A Ramsey-type result for the hypercube

Noga Alon * Radoš Radoičić † Benny Sudakov ‡ Jan Vondrák §

Abstract

We prove that for every fixed k and $\ell \geq 5$ and for sufficiently large n, every edge coloring of the hypercube Q_n with k colors contains a monochromatic cycle of length 2ℓ. This answers an open question of Chung. Our techniques provide also a characterization of all subgraphs H of the hypercube which are *Ramsey*, i.e., have the property that for every k, any k-edge coloring of a sufficiently large Q_n contains a monochromatic copy of H.

1 Introduction

Let Q_n denote the graph of the n-dimensional hypercube whose vertex set is $\{0,1\}^n$ and two vertices are adjacent if they differ in exactly one coordinate. Ramsey and Turán-type questions concerning the hypercube were mentioned in a 1984 paper by Erdős [8], but in fact had been considered even earlier, as in this paper he outlined a collection of "old unsolved problems which had been perhaps undeservedly neglected". In one of these problems he asked how many edges of an n-dimensional hypercube are necessary to imply the existence of a 4-cycle. Erdős conjectured that $(\frac{1}{2}+o(1))n2^{n-1}$ edges are enough to force the appearance of C_4. A similar question was posed for the existence of a cycle $C_{2\ell}$ for $\ell > 2$ where Erdős asked whether $o(n)2^n$ edges would suffice (see also [9]). Since Q_n is a bipartite graph, clearly only cycles of even length are in question.

It is easy to see that there are $n2^{n-2}$ edges of Q_n avoiding a C_4, e.g., for all odd values of $1 \leq k \leq n$ take those edges lying between levels $k-1$ and k. This example is not maximal and can be improved by adding some independent edges. The best example to date was obtained by Brass, Harborth and Nienborg [2]. For $n = 4^t$ it has $(n+\sqrt{n})2^{n-2}$ edges which may well be a tight bound for Erdős's

*Schools of Mathematics and Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 69978, Israel and IAS, Princeton, NJ 08540, USA. Email address: nogaa@tau.ac.il. Research supported in part by the Israel Science Foundation, by the Hermann Minkowski Minerva Center for Geometry at Tel Aviv University and by the Von Neumann Fund.

†Department of Mathematics, Rutgers University, Piscataway, NJ 08854, USA. Email address: rados@math.rutgers.edu.

‡Department of Mathematics, Princeton University, Princeton, NJ 08544, USA. Email address: bsudakov@math.princeton.edu. Research supported in part by NSF grants DMS-0355497, DMS-0106589, and by an Alfred P. Sloan fellowship.

§Department of Mathematics, MIT, Cambridge, MA 02139, USA. Email address: vondrak@math.mit.edu.

conjecture. Bialostocki proved in [1] that for any 2-edge coloring of Q_n without a monochromatic C_4, the number of edges in each color is at most $(n + \sqrt{n})2^{n-2}$. Hence, this is indeed the maximum size of a C_4-avoiding set of edges, with the additional assumption that it intersects every C_4 in at least one edge. However, this assumption appears difficult to remove. On the other hand, Chung [4] proved that any subset of $\alpha n 2^{n-1}$ edges, where $\alpha \doteq 0.623$, must contain a C_4. This remains the best upper bound to this date. For small values of n, the exact number of edges in a largest C_4-free subgraph of Q_n was determined in [7], [10]. Some further results on C_4-avoiding sets of edges which are connecting vertices of three consecutive levels of the hypercube can be found in [11].

For longer cycles $C_{2\ell}$, Erdős's question was resolved positively for even $\ell \geq 4$. In [4], Chung proved that for a fixed even $\ell \geq 4$, any subset of edges of Q_n avoiding $C_{2\ell}$ has size $o(n)2^n$. On the other hand, she showed that this is not the case for cycles of length 6 since the edges of Q_n can be colored using 4 colors so that there is no monochromatic C_6 (a similar coloring was discovered also in [3]). Therefore, a subset of $\frac{1}{4}n2^{n-1}$ edges avoiding C_6 exists. This sparked new interest in edge colorings of the hypercube without monochromatic cycles. A 3-coloring avoiding a monochromatic cycle of length 6 was found in [6]. On the other hand, it was shown in [4] that any subset of $(\sqrt{2} - 1 + o(1))n2^{n-1}$ edges must contain a C_6.

Since a coloring avoiding a monochromatic $C_{2\ell}$ using a constant number of colors is impossible for even $\ell \geq 4$ due to [4], it remains to determine whether such a coloring exists for odd $\ell \geq 5$. This question was posed by Chung in [4] (see also [5], pp. 43–44). In this paper, we prove the following theorem which answers it negatively.

Theorem 1.1 *For every fixed k and $\ell \geq 5$ and sufficiently large $n \geq n_0(k, \ell)$, every edge coloring of the hypercube Q_n with k colors contains a monochromatic cycle of length 2ℓ.*

In fact, our techniques provide a characterization of all subgraphs H of the hypercube which are *Ramsey*, i.e., have the property that for every k, any k-edge coloring of a sufficiently large Q_n contains a monochromatic copy of H. We also present examples of graphs which are not Ramsey but the number of colors required to avoid their monochromatic copies is arbitrarily large. (In contrast, every even cycle is either Ramsey or it can be avoided using 2 or 3 colors.) More details are given in Section 4.

