Randomly Twisted Hypercubes

Andrzej Dudek[*]   Xavier Pérez-Giménez[†]   Paweł Prałat[‡]   Hao Qi[§]

Douglas West[¶]   Xuding Zhu[‖]

January 29, 2018

Abstract

A *twisted hypercube* of dimension k is created from two twisted hypercubes of dimension $k-1$ by adding a matching joining their vertex sets, with the twisted hypercube of dimension 0 consisting of one vertex and no edges. We generate random twisted hypercube by generating the matchings randomly at each step. We show that, asymptotically almost surely, joining any two vertices in a random twisted hypercube of dimension k there are k internally disjoint paths of length at most $\frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right)$. Since the graph is k-regular with 2^k vertices, the number of paths is optimal and the length is asymptotically optimal.

1   Introduction

The *hypercube* Q_k is a fundamental and thoroughly studied graph and a useful interconnection network. It has vertex set $\{0,1\}^k$, with vertices adjacent when they differ in exactly one coordinate. The graph can also be formed inductively by letting Q_0 be a single vertex and creating Q_k from two copies of Q_{k-1} by adding a matching joining corresponding vertices in the two copies.

Among the most appealing features of the hypercube are its good expansion properties and small diameter. A u,v-*path* is a path with endpoints u and v. The *distance* between u and v, denoted $d(u,v)$, is the minimum length (number of edges) of a u,v-path. The *diameter* of G, denoted $d(G)$, is $\max_{u,v \in V(G)} d(u,v)$.

The diameter of the hypercube Q_k is k, although the number of vertices is 2^k. Furthermore, this measure of maximum distance is quite robust. Saad and Schultz [14] proved that joining any two vertices in Q_k there are at least $k-1$ internally disjoint paths of length at most k.

Nevertheless, one can seek even stronger results for appropriately modified graphs. For any k-regular graph with 2^k vertices, the diameter is bounded from below using a counting argument: there are at most $k(k-1)^{d-1}$ vertices at distance d from a given vertex. Hence the diameter is at

1

least $\frac{k}{\lg k}$, and the hypercube exceeds this by a factor of $\lg k$, where $\lg k = \log_2 k$. The expansion properties of Q_k are limited by the prevalence of small cycles, especially 4-cycles, which reduce the number of vertices reachable within a given distance.

To overcome this difficulty, various researchers proposed more flexible constructions. The terms "twisted hypercube" or "crossed hypercube" appear in dozens of papers, often studying restricted versions of the matchings allowed in the general recursive model we study. Often the aim was to reduce the diameter while maintaining k-regularity and vertex set $\{0,1\}^k$, although connectivity and other properties of paths and cycles have also been studied.

Definition 1. For $k \geq 1$, a *twisted hypercube* of dimension k is obtained from two twisted hypercubes of dimension $k-1$ by adding a matching joining the vertex sets of the two smaller graphs. The unique twisted hypercube of dimension 0 consists of one vertex and no edges.

A k-dimensional twisted hypercube is k-regular with 2^k vertices, so its diameter is at least $\frac{k}{\lg k}$. Esfahanian, Ni, and Sagan [4] twisted a single 4-cycle, thereby reducing the diameter to $k-1$. Using cartesian products of 3-dimensional twisted hypercubes, Chedid and Chedid [2] noted that diameter $2k/3$ can be achieved. A model proposed by Efe [3], an explicit inductive restricted version of the model we study, reduced the diameter to $\lceil (k+1)/2 \rceil$. To our knowledge, the earliest suggestion of the general recursive twisting in Definition 1 was in an unpublished private comment by Peter Slater in 1997. In response to that comment, Hartman [6] found an explicit inductive construction of a k-dimensional twisted hypercube with diameter at most $c\frac{k}{\lg k}$, were c can be any number greater than 8. Recently, Zhu [20] presented a simple inductive construction of twisted hypercubes: to obtain the k-dimensional twisted hypercube H_k, take two copies $0H_{k-1}$ and $1H_{k-1}$ of H_{k-1}, let $\kappa = \lceil \lg(k-1) - 2\lg^2(k-1) \rceil$, add a matching edge connecting $0a_1a_2 \ldots a_{k-1}$ and $1b_1b_2 \ldots b_{k-1}$ if for $1 \leq i < \kappa$, $b_i = a_i + a_{k-\kappa+i}$ (mod 2) and for $\kappa \leq i < k$, $b_i = a_i$. It was proved in [20] that H_k has diameter asymptotic to $\frac{k}{\lg k}$.

