On Hypercube Labellings and Antipodal Monochromatic Paths

Tomás Feder* Carlos Subi

Abstract

A labelling of the n-dimensional hypercube H_n is a mapping that assigns value 0 or 1 to each edge of H_n. A labelling is antipodal if antipodal edges of H_n get assigned different values. It has been conjectured that if $H_n, n \geq 2$, is given a labelling that is antipodal, then there exists a pair of antipodal vertices joined by a monochromatic path. This conjecture has been verified by hand for $n \leq 5$. In this paper we verify the conjecture in the case where the labelling is simple in the sense that no square $xyzt$ in H_n has value 0 assigned to xy, zt and value 1 assigned to yz, tx, even if the given labelling is not antipodal. The proof is based on a new property of (not necessarily antipodal) simple labellings of H_n. We also exhibit a large class of simple labellings that thus satisfy the conjecture. Finally we conjecture that even if the given labelling is not antipodal, there is always a path joining antipodal vertices that switches labels at most once, which implies the original conjecture. We establish this new conjecture for $H_n, n \leq 5$ as well.

1 Introduction

We let $H_n = (V, E)$ denote the *n-dimensional hypercube* whose 2^n vertices are the $\{0,1\}$-strings $x = x_1 \cdots x_n$ of length n and whose $n2^{n-1}$ edges join vertices that differ in only one bit position in the corresponding strings. An *labelling of* H_n is a decomposition of H_n into two graphs $G_0 = (V, E_0)$ and $G_1 = (V, E_1)$ such that $E_0 \subseteq E$ and $E_1 = E \setminus E_0$. Two vertices v, v^T in H_n are *antipodal* if the corresponding strings differ in all bit positions. Two edges $e = uv$ and $e^T = u^T v^T$ are also said *antipodal*. A labelling of H_n is called *antipodal* if for every pair of antipodal edges $e, e^T \in E$, exactly one of e, e^T is in E_0 (the other one is in E_1).

It has been conjectured [1] that for every antipodal labelling of H_n, $n \geq 2$, there exists a connected component K_0 of G_0 that contains a pair of antipodal vertices $v, v^T \in V$ (the antipodal connected component K_0^T of G^1 then also contains v, v^T). In other words, there exists a pair of antipodal vertices v, v^T that are connected by a monochromatic path. This conjecture has been verified by hand for $n \leq 5$.

A (not necessarily antipodal) labelling of H_n is called *simple* if it does not contain a square $xyzt$ such that $xy, zt \in E_0$ and $yz, tx \in E_1$. In this paper we verify the conjecture for simple labellings that are not necessarily antipodal. The main lemma shows that if G_0, G_1 is a simple labelling of H_n, K_0 is a connected component of G_0, and K_1 is a connected component of $G' = (V, E \setminus E(K_0))$, then for every pair of vertices $v, v' \in V(K_0) \cap V(K_1)$, the two vertices v, v' are in the same connected component of G_1.

An *even labelling* of H_n is a partition of the set $V' \subseteq V$ of vertices with an even number of 1s into two subsets V_0, V_1, and *induces* a labelling of H_n by letting E_0, E_1 be the edges of E that have an endpoint in V_0, V_1 respectively. Even labellings induce simple labellings, so the conjecture follows for labellings induced by even labellings.

*268 Waverley Street, Palo Alto, CA 94301. Email: tomas@theory.stanford.edu

Finally we conjecture that even if the given labelling is not antipodal, there is always a path joining antipodal vertices that switches labels at most once, and show that this conjecture generalizes the conjecture on antipodal labellings. We establish this new conjecture for $H_n, n \le 5$ as well.

2 A Property of Simple Labellings of the Hypercube

In this section we establish a basic property of simple labellings of H_n.

Theorem 1 *If G_0, G_1 is a simple labelling of H_n, K_0 is a connected component of G_0, and K_1 is a connected component of $G' = (V, E \setminus E(K_0))$, then for every pair of vertices $v, v' \in V(K_0) \cap V(K_1)$, the two vertices v, v' are in the same connected component of G_1.*

Proof. Let $G'_0 = (V, E \setminus E(K_1))$ and let $G'_1 = (V, E(K_1))$. Then G'_0, G'_1 is a labelling of H_n, and both G'_0 and G'_1 have exactly one connected component that does not consist of just an isolated vertex. We shall refer to the edges in $E(G'_0)$ and $E(G'_1)$ as 0-*edges* and 1-*edges* respectively. We shall refer to the 1-edges that are in $E(K_1) \cap E(G_1)$ as *true* 1-*edges* respectively. Notice that if a 0-edge e shares an endpoint with a 1-edge e', then e' is a true 1-edge.