Definitions and notation.

Recall that Q_n denotes the n-dimensional hypercube whose vertex set is $\{0, 1\}^n$. We refer to the n coordinates as *bits* and write vertices as n-bit words, for example $x = [10001], y = [10101]$. Edges are between vertices that differ in exactly one bit. We call the unique bit where $x_i \neq y_i$ the *flip-bit*. The vertex where the flip-bit is zero is called the *lower vertex* and the other vertex is called the *upper vertex*. For example, for the vertices x, y above, $\{x, y\}$ is an edge where x is the lower vertex, y is the higher vertex and the 3-rd bit is the flip-bit. To simplify notation, we write such an edge as $[10 * 01]$; the $*$ symbol denotes the flip-bit and we obtain the two vertices of the edge by substituting 0 or 1 in place of $*$.

2 Cycles of length 10

First, we address the question for cycles of length 10. The colorings that have been used in order to avoid monochromatic cycles of length 4 and 6 are based on two parameters: for an edge $e = \{x, y\}$ where x is the lower vertex and j is the flip-bit, define

- $w(e) = \sum_{i=1}^{n} x_i$.
- $p(e) = \sum_{i=1}^{j-1} x_i$.

The first parameter distinguishes different *levels* of vertices; each level is defined by a constant value of $\sum_{i=1}^{n} x_i$. The second parameter further distinguishes the edges between each pair of consecutive levels; we call $p(e)$ the *prefix sum* of e. To avoid monochromatic cycles of length 4 and 6, it is enough to consider colorings based on these two parameters, taken modulo a suitable number (see, e.g., [6]).

In contrast, it turns out that for cycles of length 10, no such coloring can work. The reason is the following cycle in Q_5:

$$\begin{aligned}
e_1 &= [1 * 0\ 0\ 1] \\
e_2 &= [1\ 1\ 0\ 0 *] \\
e_3 &= [1\ 1\ 0 * 0] \\
e_4 &= [* 1\ 0\ 1\ 0] \\
e_5 &= [0\ 1 * 1\ 0] \\
e_6 &= [0\ 1\ 1 * 0] \\
e_7 &= [* 1\ 1\ 0\ 0] \\
e_8 &= [1 * 1\ 0\ 0] \\
e_9 &= [1\ 0\ 1\ 0 *] \\
e_{10} &= [1\ 0 * 0\ 1]
\end{aligned}$$

Here, every odd edge e_{2i-1} goes from $* = 0$ to $* = 1$, and every even edge e_{2i} goes from $* = 1$ to $* = 0$. The reader can verify that these edges form a C_{10}. Observe that $w(e_i)$ is equal for all these edges which corresponds to the fact that the cycle is alternating between two levels of the hypercube. Regarding $p(e_i)$, it is not the same for each edge, but it depends only on the location of the flip-bit; for each pair of edges with the same flip-bit, $p(e_i)$ is the same: either 0, 1, or 2. It is not difficult to see that for any coloring of the type ($p(e_i) \bmod k$), we can insert blocks of 1's between these 5 bits so that the resulting cycle (in a higher-dimensional hypercube) is monochromatic. In the following, we show that there is a deeper reason why this kind of coloring cannot avoid monochromatic 10-cycles: in fact, for *any* coloring with a fixed number of colors, there is some form of the cycle above which turns out to be monochromatic.

Theorem 2.1 *For any fixed k and sufficiently large $n \geq n_0(k)$, every edge coloring of Q_n with k colors contains a monochromatic cycle of length* 10.

Proof. Consider an arbitrary k-edge coloring χ of Q_n, for a very large n to be chosen later. Let's consider only edges between levels $2k$ and $2k+1$, which are defined by $2k$ coordinates equal to 1 and a given flip-bit. We call these $2k+1$ bits the *support* of an edge. We can encode each edge uniquely by (S,p) where $S \subset [n]$ is the support of the edge, and $p \in \{0,1,...,2k\}$ denotes its prefix sum. In other words, p determines the relative location of the flip-bit in the support of the edge. Each pair (S,p) gets some color $\chi(S,p)$ in our coloring. Let's assign a vector $c(S) = (\chi(S,0),\ldots,\chi(S,2k))$ to each subset S, i.e. the edge colors for all possible locations of the flip-bit. We get a coloring of the complete $(2k+1)$-uniform hypergraph on $[n]$, using k^{2k+1} colors.