In addition to decreasing the diameter, one can ask for robustness. Fault-tolerance for diameter is captured by the notion of "wide diameter". A graph is k-*connected* if the deletion of fewer than k vertices cannot result in a disconnected graph or a graph with only one vertex. By Menger's Theorem [12], in a k-connected network there exist k internally disjoint paths joining any two vertices, where paths with the same endpoints are *internally disjoint* if they share no other vertices. Therefore, when G is ℓ-connected, we may define the ℓ-*wide-distance* between u and v, denoted $d_\ell(u,v)$, to be the least integer d such that G contains ℓ internally vertex-disjoint u,v-paths of length at most d. The ℓ-*wide-diameter* of G, denoted $d_\ell(G)$, is $\max_{u,v \in V(G)} d_\ell(u,v)$. Note that

$$d(G) = d_1(G) \leq d_2(G) \leq \ldots \leq d_\ell(G).$$

The wide diameter was introduced independently by Hsu and Lyuu [7] and by Flandrin and Li [5]. For general graphs, Hsu and Łuczak [8] showed that every k-regular k-connected graph on n vertices has k-wide-diameter at most $n/2$. Wide diameter was studied for various Cartesian product graphs and hypergraph variations in [1, 11, 15, 16, 18, 19].

When the ℓ-wide-diameter of a graph equals its diameter, the graph can be said to be $(\ell - 1)$-fault-tolerant for diameter. The result of Saad and Schultz [14] mentioned earlier shows that the hypercube Q_k is $(k-2)$-fault-tolerant for diameter. Their full result on the wide-diameter of Q_k is

$$d_\ell(Q_k) = \begin{cases} k & \text{if } 1 \leq \ell \leq k-1, \\ k+1 & \text{if } \ell = k. \end{cases}$$

Qi and Zhu [13] showed that for the k-dimensional twisted hypercubes constructed in [20], for $\ell \leq \lg k$, its ℓ-wide diameter is asymptotically $\frac{k}{\lg k}$. Our objective in this paper is to extend this

statement to all $\ell \leq k$. In particular, we show that both the reduction of diameter asymptotically to $\frac{k}{\lg k}$ and the stronger property of having ℓ-wide-diameter asymptotic to diameter for $\ell \leq k$ are achieved by asymptotically almost all twisted hypercubes when generated by choosing random matchings.

Definition 2. A *random k-dimensional twisted hypercube*, denoted $\hat{Q}_k$, is formed as follows. Let $\hat{Q}_0$ consist of a single vertex and no edges. For $k \geq 1$, form $\hat{Q}_k$ from two disjoint independently generated random $(k-1)$-dimensional twisted hypercubes $\hat{Q}_{k-1}$ and $\hat{Q}'_{k-1}$ by adding a random matching joining their vertex sets. In the process of forming $\hat{Q}_k$, at step i there are 2^{k-i} independently generated random matchings used to form connected subgraphs with 2^i vertices; we call these *matchings of type i*.

Let us stress again that $\hat{Q}_{k-1}$ and $\hat{Q}'_{k-1}$ are generated independently and need not be isomorphic; in fact, they are isomorphic with only a small probability when k is large. The set of edges of $\hat{Q}_k$ contains for $1 \leq i \leq k$ a union of 2^{k-i} (independent) random matchings of type i. We still treat $V(\hat{Q}_k)$ as $\{0,1\}^k$. Thus for $a_1 \cdots a_{k-i} \in \{0,1\}^{k-i}$, there is a matching of type i that matches the set $a_1 a_2 \cdots a_{k-i} 0 \times \{0,1\}^{i-1}$ to the set $a_1 a_2 \cdots a_{k-i} 1 \times \{0,1\}^{i-1}$.

As is usual when discussing randomized constructions, our results are asymptotic. That is, we always consider $k \to \infty$, so we may always assume that k is large enough to guarantee the truth of statements that depend on k being large enough. The notation $o(f(k))$ indicates a function of k whose ratio to $f(k)$ tends to 0. The notation $O(f(k))$ indicates a function of k whose magnitude is bounded by a multiple of $f(k)$ when k is sufficiently large. We also write $f(k) \sim g(k)$ if $f(k)/g(k) \to 1$ as $k \to \infty$, which is equivalent to $f(k) = (1 + o(1))g(k)$.

We say that a property or event A_k holds *asymptotically almost surely* (or *a.a.s.*) if the probability of A_k tends to 1 as k goes to infinity. We write $\hat{Q}_k$ to indicate a graph generated by the model specified in Definition 2. Finally, we denote the logarithm of x with base 2 by $\lg x$. Our main result is the following.

Theorem 3. *Asymptotically almost surely,* $d_k(\hat{Q}_k) = \frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right) \sim \frac{k}{\lg k}$.

We have noted that the diameter is always at least $\frac{k}{\lg k}$ for a k-regular graph with 2^k vertices. Since $d_1(G) \leq \cdots \leq d_k(G)$, the result implies that a.a.s. for $1 \leq \ell \leq k$ we have $d_\ell(\hat{Q}_k) \sim \frac{k}{\lg k}$, which is (asymptotically) sharp.

