternating even and odd layers except for two adjacent even layers 0 and $n_3 - 1$. Each of the paths described above consist of three subpaths: a (possibly empty) path in the source layer and possibly an adjacent layer of opposite parity, a path consisting entirely of dimension three edges from the source layer to the destination layer, and a (possibly empty) path in the destination layer and possibly an adjacent layer.

In these paths, we have used the layers 1 and $i + 1$ as the layers adjacent to layers 0 and i, respectively. However, the delay is unchanged if we were to use layers -1 or $i - 1$ instead. In a torus with an odd number of layers, we can therefore change the first or third subpath to use layers -1 or $i - 1$ should layers 0 and 1 or i and $i + 1$ be the two adjacent even layers.

We have now established that $EDS(T_{n_1,n_2,n_3}, 8) \geq 2$ for when n_1 and n_2 are congruent to $0 (\bmod 4)$. As nothing in our argument was dependent on our assumption that $n_1 \geq n_2 \geq n_3$, we can conclude that $EDS(T_{n_1,n_2,n_3}, 8) \geq 2$ whenever any two of n_1, n_2, and n_3 are congruent to $0 (\bmod 4)$.

2.4 k spanners in d-dimensional tori

In the previous subsections, we have studied spanners in tori of low dimension. In this subsection, we study spanners in tori of arbitrary dimension. Our main result is Theorem 7 which establishes the existence of a set of $k \approx d/5$ edge-disjoint spanners in a d-dimensional torus each having delay that depends only on the number of spanners k and the size of the smallest k dimensions. The proof is modelled on the proof of Theorem 8 in our earlier paper [4]. In particular, it uses the same key lemma which we state as Lemma 6 below. We first require some definitions and intermediate results.

A *d-dominating set* of vertices in graph G is a set $U \subseteq V$ such that every vertex in V is within distance d from some element of U. A *d-domatic coloring* of

G is a vertex coloring of G such that each color class constitutes a d-dominating set of G.

We start with two lemmas concerning domatic colorings and Hamilton decompositions of tori.

Lemma 3 *For any $k \geq 1$, the torus $T_{n_1,n_2,\ldots,n_{2k-1}}$ has a 1-domatic coloring with k colors.*

Proof. We first show that the $(2k-1)$-dimensional hypercube Q_{2k-1} has a 1-domatic coloring with k colors. We will then extend this coloring to the $(2k-1)$-dimensional torus $T_{n_1,n_2,\ldots,n_{2k-1}}$. It is known that the hypercube Q_d has a 1-domatic coloring with d colors when d is a power of 2 [22]. Let k' be the power of 2 that is at least k and less than $2k$. We can construct a 1-domatic coloring of Q_{2k-1} with k' colors from the coloring of $Q_{k'}$ by repeating the coloring of $Q_{k'}$ in each k'-dimensional subcube of Q_{2k-1}. The first k colors of this coloring constitute the desired 1-domatic coloring of Q_{2k-1} with k colors.

To obtain a 1-domatic coloring of $T_{n_1,n_2,\ldots,n_{2k-1}}$ with k colors, give each vertex $[i_1, i_2, \ldots, i_{2k-1}]$ of $T_{n_1,n_2,\ldots,n_{2k-1}}$ the color that was assigned to $[i_1 (\bmod 2), i_2 (\bmod 2), \ldots, i_{2k-1} (\bmod 2)]$ in the k-coloring of Q_{2k-1}.

The following result is a special case of Kötzig's conjecture, proved in [11].

Lemma 4 *For any $k \geq 1$, the torus $T_{n_1,n_2,\ldots,n_k}$ has a decomposition into k Hamilton cycles.*

We next need a lemma that shows that in a high dimension torus, we can obtain a set of edge-disjoint spanners that contains at least one spanner of small delay.

Lemma 5 *For any $k \geq 1$ and $d \geq 3k - 1$, there exists a set of k edge-disjoint spanners $S_1, S_2, \ldots, S_k$ of $T_{n_1,n_2,\ldots,n_d}$ such that S_1 has delay at most*

$$2\lfloor n_{d-k+1}/2 \rfloor + 2\lfloor n_{d-k+2}/2 \rfloor + \cdots + 2\lfloor n_d/2 \rfloor + 2.$$

Proof. We express $T_{n_1,n_2,\ldots,n_d}$ as the product of two graphs

$$H = T_{n_{d-k+1},n_{d-k+2},\ldots,n_d} \text{ and } G = T_{n_1,n_2,\ldots,n_{d-k}}.$$

By Lemma 4, we decompose H into a set of k Hamilton cycles $C_1, C_2, \ldots, C_k$, and arbitrarily choose a distinguished vertex z. Since $d - k \geq 2k - 1$, by Lemma 3, we construct a 1-domatic coloring of G with k colors. Further, we decompose G into a set of $k - 1$ edge-disjoint spanners $D_2, D_3, \ldots, D_k$. This can be done, since $d - k \geq k - 1$, and we can obtain a set of $k - 1$ edge-disjoint Hamilton cycles and dispense the remaining edges arbitrarily.