By Ramsey's Theorem for hypergraphs, for sufficiently large $n \geq n_0(k)$, there is a subset of coordinates $T \subset [n]$ of size $2k+3$ such that all $(2k+1)$-subsets $S \subset T$ have the same color $c(S) = c^*$. Now, since c^* has $2k+1$ coordinates colored by k colors, there must be 3 indices $p_1, p_2, p_3 \in \{0,...,2k\}$ such that $c^*_{p_1} = c^*_{p_2} = c^*_{p_3}$. This means that all edges (S,p_i) where $S \subset T, |S| = 2k+1$ and $i = 1,2,3$, have the same color. We show that this set of edges contains a monochromatic cycle of length 10, which can be obtained from the cycle above by inserting blocks of 1's of suitable length in front of the first bit, between the first and second bit, and between the third and fourth bit. Since we want the prefix sum $p(e_i)$ for each edge to be equal to p_1, p_2 or p_3, we choose these blocks as $\alpha = 1^{p_1}$ (a string of p_1 ones), $\beta = 1^{p_2-p_1-1}$, $\gamma = 1^{p_3-p_2-1}$ and $\delta = 1^{2k-p_3}$. The cycle looks like this: (only the coordinates of T are shown, the rest is zero)

$$\begin{aligned}
e_1 &= [\,\alpha\ 1\ \beta * 0\ \gamma\ 0\ 1\ \delta\,] \\
e_2 &= [\,\alpha\ 1\ \beta\ 1\ 0\ \gamma\ 0 * \delta\,] \\
e_3 &= [\,\alpha\ 1\ \beta\ 1\ 0\ \gamma * 0\ \delta\,] \\
e_4 &= [\,\alpha * \beta\ 1\ 0\ \gamma\ 1\ 0\ \delta\,] \\
e_5 &= [\,\alpha\ 0\ \beta\ 1 * \gamma\ 1\ 0\ \delta\,] \\
e_6 &= [\,\alpha\ 0\ \beta\ 1\ 1\ \gamma * 0\ \delta\,] \\
e_7 &= [\,\alpha * \beta\ 1\ 1\ \gamma\ 0\ 0\ \delta\,] \\
e_8 &= [\,\alpha\ 1\ \beta * 1\ \gamma\ 0\ 0\ \delta\,] \\
e_9 &= [\,\alpha\ 1\ \beta\ 0\ 1\ \gamma\ 0 * \delta\,] \\
e_{10} &= [\,\alpha\ 1\ \beta\ 0 * \gamma\ 0\ 1\ \delta\,]
\end{aligned}$$

It can be seen that for edges e_4 and e_7, the prefix sum is $|\alpha| = p_1$, for edges e_1, e_5, e_8, e_{10}, the prefix sum is $|\alpha| + |\beta| + 1 = p_2$ and for e_2, e_3, e_6 and e_9, the prefix sum is $|\alpha| + |\beta| + |\gamma| + 2 = p_3$. Thus each of these edges is encoded by (S,p_1), (S,p_2) or (S,p_3) for some $S \subset T, |S| = 2k+1$, and therefore they all have the same color. □

3 Cycles of length $2\ell \geq 12$

In this section, we extend the proof for 10-cycles to all even cycles of length at least 12. All we have to do is find a cycle of length 2ℓ with properties similar to the 10-cycle shown above.

Lemma 3.1 *For any $\ell \geq 6$, Q_ℓ contains a cycle of length 2ℓ in which each edge has a support of size 3, such that for some $1 < a < b < \ell$,*

1. *Each edge with a flip-bit located in $\{1, \ldots, a\}$ has prefix sum $p(e) = 0$.*
2. *Each edge with a flip-bit located in $\{a+1, \ldots, b\}$ has $p(e) = 1$.*
3. *Each edge with a flip-bit located in $\{b+1, \ldots, \ell\}$ has $p(e) = 2$.*

Proof. For any $\ell \geq 6$, we define a cycle on vertices with ℓ bits, consisting of edges $(e_1, e_2, \ldots, e_{2\ell})$, associated with a permutation $\pi \in S_\ell$:

- Each edge in the cycle has a support of size 3.
- Every odd edge e_{2i-1} has bits $\pi(i), \pi(i+1)$ equal to 1, and $\pi(i+2)$ is the flip-bit (going *up*).
- Every even edge e_{2i} has bits $\pi(i+1), \pi(i+2)$ equal to 1, and $\pi(i)$ is the flip-bit (going *down*).

To simplify notation, we consider π as a periodic function, i.e. $\pi(i+\ell) = \pi(i)$ for any i.

It is easy to verify that this is indeed a cycle of length 2ℓ. We need to find a permutation such that the cycle satisfies the requirements of the lemma. Observe that for a given i, there are exactly two edges with flip-bit $\pi(i)$. The other non-zero bits on these two edges are $\pi(i-1), \pi(i-2)$ for one edge and $\pi(i+1), \pi(i+2)$ for the other edge. Thus the prefix sum $p(e)$ for each edge is determined by the two nearest elements in the permutation, on either side.