We also remark that the a.a.s. claim in Theorem 3 also holds with probability $1 - o(e^{-\alpha k})$, where α is any positive constant. This can be easily deduced from Remark 8 in Section 3 and some straightforward modifications of the argument. However, we do not attempt to optimize this probability bound.

2   The Pairing Model

Twisted hypercubes are regular graphs. Analysis of random twisted hypercubes is greatly aided by modifying a model that has proved to be highly successful in analyzing random regular graphs. This is the so-called *pairing model* (also known as the *configuration model*).

We view each vertex of $\hat{Q}_k$ as a set of k *points*. The points in a vertex are assigned distinct labels (types) from $1, \ldots, k$, corresponding to the k matchings that form $\hat{Q}_k$. In a matching of type i, we are restricted for each $a \in \{0,1\}^{k-i}$ to match the points of type i in the vertices of $a_1 a \cdots a_{k-i} 0 \times \{0,1\}^{i-1}$ with the points of type i in the vertices of $a_1 \cdots a_{k-i} 1 \times \{0,1\}^{i-1}$ (which we call the *permissible* points), and we do so at random. This can be done in many different ways,

3

some of which turn out to be very convenient. In particular, the edges can be chosen sequentially. As matchings are independent, one can generate part of one matching and move to generating part of another one, if needed. The graph $\hat{Q}_k$ is then obtained by merging the k points in each vertex; each point contributes one incident edge.

In the usual pairing model, all points in all vertices are available as neighbors, and hence loops and multiple edges may arise. In the restricted version we have specified here, the resulting graph is always simple. Although odd cycles may occur as soon as dimension 3 is reached, always $\hat{Q}_k$ has girth at least 4.

The most important advantage of the pairing model is the ability to view the edges as generated sequentially. A matching between sets S_1 and S_2 of points can be selected uniformly at random in many ways. In particular, at any stage the first point in the next random pair chosen can be selected using any rule whatsoever, as long as the second point in that pair is chosen uniformly at random from the remaining unmatched points in the opposite set. For example, one can insist that the next point chosen is the next one remaining in any pre-specified ordering of the points in $S_1 \cup S_2$. Alternatively, it can be restricted to come from a vertex containing one of the points in the previous pair chosen (if any such points remain available).

We use this idea several times. For example, consider the process of generating $\hat{Q}_4$. We start with 2^4 vertices, where each vertex consists of four points. No edge is generated yet and our goal is to explore the graph around an arbitrarily chosen vertex v containing points a, b, c, d (see Figure 1). Point a must be matched to e in the matching of type 1. Since points must be matched to permissible points of the same type, point b is matched randomly to one of the two f points, c is matched to one of the four g points, and d to one of the eight h points. The neighborhood of v has been explored. After that, we may want to continue exploring the graph from one of the neighbors of v or start from a new vertex, depending on the application. For more details about the model, see the survey by Wormald [17] on random regular graphs, for example.


Figure 1: Exploring neighborhoods using the pairing model

We will use a well-known bound on tail probabilities.

Theorem 4 (**Chernoff Bound** (see [9], Theorem 2.1)). *Let X be distributed as a binomial random variable with n trials and success probability p, so $\mathbb{E}[X] = \mu = pn$. If $0 < \delta < 1$, then*

$$\mathbb{P}[X < (1-\delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2}\right). \tag{1}$$

If $\delta > 0$, then

$$\mathbb{P}[X > (1+\delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2 + \delta}\right). \tag{2}$$

3 Further Tools

The main task in this section is to show that a.a.s. $\hat{Q}_{k-1}$ has good expansion properties. In particular, we want to show that a.a.s. from every vertex of Q_{k-1} we can reach many vertices in a relatively short distance. Specifically, we want to reach $k2^{k/2}$ vertices within distance $\frac{k}{2\lg k} + O\left(\frac{k}{\lg^2 k}\right)$. This is asymptotically best possible within $\hat{Q}_{k-1}$, since Q_{k-1} is $(k-1)$-regular, and hence the number of vertices within distance d of a given vertex is at most $1 + (k-1)\sum_{j=0}^{d-1}(k-2)^j$. The sum is $(1 + o(1))k^d$. Thus to reach $k2^{k/2}$ vertices we need $d > \frac{k}{2\lg k}$.

Definition 5. For $v \in V(\hat{Q}_{k-1})$, say that a set $V' \subseteq V(\hat{Q}_{k-1})$ is *v-proper* if $v \in V'$, each vertex of V' has distance at most $\frac{k}{2\lg k} + O\left(\frac{k}{\lg^2 k}\right)$ from v in the subgraph of $\hat{Q}_{k-1}$ induced by V', and $|V'| \geq 2^{k/2}k$.