Let P_0 be a path $v = v_0, v_1, \ldots, v_\ell = v'$ of 0-edges in K_0 and let P_1 be a path $v' = v_\ell, v_{\ell+1}, \ldots, v_k = v$ of 1-edges. We need to find a path P'_1 from v' to v of true 1-edges. We shall represent the problem as a topological problem in the plane. Assume v is the all-0s bit string, and for any vertex w in H_n, let $|w| = d(v, w)$ denote the number of 1s in the bit string for w. For $0 \le i \le k$, represent v_i in the plane by the point with coordinates $(i, |v_i|)$, with $i + |v_i|$ even. In particular v_0 and v_k are represented by $(0,0)$ and $(k, 0)$, with k even, while for $0 < i < k$, the vertex v_i is represented by $(i, |v_i|)$ with $|v_i| \ge 1$. The edges $v_i v_{i+1}$ are represented by diagonal segments $(i, |v_i|)(i+1, |v_{i+1}|)$ where $|v_{i+1} = |v_i| \pm 1$.

We let $S = \{(i, j) : 0 \le i \le k, 0 \le j \le |v_i| \text{ and } i + j \text{ is even}\}$. Note that for all v_i, the point $(i, |v_i|)$ representing v_i is in S. We shall assign a vertex in H_n to each point in S, extending the assignment of v_i to $(i, |v_i|)$, and in a way such that if (i, j) is assigned w then $|w| = j$, and if $(i, j)(i \pm 1, j - 1)$ are assigned w, w' respectively, then $ww' \in E(H_n)$. We assign (i, j) to meet these requirements in decreasing order of j. Suppose we have assigned all $(i', j') \in S$ with $j' > j$ and wish to assign (i, j). In particular, we have assigned $(i-1, j+1), (i, j+2), (i+1, j+1)$ a path w', w'', w''' with $|w'| = |w''| - 1 = |w'''| = j + 1$. If $w' \ne w'''$, then there is a unique vertex w such that w, w', w'', w''' is a square in H_n, and we assign w to (i, j), with $|w| = j$. If $w' = w'''$, then we assign to (i, j) any w such that $ww' \in E(H_n)$ and $|w| = j$. At the last stage, this process assigns, to all $(i, 0)$ with $0 \le i \le k$ and i even, the vertex v corresponding the the all-0s bit string (recall k is even).

The path $v_0, v_1, \ldots, v_\ell, v_{\ell+1}, \ldots, v_k$ is thus represented in the plane by a non-crossing closed curve C that follows the path of $(i, |v_i|)$ from $(0, 0)$ to $(k, 0)$, and then back along the x-axis from $(k, 0)$ to $(0, 0)$. We wish to find a path of true 1-edges from $v' = v_\ell$ to v by a path in the plane from $(\ell, |v_\ell|)$ to some $(i, 0)$ with $0 \le i \le k$ and i even that moves at each step from some $(i', j') \in S$ to some $(i' \pm 1, j' \pm 1) \in S$.

Notice that the non-crossing closed curve C goes along 0-edges from $(0, 0)$ to $(\ell, |v_\ell|)$, along 1-edges from $(\ell, |v_\ell|)$ to $(k, 0)$, then from $(k, 0)$ back to $(0, 0)$ along the x-axis. We shall gradually shrink the curve C to obtain non-crossing closed curves D that successively have a smaller area inside. For the induction, suppose we have obtained a closed curve D, initially $D = C$. We allow the curve D to touch itself on the outside, but not on the inside. If D contains some point on the x-axis, then it contains just one segment $(i',0)$ back to $(i,0)$ on the x-axis, with $i \leq i'$ even, and D goes from $(i,0)$ to $(i',0)$ by following segments $(i,j)(i\pm 1, j\pm 1)$, first along 0-edges to get from $w = v$ represented by $(i,0)$ to w' represented by some $(i,|w'|)$, and then along 1-edges to get from w' to $w = v$ represented by $(i',0)$. If D contains no point on the x-axis, then it ony follows segments $(i,j)(i\pm 1, j\pm 1)$, first along 0-edges to get from w to w' and then along 1-edges to get from w' to w. The aim is to get from w' to w inside D by following true 1-edges.

Let P be the path of 0-edges from w to w', and let Q be the path of 1-edges from w' to w, along D. We shall find a square R inside of D given by $xyzt = (i-1,j)(i,j+1)(i+1,j),(i,j-1)$ such that w' is one of x,y,z,t, say $w' = y$, with yx a 0-edge and yz a 1-edge, so that the path of 0-edges from w to $w' = y$ to x and the path of 1-edges from z to $y = w'$ to w meet at w' but do not cross. The edge yz is a true 1-edge, as it meets a 0-edge. The edges zt and tx cannot be a 0-edge and a 1-edge respectively simultaneously, since the labelling is simple. The series of edges along $yxtzy$ are thus 0-edges followed by 1-edges according to the possible sequences $0001, 0011, 0111$, with all the 1-edges in any of these sequences being true 1-edges. Thus following P, then $yxtzy$, then Q, gives a smaller closed curve D' for the induction.