First, consider $\ell \geq 6$ divisible by 3 and set $\ell = 3a$, $b = 2a$. Take an arbitrary permutation of type $(A, B, C, A, B, C, \ldots, A, B, C)$, where each A stands for some element in $\{1, \ldots, a\}$, each B for an element in $\{a+1, \ldots, b\}$ and each C for an element in $\{b+1, \ldots, \ell\}$. It can be seen that for each A, the two nearest elements in the permutation, on either side, are B, C, which defines the location of the other two non-zero bits. Such an edge looks like this:

$*$	1	1

where the three blocks correspond to bits of type A, B and C. Therefore, in this case $p(e) = 0$. Similarly for each B, the two nearest elements on each side are A, C and the prefix sum in both cases is $p(e) = 1$:

1	$*$	1

For each C, the two nearest elements on each side are A, B and the prefix sum is $p(e) = 2$:

1	1	$*$

Next, we handle the case of $\ell = 3a + 1$. We insert an element of a new type X, located between B and C; that is, X stands for $2a + 1$ and the range for C is shifted to $\{2a+2, \ldots, \ell\}$. We take a permutation of type $(A, X, B, C,\ A, B, C, \ldots,\ A, B, C)$. Note that for each flip-bit of type A, the other two non-zero bits are of types B, C or B, X and the prefix sum is $p(e) = 0$. For flip-bits of type B, the non-zero bits are of types A, C or A, X; in either case, $p(e) = 1$. For a flip-bit of type X, the non-zero bits are of types A, C or B, C; again, $p(e) = 1$. Finally, for flip-bits of type C, the two non-zero bits are of types A, B or X, B, and $p(e) = 2$. So the lemma holds with $b = 2a + 1$.

For $\ell = 3a + 2 > 5$, X stands for an element in $\{2a + 1, 2a + 2\}$ and the range for C is shifted to $\{2a + 3, \ldots, \ell\}$. We take a permutation of type $(A, X, B, C,\ A, X, B, C, \ldots,\ A, B, C)$. The same analysis yields that the prefix sums are 0 for flip-bits of type A, 1 for flip-bits of type B or X, and 2 for flip-bits of type C. Therefore lemma holds with $b = 2a + 2$. □

Theorem 3.2 *For any fixed k and $\ell \geq 6$ and sufficiently large $n \geq n_0(k, \ell)$, every edge coloring of Q_n with k colors contains a monochromatic cycle of length 2ℓ.*

Proof. Given a coloring $\chi : E(Q_n) \to [k]$, consider only edges with support of size $|S| = 2k + 1$. Just like in the proof of Theorem 2.1, encode edges by their support and prefix sum (S, p), and define a coloring $c(S) = (\chi(S, 0), \ldots, \chi(S, 2k))$ of the complete $(2k + 1)$-uniform hypergraph on $[n]$ using k^{2k+1} colors. By Ramsey's theorem, for sufficiently large $n \geq n_0(k, \ell)$, there is a subset $T \subset [n]$ of size $2k + \ell - 2$ such that $c(S) = c^*$ for all $S \subset T, |S| = 2k + 1$. By the pigeonhole principle, there are three elements $p_1, p_2, p_3 \in \{0, 1, \ldots, 2k\}$ such that $c^*_{p_1} = c^*_{p_2} = c^*_{p_3}$, i.e. all edges (S, p_i) for $S \subset T, |S| = 2k + 1$ and $i = 1, 2, 3$ have the same color.

Now we take the cycle $\mathcal{C}$ provided by Lemma 3.1 and embed it in the monochromatic subgraph that we found in Q_n. The ℓ-bit representation of $\mathcal{C}$ consists of three blocks defined by the parameters $1 < a < b < \ell$. As Lemma 3.1 guarantees, the prefix sum of each edge is either 0, 1 or 2, depending on the block in which the flip-bit of the edge appears. We insert strings of 1's between these blocks, in order to convert the prefix sums to p_1, p_2 and p_3: $\alpha = 1^{p_1}$ in front of the first bit, $\beta = 1^{p_2 - p_1 - 1}$ after the first a bits, $\gamma = 1^{p_3 - p_2 - 1}$ after b bits and $\delta = 1^{2k - p_3}$ at the end. We obtain a cycle embedded in $Q_{2k+\ell-2}$ where the prefix sum for each edge is p_1, p_2 or p_3. Finally, we embed this subcube $Q_{2k+\ell-2}$ in Q_n by laying its $(2k + \ell - 2)$-bit representation on the subset of coordinates $T \subset [n]$; all other coordinates are fixed to be zero. The edges of $\mathcal{C}$ thus embedded in Q_n have their support in T and prefix sums equal to p_1, p_2 or p_3; therefore the cycle is monochromatic. □

4 Ramsey subgraphs of the hypercube

4.1 A full characterization

In this section we consider more generally the question of finding monochromatic subgraphs in large edge-colored hypercubes. Call a graph H *k-Ramsey* if every k-edge coloring of a sufficiently large hypercube contains a monochromatic copy of H. Call H *Ramsey*, if it is k-Ramsey for every k. Therefore, Theorem 1.1 asserts that every even cycle of length at least 10 is Ramsey. Our technique here provides a characterization of all subgraphs of the hypercube which are Ramsey. This is stated in the following theorem.