Lemma 6. *Given $h \leq 2k$, let $v_1, \ldots, v_h$ be distinct vertices in $\hat{Q}_{k-1}$. With probability $1 - o(2^{-2k})$, there exist h pairwise disjoint vertex sets $V_1, \ldots, V_h$ such that V_j is v_j-proper, for $1 \leq j \leq h$.*

Proof. Fixing k, let r_0 be the largest integer r such that $k^r \leq 2^{0.1k}$, and let r_1 be the smallest integer r such that $(k/4)^r/4 \geq 2^{k/2}k$. Clearly,

$$r_0 = \frac{0.1k}{\lg k} + O(1) \qquad \text{and} \qquad r_1 = \frac{k}{2\lg k} + O\left(\frac{k}{\lg^2 k}\right) \sim \frac{k}{2\lg k}.$$

For technical reasons, we consider a subgraph of $\hat{Q}_{k-1}$, denoted by $\overline{Q}_{k-1}$, that consists of the matchings of type more than $0.51k$. In order to analyze this subgraph, we modify the pairing model so that instead of $k-1$ points associated with each vertex there are q points, where $q = k-1-\lfloor 0.51k \rfloor$. Note that $0.49k-1 \leq q \leq 0.49k$. Since $\overline{Q}_{k-1} \subseteq \hat{Q}_{k-1}$, the claim follows immediately from the same claim for $\overline{Q}_{k-1}$. Therefore, we restrict our attention to $\overline{Q}_{k-1}$.

Let $U = \{v_1, \ldots, v_h\}$, and consider a breadth-first search (BFS) in $\overline{Q}_{k-1}$ starting from the set U. The search builds a forest with components $T_1, \ldots, T_h$ such that v_j lies in T_j for all j. Since we are using the pairing model to generate $\overline{Q}_{k-1}$, we describe the process in terms of points rather than vertices. (See Section 2 for details on the pairing model used.) In each step of the process, we partition the set of points into three sets: points already matched (*saturated*), unmatched points associated with vertices containing some point reached by the BFS (*discovered*), and the remaining points (*undiscovered*).

Initially, all the points associated with vertices in U are labeled discovered and added to a queue L; the remaining points are undiscovered. At each step, the next discovered point p is removed from L and matched to a point p' of the same type as p that is taken uniformly at random from the set of currently unsaturated points among those permitted to be matched to p. These points may be discovered or undiscovered; recall that the ith point associated with each vertex is used only in a matching of type i. If p' is undiscovered (we reach a new vertex), then the remaining points associated with the vertex containing p' change from undiscovered to discovered and are added to the end of L. Otherwise p' was discovered earlier, and all points associated with the vertex containing p' are already discovered or saturated. In this case we call the edge pp' a *bad* edge. Finally, both p and p' are labeled as saturated.

For $0 \leq i \leq r_1$, let $F(i)$ be the forest obtained from the BFS once the balls of radius i around U are exposed; this is the graph consisting of all edges reaching vertices at distance at most i in $\overline{Q}_{k-1}$ from some vertex of U, with all bad edges removed. For $1 \leq j \leq h$ and $0 \leq i \leq r_1$, let $V_j(i)$ be the set of vertices at distance i from v_j in $F(i)$, and let $S_j(i)$ be the set of points associated with vertices in $V_j(i)$ that are labeled discovered when all vertices in $V_j(i)$ are reached by the BFS. In particular, $V_j(0) = \{v_j\}$ and $|S_j(0)| = q$. Since bad edges are removed during the construction of $F(i)$, the vertices of U must lie in distinct components. Therefore, the sets $V_1(i), \ldots, V_h(i)$ are pairwise disjoint, and the same holds for $S_1(i), \ldots, S_h(i)$. Let $V_j = \bigcup_{i=0}^{r_1} V_j(i)$. In order to prove the claim, we only need to show that with probability $1 - o(2^{-2k})$ each V_j has size at least $2^{k/2}k$, since the other specified requirements on V_j hold by construction.

We analyse $F(i)$ in two phases. During the first, we obtain the desired expansion up to level r_0. We compute an upper bound on the number m of edges encountered during this phase:

$$m \leq hq \sum_{i=0}^{r_0-1} (q-1)^i \leq 2kq^{r_0} \leq 2k(0.49k)^{r_0} \leq 2k2^{0.1k} \leq 2^{0.11k}.$$