Each copy of H in $H \times G$ corresponds to a vertex of a particular color i in the domatic coloring of G. In each copy of H, we place the edges of $C_1, C_2, \ldots, C_k$ into spanners $S_1, S_2, \ldots, S_k$, respectively, except for C_1 and C_i. We place the edges of C_1 in S_i and edges of C_i in S_1. In the copy of G corresponding to the distinguished vertex z, we place the edges of $D_2, D_3, \ldots, D_k$ into spanners $S_2, S_3, \ldots, S_k$, respectively. In every other copy of G, we place all edges into S_1.

Consider a spanner S_i, $2 \le i \le k$, and two arbitrary vertices u and v. There is a path from u to v in S_i that starts at u, proceeds within a copy of H (using edges either of C_i or C_1) to a copy of z, then proceeds within a copy of G (using edges of $D_2, D_3, \ldots, D_k$) to another copy of z, and then proceeds within a copy of H (using edges either of C_i or C_1) to v. Thus, S_i is connected and, therefore, a spanner.

Now consider a pair of vertices u and v in S_1. If u is a copy of z, let u' be a vertex adjacent to u in S_1 and otherwise let $u' = u$. Similarly, if v is a copy of z, let v' be a vertex adjacent to v in S_1, and otherwise, let $v' = v$. Let $u' = [h_1, g_1]$ and $v' = [h_2, g_2]$. Let $u'' = [h_2, g_1]$, that is u'' is a copy of u' in the copy of H containing v'. We will construct a path from u to v that commences at u, and passes through u', u'', and v' in order and then arrives at v. The subpaths from u to u' and from v' to v are straightforward. The subpath from u'' to v' requires further elucidation. Let H' be the copy of H containing u'' and v', and let z' be the copy of z in H'. Note that the distance from u'' to v' in H' is at most $\ell = \lfloor n_{d-k+1}/2 \rfloor + \lfloor n_{d-k+2}/2 \rfloor + \cdots + \lfloor n_d/2 \rfloor$.

Now consider a shortest path from u'' to v' in H'; since either u'' and v' are adjacent or there are at least two such vertex-disjoint paths, there is a shortest path P that does not contain z'. We construct a path from u'' to v' in S_1 by replacing each edge of P that is not already in S_1 by a path of three edges in S_1. Let $e = (x, y)$ be an edge of P not in S_1. In H', e belongs to some cycle C_j. In the domatic coloring of G, there is a vertex of color j adjacent to the vertex corresponding to H'. Let x' and y' be the vertices corresponding to x and y, respectively, in the copy of H corresponding to this vertex of color j. By construction, (x', y') is in S_1 and since neither x nor y is z', both (x, x') and (y, y') are edges in S_1. We use the path (x, x', y', y) to replace the edge (x, y) in the path P.

If u and v are both copies of z, then the vertices u' and v' can be chosen so that they are corresponding to each other, so the distance from u'' to v' is 0. Moreover, the distance in S_1 from u to u' is 1, the distance in S_1 from u' to u'' is $d_{H\times G}(u, v)$, and the distance in S_1 from v' to v is 1. Thus, in this case, the delay is 2.

If exactly one of the vertices u and v, without loss of generality v, is a copy of z, then the distance in S_1 from u to u' is 0, the distance in S_1 from u' to u'' is $d_G(u, v)$, and the distance in S_1 from v' to v is 1. The path from u'' to v' in S_1 has length at most $3d_H(u'', v')$, which is at most $d_H(u'', v') + 2\ell$. Finally, this quantity is at most $d_H(u, v) + 2\ell + 1$ in this case. This gives a total distance from u to v of at most $d_{H\times G}(u, v) + 2\ell + 2$ and a delay of at most $2\ell + 2$.

If neither u nor v are copies of z, then the distance in S_1 from u to u' is 0, the distance in S_1 from u' to u'' is $d_G(u, v)$, the distance in S_1 from u'' to v' is at most $d_H(u'', v') + 2\ell$ as in the previous case, and the distance in S_1 from v' to v is 0. Since, in this case, $d_H(u'', v') = d_H(u, v)$, this gives a total distance from u to v of at most $d_{H\times G}(u, v) + 2\ell$ and a delay of at most 2ℓ.

Finally, we will use the following lemma (which is Corollary 2 from [4]) to combine the previous lemmas. We use the notation $\log k$ to denote the base-2 logarithm of x.

Lemma 6 *[4] Let H be a graph with k edge-disjoint spanners such that H_1 is a delay c spanner. Let G be a graph with an r-domatic coloring with k colors. Then $EDS(H \times G, 2r + 48k \log k + c) \geq k$.*

We are now ready to state our result.