Theorem 4.1 *Let H be a fixed subgraph of a hypercube. Then H is Ramsey if and only if there exists an embedding of H between two levels of a hypercube such that in this embedding all edges $e \in E(H)$ with the same flip-bit have the same prefix sum $p(e)$.*

Sketch of proof. Assume there is an embedding as above, between levels t and $t+1$ of a hypercube Q_m. Then given a k-edge coloring of a large Q_n, we apply Ramsey's Theorem for hypergraphs, as in the proof of Theorems 2.1 and 3.2, to obtain a sufficiently large subcube Q_s in which the color of each edge e with support of size w is determined by the value of $p(e) \in \{0, 1, \ldots, w\}$. We choose w large enough so that it is possible to find $M \subset \{0, 1, \ldots, w\}, |M| = m$, such that any two elements $i, j \in M$ are at least t apart, and the edges whose prefix sums are in M all get the same color. Then we can take our embedding of H and add suitable blocks of 1's between the bits so that all the prefix sums fall in M. Finally, we add a block of 0's to embed H in Q_s so that the color of each edge is determined by $p(e) \in M$ and consequently this copy of H is monochromatic.

Conversely, assume that for every embedding of H between two levels there are two edges with the same flip bit and different prefix sums. Consider the coloring $\chi(e) = (w(e) \mod 2, p(e) \mod \lceil d/2 \rceil)$ where d denotes the diameter of H. Then a copy of H could be possibly monochromatic only if it lies between two levels, but then there are two edges e_1, e_2 with the same flip bit and different prefix sums. The prefix sums cannot differ by a multiple of $\lceil d/2 \rceil$ because then the suffix sums would differ by the same amount and together with the flip-bit we would get two vertices at distance more than d. Therefore these two edges have different colors. □

We remark that although the above result characterizes all Ramsey subgraphs of the hypercube, this characterization is not very efficient. Still, it can be checked in time that depends only on the size of the small graph H. This is because it suffices to check embeddings of H in a hypercube Q_m, with m being the number of edges of H.

4.2 The number of necessary colors

Considering our characterization of Ramsey subgraphs in the hypercube, we can ask what is the number of colors necessary to avoid a monochromatic H, given that H is not Ramsey. We have seen that C_4 is not Ramsey, and a monochromatic C_4 can be avoided using only 2 colors. C_6 is not Ramsey either, but in fact it is 2-Ramsey and we need 3 colors to avoid a monochromatic C_6. Note that the number of colors needed in both cases is equal to the diameter of the subgraph in question. The proof of Theorem 4.1 shows that for any H of diameter d which is not Ramsey, we can also say that H is not k-Ramsey for any $k > d$.

Here, we show that for any k, there exists a graph $H_{m,k}$ which is k-Ramsey but not $(k+1)$-Ramsey. The diameter of $H_{m,k}$ is $O(k)$ which means that the number of colors required to avoid a monochromatic subgraph of diameter d can be indeed $\Omega(d)$.

Construction. For $k > 0, m > k$, let $Q_{m,k}$ denote the subgraph of Q_m which contains all vertices on levels k and $k+1$ and all edges between them whose prefix sum is $p(e) = 0$ or k. We represent vertices by their support, i.e. the subset of coordinates equal to 1. We define $H_{m,k}$ as the subgraph of $Q_{m,k}$ induced by all vertices at distance at most $2k+1$ (in $Q_{m,k}$) from the vertex represented by $K = \{1, 2, \ldots, k\}$.

The structure of $Q_{m,k}$ is the following: every vertex on the upper level $k+1$ has degree 2. If this vertex is given by a $(k+1)$-subset $A = \{a_1 < a_2 < \ldots < a_{k+1}\}$, then its two neighbors on level k are given by $A_1 = \{a_1, a_2, \ldots, a_k\}$ and $A_2 = \{a_2, a_3, \ldots, a_{k+1}\}$. On the other hand, vertices on the lower level have a larger degree and their neighbors are obtained by adding any element which is smaller or larger than everything in the subset. Thus edges in terms of subsets correspond to adding/removing a minimum or maximum element.

The Ramsey properties that we prove hold equally for $Q_{m,k}$ and $H_{m,k}$. However, note that $Q_{m,k}$ is not a connected graph (for example, $\{1, 2, \ldots,\ k-1, m\}$ represents an isolated vertex). $H_{m,k}$ is connected and by definition, its diameter is $O(k)$.

Lemma 4.2 *For any $m > k$, $H_{m,k}$ is $k-$Ramsey, i.e. for any k-edge-coloring of a sufficiently large hypercube, there is a monochromatic copy of $H_{m,k}$.*

Proof. We prove in fact that $Q_{m,k}$ is $k-$Ramsey. First, we show that for any $t \in \{1, 2, \ldots, k\}$, there is $r(t)$ such that $Q_{m,k}$ can be embedded in $Q_{r(t),t}$. That is, we would like to have prefix sums 0 and t instead of 0 and k. For that purpose, consider all $(k-t+1)$-subsets of $[m]$ and index them lexicographically. The indices go from 1 up to $r(t) = \binom{m}{k-t+1}$ and the index of a subset $A \subseteq [m], |A| = k-t+1$ is denoted by $\phi(A)$. Define a mapping from the k-th level of Q_m to the t-th level of $Q_{r(t)}$ as follows. For each subset $A = \{a_1 < a_2 < \ldots < a_k\} \subset [m]$, let $A_1 = \{a_1, \ldots, a_{k-t+1}\}$, $A_2 = \{a_2, \ldots, a_{k-t+2}\}, \ldots,\ A_t = \{a_t, \ldots, a_k\}$. We map the subset A to