Note that $q \sum_{r=0}^{r_0-1}(q-1)^r$ is the number of edges in a rooted q-ary tree of depth r_0, which corresponds to the optimal situation with no bad edge present. For $0 \leq \ell \leq m$, we wish to bound the probability p_ℓ that the next edge after exposing ℓ edges is bad. Recall that there are initially hq discovered points, and each exposed edge discovers at most $q - 1$ new points. Moreover, since we only use matchings of type more than $0.51k$, each matching has at least $2^{0.51k-1}$ points. Therefore,

$$p_\ell \leq \frac{hq + \ell(q-1)}{2^{0.51k-1} - \ell} \leq \frac{(h+\ell)\,0.49k}{2^{0.51k-1} - \ell} \leq \frac{(h+\ell)k}{2^{0.51k}},$$

since $\ell \leq 2^{0.11k} = o(2^{0.51k-1})$. Since the bound on p_ℓ increases with ℓ, we can apply it for each step among the first m edges. Therefore, to bound the probability p of having at least c bad edges among the first m edges, we can use the Union Bound over the possible subsets of size c to compute

$$p \leq \binom{m}{c} \left(\frac{(h+m)k}{2^{0.51k}} \right)^c \leq \left(\frac{e(h+m)mk}{c2^{0.51k}} \right)^c \leq \left(\frac{2ek}{c2^{0.29k}} \right)^c,$$

since $m \leq 2^{0.11k}$. We choose $c = 7$ so that $0.29c > 2$, and then the probability of having at least c bad edges during this phase is $o(2^{-2k})$. The worst scenario for expansion from $S_j(i)$ to $S_j(i+1)$ is when c bad edges are within $S_j(i)$; the remaining points generate $q-1$ points that fall into $S_j(i+1)$. Thus, for $1 \leq j \leq h$,

$$
\begin{aligned}
|S_j(0)| &= q \geq 0.49k - 1 \geq k/4 &\qquad(3)\\
|S_j(1)| &\geq (|S_j(0)| - 2c) \cdot (q-1) \geq (|S_j(0)| - 2c) \cdot (0.49k - 2) \geq (k/4)^2\\
&\cdots \geq \cdots\\
|S_j(r_0)| &\geq (|S_j(r_0-1)| - 2c) \cdot (q-1) \geq (|S_j(r_0-1)| - 2c) \cdot (0.49k - 2) \geq (k/4)^{r_0+1}.
\end{aligned}
$$

During the second phase, we control the expansion continuing from level r_0 up to level r_1. It follows from the definition of r_0 that $k^{r_0+1} \geq 2^{0.1k}$. Therefore,

$$(k/4)^{r_0+1} \geq 2^{0.1k - O(k/\lg k)} \geq 2^{0.05k}.$$

For $1 \leq j \leq h$ and $r_0 \leq i \leq r_1 - 1$, conditioned on the event $|S_j(i)| \geq (k/4)^{i+1} \geq 2^{0.05k}$, we want to have $|S_j(i+1)| \geq |S_j(i)|(k/4)$ with probability $1 - o(1/(2^{2k}k^2))$. This will imply $|S_j(r_1)| \geq (k/4)^{r_1+1}$, and thus

$$|V_j(r_1)| \geq |S_j(r_1)|/k \geq (k/4)^{r_1}/4 \geq 2^{k/2}k$$

6

with probability at least $1 - o(1/(2^{2k}k))$. Since $V_j = \bigcup_{i=0}^{r_1} V_j(i) \supseteq V_j(r_1)$, the result then follows from applying the Union Bound over all h choices of j.

Fix i and j, and assume $|S_j(i)| \geq 2^{0.05k}$. Let X be the random variable counting the bad edges created when the edges of the matching for points in $S_j(i)$ are exposed. The number of points discovered when BFS is performed up to level $i+1$ (from all h initial vertices) is at most hk^{i+1}, which is bounded by $2k^{i+2}$. The probability of generating a bad edge changes at each step of the BFS, but for all edges added when moving from level i to level $i+1$ this probability is bounded by p_i', where

$$p_i' \leq \frac{2k^{i+2}}{2^{0.51k-1} - 2k^{i+2}} = O\left(\frac{k^{i+2}}{2^{0.51k}}\right).$$

As a result, X is stochastically bounded above by $\mathrm{Bin}(|S_j(i)|, p_i')$. Thus

$$\mathbb{E}[X] \leq |S_j(i)| \cdot p_i' = O\left(|S_j(i)| \cdot \frac{k^{i+2}}{2^{0.51k}}\right) = o\big(|S_j(i)|\big).$$

It therefore follows from the Chernoff Bound that

$$\mathbb{P}\big(X \geq |S_j(i)|/100\big) \leq \exp\big(-\Omega(|S_j(i)|)\big) = o(1/(2^{2k}k^2)).$$

Therefore, with the desired probability

$$|S_j(i+1)| \geq (|S_j(i)| - 2X)(q-1) \geq |S_j(i)|(49/50)(0.49k - 2) \geq |S_j(i)|(k/4). \tag{4}$$

This completes the proof of the lemma. $\qquad\square$

In order to prove our main result, we slightly strengthen the previous lemma by allowing a constant number of vertices to be removed from the graph. The proof can be easily adapted from the proof of Lemma 6, so we will only sketch the main differences. Before we state the corollary, let us say a few words about the conditional probability space we work with. Given a set of vertices W, we reveal the neighbors of W by exposing the appropriate matchings. Then we fix any possible outcome of that event (even if it is highly unlikely one), restrict the probability space to that situation, and continue the process for a given $U \subseteq V(\hat{Q}_{k-1}) - W$.