Theorem 7 *For any $k \geq 1$ and $d \geq 5k - 2$,*

$$EDS(T_{n_1,n_2,\ldots,n_d}, 48k \log k + 2\lfloor n_{d-k+1}/2 \rfloor + 2\lfloor n_{d-k+2}/2 \rfloor + \cdots + 2\lfloor n_d/2 \rfloor + 4) \geq k.$$

Proof. We express $T = T_{n_1,n_2,\ldots,n_d}$ as $T = H \times G$, where

$$G = T_{n_1,n_2,\ldots,n_{2k-1}} \text{ and } H = T_{n_{2k},n_{2k+1},\ldots,n_d}.$$

By Lemma 3, the graph G has a 1-domatic coloring with k colors, and since $d - 2k + 1 \geq 3k - 1$, by Lemma 5, the graph H has k factors from which one has delay at most $2\lfloor n_{d-k+1}/2 \rfloor + 2\lfloor n_{d-k+2}/2 \rfloor + \cdots + 2\lfloor n_d/2 \rfloor + 2$. Finally, we apply Lemma 6.

References

1. Althöfer, I., G. Das, D. Dobkin, D. Joseph, and J. Soares, On sparse spanners of weighted graphs, *Discr. & Comp. Geometry*, **9**, 1993, 81–100.
2. Awerbuch, B., A. Baratz, and D. Peleg, Efficient Broadcast and Light-Weight Spanners, Technical Report CS92-22, the Weizmann Institute of Science, Rehovot, Israel, 1992.
3. Awerbuch, B. and D. Peleg, Sparse partitions, *31st IEEE Symp. on Foundations of Computer Science*, 503–513, October 1990.
4. Fertin, G., Liestman, A. L., Shermer, T. C., and Stacho, L., Edge-disjoint spanners in Cartesian products of graphs, submitted.
5. Chandra, B., G. Das, G. Narasimhan, and J. Soares, New sparseness results on graph spanners, *Proc. 8th ACM Symp. on Computational Geometry*, 1992.
6. Heydemann, M.-C., J. Peters, and D. Sotteau, Spanners of Hypercube-derived Networks, *SIAM J. Discr. Math.* **9**, 1996, 37–54.
7. Kortsarz, G. and D. Peleg, Generating Sparse 2-Spanners, *J. Algorithms*, **17**, 1994, 222–236.
8. Kortsarz G. and D. Peleg, Generating Low-Degree 2-Spanners, *SIAM J. on Computing*, **27**, 1998, 1438–1456.
9. Laforest, C., A. L. Liestman, T. C. Shermer, and D. Sotteau, Edge-disjoint spanners of complete bipartite graphs, *Discrete Mathematics* **234**, 2001, 65–76.
10. Laforest, C., A. L. Liestman, D. Peleg, T. C. Shermer, and D. Sotteau, Edge Disjoint Spanners of Complete Graphs and Complete Digraphs, *Discr. Math.* **234**, 2001, 65–76.
11. Lian, G. C., Hamiltonian decomposition of the Cartesian product of n graphs–proof of conjecture of A. Kötzig, Acta Math. Appl. Sinica **11**, 1998, 123–126.
12. Liestman, A. L. and T. C. Shermer, Additive Graph Spanners, *Networks* **23**, 1993, 343–364.
13. Liestman, A. L. and T. C. Shermer, Additive Spanners for Hypercubes, *Par. Proc. Lett.* **1**, 1991, 35–42.
14. Liestman, A. L. and T. C. Shermer, Degree-Constrained Network Spanners With Non-constant Delay, *SIAM J. Discr. Math.* **8**, 1995, 291–321.
15. Liestman, A. L. and T. C. Shermer, Grid Spanners, *Networks* **23**, 1993, 123–133.
16. Liestman, A. L., T. C. Shermer, and C. R. Stolte, Degree-Constrained Spanners for Multi-dimensional Grids, *Discr. Appl. Math.* **68**, 1996, 119–144.

17. Mansour, Y. and D. Peleg, An Approximation Algorithm for Minimum-Cost Network Design, Technical Report CS94-22, the Weizmann Institute of Science, 1994.
18. Peleg, D. and A.A. Schäffer, Graph Spanners, *J. Graph Theory* **13**, 1989, 99–116.
19. Peleg, D. and J. D. Ullman, An Optimal Synchronizer for the Hypercube, *SIAM J. Computing* **18**, 1989, 740–747.
20. Richards, D. and A. L. Liestman, Degree-Constrained Pyramid Spanners, *J. Par. Distr. Comp.* **25**, 1995, 1–6.
21. West, D.B., Introdusction to Graph Theory - Second Edition, Prentice Hall, 1996, 2001, ISBN 0-13-014400-2.
22. Zelinka, B., Nektere ciselne invarianty grafu (Some numerical invariants of graphs), Ph.D. Dissertation, Prague, 1988.