$$f(A) = \{\phi(A_1), \phi(A_2), \ldots, \phi(A_t)\}.$$

Similarly, we define a mapping from the $(k+1)$-th level of Q_m to the $(t+1)$-th level of $Q_{r(t)}$. We cover $B = \{b_1 < b_2 < \ldots < b_{k+1}\}$ by $t+1$ subsets $B_1, B_2, \ldots, B_{t+1}$ where $B_i = \{b_i, \ldots, b_{k-t+i}\}$ and we set $f(B) = \{\phi(B_1), \phi(B_2), \ldots, \phi(B_{t+1})\}$. The edges of $Q_{m,k}$ incident with this vertex are obtained by removing either b_1 or b_{k+1} which produces two neighbors on the lower level. Observe that the two neighbors map to $f(B \setminus \{b_1\}) = \{\phi(B_2), \ldots, \phi(B_{t+1})\}$ and $f(B \setminus \{b_{k+1}\}) = \{\phi(B_1), \ldots, \phi(B_t)\}$, which are neighbors of $f(B)$ in $Q_{r(t)}$. Moreover, the lexicographic ordering ensures that $\phi(B_1) < \phi(B_2) < \ldots < \phi(B_{t+1})$ and the prefix sums of these two edges are 0 and t. Thus the edges of $Q_{m,k}$ map to edges of $Q_{r(t),t}$.

Now consider any k-edge-coloring of Q_n. We choose $n \geq n_0(m,k)$ large enough so that applying Ramsey's theorem (as in the proof of Theorems 2.1 and 3.2), we obtain a subcube $Q_s, s = 2^m$, where the coloring of edges between levels k and $k+1$ depends only on the prefix sum. Since the available prefix sums are $0, 1, \ldots, k$, there must be two prefix sums $p_1 < p_2$ which get the same color. Let $t = p_2 - p_1$ and construct an embedding of $Q_{m,k}$ in $Q_{r(t),t}$. Recall that this embedding is between levels t and $t+1$ so that all the prefix sums are 0 or t. Since $r(t) \ll s$, we still have enough space to add a block of 1^{p_1} in front of the bit representation of $Q_{r(t)}$, a block of 1^{k-p_2} at the end, and another block of 0's at the end, so that we get an embedding of $Q_{r(t),t}$ (and thus also of $Q_{m,k}$) between levels k and $k+1$ of Q_s such that the prefix sums of all edges are p_1 or p_2. This gives a monochromatic copy $Q_{m,k}$ and since $H_{m,k} \subset Q_{m,k}$ it gives a a monochromatic copy of $H_{m,k}$ as well. □

Thus at least $k+1$ colors are necessary to avoid a monochromatic copy of $H_{m,k}$. We show that $k+1$ colors are also sufficient, and the right coloring is the natural choice of $(p(e) \bmod k+1)$. However, first we note a simple property of $H_{m,k}$ which will be useful in the proof.

Lemma 4.3 *$H_{m,k}$ contains all the vertices represented by subset $A \subseteq \{k+1, k+2, \ldots, m\}$ of size $|A| = k$ or $k+1$ and the distance between $K = \{1, 2, \ldots, k\}$ and A in $H_{m,k}$ is $k + |A|$.*

Proof. Recall that $H_{m,k}$ contains the vertex represented by $K = \{1, 2, \ldots, k\}$ together with each vertex whose distance from K in $Q_{m,k}$ is at most $2k+1$. Consider a subset $A \subseteq \{k+1, k+2, \ldots, m\}$, of size k or $k+1$. We can transform K into A by adding elements of A and removing elements of K alternately, starting from the smallest and ending with the largest. Since we always remove the minimum element and add the maximum element, this corresponds to a path in $Q_{m,k}$ of length $|K| + |A| \le 2k+1$. Therefore A also represents a vertex of $H_{m,k}$. □

Lemma 4.4 *For m sufficiently large, $H_{m,k}$ is not $(k+1)$-Ramsey. In particular, for any n, there is no monochromatic copy of $H_{m,k}$ in the hypercube Q_n in which edge e is colored by $(p(e) \bmod k+1)$.*

Proof. Consider any embedding of $H_{m,k}$ in Q_n, represented by a function $g : 2^{[m]} \to 2^{[n]}$. We consider m very large, so that we can use Ramsey's theorem repeatedly to select a subgraph of $H_{m,k}$ with specific properties. In the first step, consider the subset of the lower vertices of $H_{m,k}$, represented by k-subsets of $X = \{k+1, \ldots, m\}$ (see Lemma 4.3). Since in $H_{m,k}$, all these vertices are within distance $2k$ from $K = \{1, 2, \ldots, k\}$, this must also be the case in the new embedding. The images of these vertices can occupy at most $2k+1$ different levels of Q_n. Define the color of $A \in \binom{X}{k}$ by $|g(A)|$ which can take at most $2k+1$ different values. By Ramsey's theorem, there is a large subset $X' \subseteq X$ such that $\binom{X'}{k}$ maps to one level, i.e. $|g(A)|$ is constant for all $A \in \binom{X'}{k}$.