Corollary 7. *Let W be a set of m vertices in $\hat{Q}_{k-1}$, where m is fixed and does not depend on k. Expose all edges with endpoints in W, and condition the probability space on these edges. Let U be a set of distinct vertices $\{v_1, \ldots, v_h\}$ in $\hat{Q}_{k-1} - W$, where $1 \leq h \leq 2k$. With probability $1 - o(2^{-2k})$, there exist pairwise disjoint sets $V_1, \ldots, V_h$ in $V(\hat{Q}_{k-1}) - W$ such that V_j is v_j-proper, for $1 \leq j \leq h$.*

Proof (sketch). First expose the neighbors of the vertices in W with respect to $\overline{Q}_{k-1}$, saturating the corresponding points (for each edge we saturate two points, one for each endvertex). Next run the same BFS process as analyzed in the proof of Lemma 6. The only difference is that we may have some additional saturated points along the way (possibly in U). However, each vertex of $\overline{Q}_{k-1} - W$ has at most m saturated points, one for each vertex in W. Therefore, the same analysis is valid, replacing the factor $q-1$ in (3) and (4) with $q-1-m$, which is asymptotic to $0.49k$ (recall that m is fixed and $q = k - 1 - \lfloor 0.51k \rfloor$). $\qquad\square$

Remark 8. *Note that the statements of Lemma 6 and Corollary 7 are still valid if we replace the probability bound $1 - o(2^{-2k})$ by $1 - o(e^{-\alpha k})$, where α is any positive constant. To show this, the only modification required in the argument of Lemma 6 is to pick c to be large enough rather than setting $c = 7$.*

We will also need the following property of each possible $\hat{Q}_k$.

Lemma 9. *If $k \geq 3$, then $\hat{Q}_k$ has no two vertices with identical neighborhoods.*

Proof. Let A and B denote the two instances of $\hat{Q}_{k-1}$ joined by the last matching (the one of type k) to create $\hat{Q}_k$. We may let A be the graph induced by all vertices whose last coordinate is 0, while B is induced by those with last coordinate 1. We call A and B the *sides* of $\hat{Q}_k$. Note that A and B are *any* graphs that may arise during the random process.

By construction, every vertex has exactly one neighbor in the opposite side and $k - 1$ neighbors in its own side. If distinct vertices u and v lie in opposite sides, say $u \in A$ and $v \in B$, then u has more neighbors than v in A, so u and v cannot have the same neighborhoods. If u and v lie in the same side, say A, then they have distinct neighbors in B via the last matching. $\square$

4   Proof of the Main Result

Now, we are ready to prove our main result. We repeat the statement for convenience.

Theorem 3. *A.a.s. $d_k(\hat{Q}_k) = \frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right) \sim \frac{k}{\lg k}$.*

Proof. Let A and B be the sides of $\hat{Q}_k$, as defined in the proof of Lemma 9. We must show that any two vertices u and v are joined by k internally disjoint paths of length at most $\frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right)$. We consider several cases. In all cases we may assume by symmetry that $u \in A$.

In each case, we will grow trees in the two sides of $\hat{Q}_k$ and add edges to link paths in pairs of trees. Growing trees in the lemma and later adding at most three edges for each pair of trees are separate processes; the first is recursive with no linking step, while the linking is done once at the end. Hence the added 3 remains an additive constant to the length and does not generate a linear term.

Case 1: *Vertices u and v lie in opposite sides of $\hat{Q}_k$ and are adjacent.* Since $u \in A$ and $v \in B$, they are adjacent only by the last matching. The edge uv provides one u, v-path, so we only need to find $k - 1$ more. Let $u_1, \ldots, u_{k-1}$ be the neighbors of u in A, and let $v_1, \ldots, v_{k-1}$ be the neighbors of v in B. Applying Corollary 7 to side A (with $U = \{u_1, \ldots, u_{k-1}\}$ and $W = \{u\}$) and to side B (with $U = \{v_1, \ldots, v_{k-1}\}$ and $W = \{v\}$), with probability $1 - o(2^{-2k})$ there are pairwise disjoint sets $U_1, \ldots, U_{k-1} \subseteq A$ not containing u and $V_1, \ldots, V_{k-1} \subseteq B$ not containing v such U_j is u_j-proper and V_j is v_j-proper, for $1 \leq j \leq k - 1$. See Figure 2.