Choose a fixed subset $L \subset X'$ by taking the k smallest elements of X' and denote by $Y = X' \setminus L$ the remaining elements in X'. Again, all vertices represented by $A \in \binom{Y}{k}$ are at distance at most $2k$ from L and therefore the same holds in the new embedding. Now, $|g(A)| = |g(L)|$ for all $A \in \binom{Y}{k}$ and due to the distance condition, we have $|g(L) \setminus g(A)| = |g(A) \setminus g(L)| \le k$. We set $g_0(A) = g(L) \setminus g(A)$ and $g_1(A) = g(A) \setminus g(L)$. By Ramsey's theorem, we can find a large subset $Y' \subseteq Y$ such that $|g_0(A)| = |g_1(A)| = k'$ for all $A \in \binom{Y'}{k}$, where $1 \le k' \le k$.

Consider the upper vertices represented by $(k+1)$-subsets of Y'. For $B = \{b_1 < b_2 < \ldots < b_{k+1}\} \subset Y'$, set $A_1 = \{b_1, \ldots, b_k\}$ and $A_2 = \{b_2, \ldots, b_{k+1}\}$ to be its two neighbors in $H_{m,k}$. A_1 and A_2 represent lower vertices at distance 2 in $H_{m,k}$, so likewise, $g(A_2)$ can be obtained from $g(A_1)$ by swapping one element for another. Since $|g_0(A_1)| = |g_0(A_2)| = |g_1(A_1)| = |g_1(A_2)|$, this means that either $|g_0(A_1) \triangle g_0(A_2)| = 2$ and $|g_1(A_1) \triangle g_1(A_2)| = 0$, or vice versa. Denote by $q \in \{0,1\}$ which of these cases occurs; i.e., assume $g_q(A_2) = g_q(A_1) \cup \{x_2\} \setminus \{x_1\}$, while $g_{1-q}(A_1) = g_{1-q}(A_2)$. Also, it could be the case that either $g(B) = g(A_1) \cup g(A_2)$ or $g(B) = g(A_1) \cap g(A_2)$. Defining $g_0(B) = g(L) \setminus g(B)$ and $g_1(B) = g(B) \setminus g(L)$, we get either $g_q(B) = g_q(A_1) \cup g_q(A_2)$ or $g_q(B) = g_q(A_1) \cap g_q(A_2)$ (while $g_{1-q}(B) = g_{1-q}(A_1) = g_{1-q}(A_2)$). We denote by $r = 0,1$ which of these cases occurs. Finally, denote by p_1, p_2 the relative locations of x_1, x_2, i.e. the number of elements preceding them, in $g_q(A_1) \cup g_q(A_2)$. We assign the color (p_1, p_2, q, r) to the subset $B \in \binom{Y'}{k+1}$. We have $0 \le p_1 \ne p_2 \le k'$, and the number of colors is at most $4k'(k'+1) \le 4k(k+1)$. By Ramsey's theorem, we find a subset $Z \subset Y'$ of size $2k+1$ such that $\binom{Z}{k+1}$ is monochromatic. This means that for any $B \in \binom{Z}{k+1}$ and its two neighbors A_1, A_2 in $H_{m,k}$, we have $|g_q(A_1) \triangle g_q(A_2)| = 2$ for the same $q \in \{0,1\}$, $g_q(B)$ is always either the union or the intersection of $g_q(A_1)$ and $g_q(A_2)$, and the relative locations of the swapped elements x_1, x_2 in $g_q(A_1) \cup g_q(A_2)$ are always the same p_1, p_2.

Denote the elements of Z by $z_0 < z_1 < z_2 < \ldots < z_{2k}$ and consider a path in $H_{m,k}$ containing vertices $A_0 = \{z_0, \ldots, z_{k-1}\}$, $B_0 = \{z_0, \ldots, z_k\}$, $A_1 = \{z_1, \ldots, z_k\}$, $B_1 = \{z_1, \ldots, z_{k+1}\}$, $\ldots$, $A_{k+1} = \{z_{k+1}, \ldots, z_{2k}\}$. By the properties of Z, it holds that $|g_q(A_i) \triangle g_q(A_{i+1})| = 2$ for $i = 0, \ldots, k$. Also, we have $g_{1-q}(A_0) = g_{1-q}(A_1) = \ldots = g_{1-q}(A_{k+1})$. Since $g(A_i) = g(L) \cup g_1(A_i) \setminus g_0(A_i)$, the changes in $g(A_i)$ are determined by changes in $g_q(A_i)$, and the prefix sums of edges along the path are determined by the locations of elements being swapped between $g_q(A_i)$ and $g_q(A_{i+1})$. In the sequence $(g_q(A_0), \ldots, g_q(A_{k+1}))$, the next subset is always obtained by swapping one element for a new element; this happens $k+1$ times, and the size of each set is $k' \le k$. Therefore, scanning the sequence from left to right, there must be an element x^* that appears and then again disappears from the subsets. When x^* appears in $g_q(A_{i+1}) \setminus g_q(A_i)$, its location in $g_q(A_i) \cup g_q(A_{i+1})$ is p_2; when $x^* \in g_q(A_{i'}) \setminus g_q(A_{i'+1})$, its location in $g_q(A_{i'}) \cup g_q(A_{i'+1})$ is p_1.

We have to be careful since the prefix sums of the two corresponding edges are not simply p_1 and p_2. First, there is a contribution from $g(L)$ and $g_{1-q}(A_i) = g_{1-q}(A_{i+1})$, which is always constant and does not influence differences between prefix sums for the same flip-bit. In addition, we have the non-constant contribution from $g_q(A_i)$ and $g_q(A_{i+1})$ which depends on $r \in \{0,1\}$, i.e. whether the intermediate vertex is obtained by taking $g_q(B_i) = g_q(A_i) \cup g_q(A_{i+1})$ or $g_q(A_i) \cap g_q(A_{i+1})$. Due to our Ramsey argument, we know that the same case occurs everywhere along the path. In the first case, the prefix sums are indeed given by p_1 and p_2, modulo the constant contribution from $g(L)$ and $g_{1-q}(A_i)$. Then they differ by exactly $|p_2 - p_1| \le k$. In the second case, when the intermediate vertices are given by intersections, the prefix sums differ only by $|p_2 - p_1| - 1$, due to the fact that there is another element missing in $g_q(A_i) \cup g_q(A_{i+1})$ when x^* is added/removed. This affects the prefix sum for one of the two edges, whichever is the larger of p_1, p_2. However, in this case we cannot have $|p_2 - p_1| = 1$, since this would correspond to a situation where the intermediate vertex $g(B_i)$ is always the same; in other words, the path would be embedded as a star. Therefore, in either case, the prefix sums differ by a number between 1 and k and so these two edges have different colors under our $(k+1)$-coloring. □

5 Concluding remarks

We have proved that for any fixed $\ell \geq 5$, every edge coloring of a sufficiently large hypercube with a fixed number of colors contains a monochromatic cycle of length 2ℓ. For odd ℓ, this answers an open question of Chung. For even $\ell \geq 4$, in fact, she proved a stronger result, namely that any $C_{2\ell}$-free subgraph of Q_n has only an $o(1)$-fraction of the edges of Q_n. It still remains open whether this is also the case for cycles of length 2ℓ for odd $\ell \geq 5$.

Finally, we note that the cycle of length 10 in Section 2 is not chordless; vertices [11000] and [11100] are connected. Curiously, there exists a 4-edge coloring of Q_n that avoids monochromatic *chordless* cycles of length 10: the coloring ν defined by $\nu(e) = (w(e) \mod 2, p(e) \mod 2)$ works. This is proved by a somewhat tedious case analysis, which is omitted. It is interesting to note that for $C_{2\ell}$, with $\ell \geq 6$, the cycles provided by Lemma 3.1 are chordless. Therefore, for each such ℓ, any k-edge coloring of a sufficiently large hypercube contains a monochromatic *induced* cycle of length 2ℓ.

References

[1] A. Bialostocki, Some Ramsey type results regarding the graph of the n-cube, *Ars Combin.* 16 (1983), 39–48.

[2] P. Brass, H. Harborth and H. Nienborg, On the maximum number of edges in a C_4-free subgraph of Q_n, *J. Graph Theory* 19 (1995), 17–23.

[3] A. E. Brouwer, I. J. Dejter and C. Thomassen, Highly symmetric subgraphs of hypercubes, *J. Algebraic Combin.* 2 (1993), 25–29.

[4] F. Chung, Subgraphs of a hypercube containing no small even cycles, *J. Graph Theory* 16 (1992), 273–286.

[5] F. Chung and R. Graham, **Erdős on graphs. His legacy of unsolved problems**, A K Peters Ltd., Wellesley, MA, 1998.

[6] M. Conder, Hexagon-free subgraphs of hypercubes, *J. Graph Theory* 17 (1993), 477–479.

[7] M. R. Emamy-K., P. Guan and P. I. Rivera-Vega, On the characterization of the maximum squareless subgraphs of 5-cube, in: Proceedings of the Twenty-third Southeastern International Conference on Combinatorics, Graph Theory, and Computing (Boca Raton), *Congr. Numer.* 88 (1992), 97–109.

[8] P. Erdős, On some problems in graph theory, combinatorial analysis and combinatorial number theory, in: Graph Theory and Combinatorics (Cambridge, 1983), Academic Press, London, 1984, 1–17.

[9] P. Erdős, Some of my favourite unsolved problems, in: **A tribute to Paul Erdős**, Cambridge Univ. Pres, 1990, 467–478.

[10] H. Harborth and H. Nienborg, Maximum number of edges in a six-cube without four-cycles, *Bull. Inst. Combin. Appl.* 12 (1994), 55–60.

[11] R.H. Schelp and A. Thomason, On quadrilaterals in layers of the cube and extremal problems for directed and oriented graphs, *J. Graph Theory* 33 (2000), 66–82.