Figure 2: Case 1: vertices u and v lie in opposite sides of $\hat{Q}_k$ and are adjacent.

Now we consider the last matching, joining A and B. For $1 \leq j \leq k-1$, the probability p that no vertex in U_j is matched to a vertex in V_j is bounded using

$$p \leq \prod_{i=1}^{2^{k/2}k} \left(1 - \frac{2^{k/2}k}{2^{k-1}-i+1}\right) \leq \left(1 - \frac{2^{k/2}k}{2^{k-1}}\right)^{2^{k/2}k} \leq \exp\left(-2k^2\right) = o(2^{-2k}/k).$$

Taking the Union Bound over all choices of j, we have matched U_j to V_j for all j with probability $1 - o(2^{-2k})$. Therefore, with at least that probability, we have k internally disjoint u,v-paths of length at most $\frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right)$.

Case 2: *Vertices u and v lie in opposite sides of $\hat{Q}_k$ and have no common neighbor.* That is, after exposing all matching edges at u and v, it turns out that u and v are non-adjacent and have no common neighbors; no other points are exposed yet. Let $u_1, \ldots, u_{k-1}$ be the neighbors of u in A, and let $v_1, \ldots, v_{k-1}$ be the neighbors of v in B. Let v_k be the neighbor of u in B, and let u_k the neighbor of v in A. The rest of the argument follows as in Case 1, using $1 \leq j \leq k$ instead of $1 \leq j \leq k-1$. See Figure 3.


Figure 3: Case 2: vertices u and v lie in opposite sides of $\hat{Q}_k$ and have no common neighbor.

Case 3: *Vertices u and v lie in opposite sides of $\hat{Q}_k$ and are nonadjacent but have a common neighbor.* All edges at u and v are exposed, and a common neighbor is found. The number of common neighbors is 1 or 2, since there can be at most one in each side. The case with a common neighbor in each side is easy. Define u_j and v_j as in Case 1, and let $u_{k-1} \in A$ and $v_{k-1} \in B$ be the common neighbors of u and v. The u,v-paths via u_{k-1} and v_{k-1} are sufficiently short. To find the remaining paths, proceed as before with $U = \{u_1, \ldots, u_{k-2}\}$ and $W = \{u, u_{k-1}\}$ on side A and with $U = \{v_1, \ldots, v_{k-2}\}$ and $W = \{v, v_{k-1}\}$ on side B.

Otherwise, by symmetry we may let $u_{k-1} \in A$ be the common neighbor of u and v. Let v'_{k-1} be the neighbor of u in B. Some vertex in $\{v_1, \ldots, v_{k-1}\}$ (label it v_{k-1}) is not matched into $\{u_1, \ldots, u_{k-2}\}$ by the matching of type k. Let u'_{k-1} be the neighbor of v_{k-1} in A. Now proceed as before with $U = \{u_1, \ldots, u_{k-2}, u'_{k-1}\}$ and $W = \{u, u_{k-1}\}$ on side A, and with $U = \{v_1, \ldots, v_{k-2}, v'_{k-1}\}$ and $W = \{v, v_{k-1}\}$ on side B. See Figure 4.

Case 4: *Vertices u and v lie on the same side of $\hat{Q}_k$.* Suppose first that u and v are adjacent. Since $\hat{Q}_k$ is triangle-free, u and v have no common neighbors. Let $u_1, \ldots, u_{k-2}$ be the neighbors of u in A other than v, let $v_1, \ldots, v_{k-2}$ be the neighbors of v in A other than u, and let u' and v' be the respective neighbors of u and v in B. Finally, let v'_i be the neighbor of v_i in B, for $1 \leq i \leq k-3$.

We apply Corollary 7 in A using $U = \{u_1, \ldots, u_{k-2}, v_1, \ldots, v_{k-2}\}$ and $W = \{u, v\}$. With probability $1 - o(2^{-2k})$, we obtain $2k-4$ pairwise disjoint sets $U_1, \ldots, U_{k-2}, V_1, \ldots, V_{k-2}$ (also

9

Figure 4: Case 3: vertices u and v lie in opposite sides of $\hat{Q}_k$ and are at distance 2.

disjoint from $\{u, v\}$) such that U_j is u_j-proper and V_j is v_j-proper, for $1 \le j \le k-2$. (In this application of Corollary 7, the sets $V_1, \ldots, V_{k-3}$ will not be used further. We include $v_1, \ldots, v_{k-3}$ in U because we must keep them out of $U_1, \ldots, U_{k-2}$ and V_{k-2} in order to build paths through them. We cannot place them in W, because Corollary 7 applies only for W of constant size, not growing with k.)

We also use Corollary 7 with $W = \emptyset$ (or simply Lemma 6) in B with $U = \{v_1', \ldots, v_{k-3}', v', u'\}$ to obtain pairwise disjoint sets $V_1', \ldots, V_{k-3}', V', U'$ having the desired properties to those above (with probability $1 - o(2^{-2k})$). We then proceed as in the previous cases by exposing the last matching (between A and B) to discover edges between U_j and V_j' for $1 \le j \le k-3$, between U_{k-2} and V', and between V_{k-2} and U'. See Figure 5.


Figure 5: Case 4: vertices u and v lie on the same side of $\hat{Q}_k$.

Finally, suppose that u and v are not adjacent. Let ℓ be their number of common neighbors in A. By Lemma 9, u and v do not have the same neighborhoods in A, so $0 \le \ell \le k-2$. The common neighbors immediately yield ℓ short u, v-paths. Let $u_1, \ldots, u_{k-1-\ell}$ be the remaining neighbors of u in A, and let $v_1, \ldots, v_{k-1-\ell}$ be the remaining neighbors of v in A; these sets are nonempty. We proceed as in the previous argument, with $k-1-\ell$ in place of $k-2$.

We conclude that given any two vertices u and v, with probability at least $1 - o(2^{-2k})$ we find k internally disjoint u, v-paths of length at most $\frac{k}{\lg k} + O\left(\frac{k}{\lg^2 k}\right)$. Taking the Union Bound over all $O(2^{2k})$ pairs finishes the proof. $\qquad\square$

10

References

[1] I. Banič, J. Žerovnik, Wide diameter of Cartesian graph bundles, Discrete Math. 310 (2010) 1697–1701.

[2] F.B. Chedid, R.B. Chedid, A new variation on hypercubes with smaller diameter, Inform. Process. Lett. 46 (1993) 275–280.

[3] K. Efe, A variation on the hypercube with lower diameter, *IEEE Trans. Comput.* 40 (11) (1991), 1312–1316.

[4] A.H. Esfahanian, L.M. Ni, B.E. Sagan, The twisted N-cube with application to multiprocessing, IEEE Trans. Comput. 40 (1991) 88–93.

[5] E. Flandrin, H. Li, Mengerian properties, hamiltonicity, and claw-free graphs, *Networks* 24 (3) (1994), 177–183.

[6] C.M. Hartman, Extremal Problems in Graph Theory, Ph.D. Thesis, University of Illinois, 1997.

[7] D.F. Hsu, Y.D. Lyuu, A graph theoretical study of transmission delay and fault tolerance, *Proceedings of the 4th ISMM International Conference on Parallel and Distributed Computing and Systems*, 20–24, 1991.

[8] D.F. Hsu, T. Łuczak, On the k-diameter of k-regular k-connected graphs, *Discrete Math.* 133 (1) (1994), 291–296.

[9] S. Janson, T. Łuczak, A. Ruciński, *Random Graphs*, Wiley, New York, 2000.

[10] S.C. Liaw, G.J. Chang, Wide diameters of butterfly networks, Taiwan. J. Math. 3 (1) (1999) 83–88.

[11] M. Ma, D.B. West, J.-M. Xu, The vulnerability of the diameter of enhanced hypercubes, Theoretical Computer Science 694 (2017), 60–65.

[12] K. Menger, Zur allgemeinen Kurventheorie, Fund. Math. 10 (1927) 95–115.

[13] H. Qi, X. Zhu, The wide-diameter of $Z_{n,k}$, *Discrete Appl. Math.* 219 (2017), 193–201.

[14] Y. Saad, M.H. Schultz, Topological properties of hypercubes, *IEEE Transac. Comput.* 37 (7) (1988), 867–872.

[15] E. Sinó, J.L.A. Yebra, The vulnerability of the diameter of folded n-cubes, Discrete Math. 174 (1997) 317–322.

[16] T.-H. Tsai, Y.-C. Chen, J.J.M. Tan, Topological properties on the wide and fault diameters of exchanged hypercubes, IEEE Trans. Parallel Distrib. Syst. 25 (2014), 3317–3327.

[17] N.C. Wormald, Models of random regular graphs, Surveys in Combinatorics, 1999, J.D. Lamb and D.A. Preece, eds. London Mathematical Society Lecture Note Series, vol 276, pp. 239–298, Cambridge University Press, Cambridge, 1999.

[18] J.-M. Xu, Wide-diameter of Cartesian product graphs and digraphs, J. Combin. Optim. 8 (2) (2004), 171–181.

[19] J.-S. Yang, J.-M. Chang, K.-J. Pai, H.-C. Chan, Parallel construction of independent spanning trees on enhanced hypercubes, IEEE Trans. Parallel Distrib. Syst. 26 (2015) 3090–3098.

[20] X. Zhu, A hypercube variant with small diameter, *J. Graph Theory* 85(3) (2017), 651–660.