The only difficulty for the induction occurs if one or more of x,t,z are already in D, as the curve is not allowed to touch itself inside. Let the path from w' back to w' along Q,w,P be given by four segments $T_1, z', T_2, t', T_3, x', T_4$, possibly satisfying any subset of the conditions $z = z', t = t', x = x'$.

The cases of the three sequences are similar, so we shall assume the sequence is 0001. The worst case is $z = z', t = t', x = x'$, so we shall assume these three equalities hold as well. If w occurs in T_1 then consider the cycle w', T_1, z for the induction, and we shall find a path of true 1-edges from z to w inside the cycle. If $w = y$ we are done. If w occurs in T_2, t, then consider the cycle z, T_2, t for the induction, and we shall find a path of true 1-edges from z to w inside the cycle. If w occurs in T_3, x, T_4, we first consider the cycle z, T_2, t for the induction and find a path of true 1-edges from z to t inside the cycle. If w occurs in T_3, x, we finish by considering the cycle t, T_3, x and find a path of true 1-edges from t to w inside the cycle. Finally if w occurs in T_4, we continue with the cycle t, T_3, x by finding a path of true 1-edges from t to x, and finish with the cycle x, T_4, w' by finding a path of true 1-edges from x to w.

It remains to find the square R. Suppose P enters w' along the 0-edge uw', and Q leaves w' along the 1-edge $w'u'$. Up to rotation, there are three cases: (1) $uw'v$ is $(i-1,j)(i,j+1)(i+1,j)$ with D underneath w'; then we may set $xyzt = uw'vt$ with $t = (i,j-1)$. (2) $uw'v$ is $(i-1,j-1)(i,j)(i+1,j+1)$ with D underneath w'; then if $e = (i,j)(i+1,j-1)$ is a 0-edge we may set $xyzt = (i+1,j-1)(i,j)(i+1,j+1)(i+2,j)$, and if e is a 1-edge we may set $xyzt = (i-1,j-1)(i,j)(i+1,j-1)(i,j-2)$. (3) $uw'v$ is $(i-1,j+1)(i,j)(i+1,j+1)$ with D underneath w'; then if $e' = (i,j)(i-1,j-1)$ is a 0-edge and $e'' = (i,j)(i+1,j-1)$ is a 1-edge we may set $xyzt = (i-1,j-1)(i,j)(i+1,j-1)(i,j-2)$, if e' is a 1-edge we may set $xyzt = (i-1,j+1)(i,j)(i-1,j-1)(i-2,j)$, and finally if e'' is a 0-edge we may set $xyzt = (i+1,j-1)(i,j)(i+1,j+1)(i+2,j)$. This completes the induction for shrinking D. □

3 Antipodal Paths in Simple Labellings of the Hypercube

In this section we first prove the conjecture for simple antipodal labellings of $H_n, n \geq 2$, and then extend the result to simple labellings of H_n that are not necessarily antipodal. The proofs are direct applications of Theorem 1.

Theorem 2 *For every simple antipodal labelling G_0, G_1 of H_n, $n \geq 2$, there exists a connected component K_0 of G_0 that contains a pair of antipodal vertices $v, v^T \in V$ (the antipodal connected* *component K_0^T of G_1 then also contains v, v^T).*

Proof. Let P be a path from v to the antipodal vertex v^T, and let P^T be the antipodal path to P from v^T back to v. Consider the cycle Q consisting of P followed by P^T. The cycle Q consists of k sequences of edges labelled 0 and k sequences of edges labelled 1 in alternation. The cycle Q can thus be viewed as an alternation of components $A_1, B_1, A_2, B_2, \ldots, A_k, B_k$, where the A_i are labelled 0 and the B_i are labelled 1. We have that A_1 is antipodal to B_ℓ, for $\ell = (k+1)/2$, so k must be odd.

Suppose two A_i are the same. We may assume without loss of generality that these A_i are A_1 and A_i for $i \leq \ell$. The sequence can then be replaced by $A_1 = A_i, B_i, A_{i+1}, B_{i+1}, \ldots, A_\ell, B_\ell = B_{\ell+i-1}, A_{\ell+i}, B_{\ell+i}, \ldots, A_k, B_k$, thus reducing k to $k' = k - 2(i-1)$. Repeating this reduction enough times we eventually have that all components in the alternation $A_1, B_1, A_2, B_2, \ldots, A_k, B_k$, are distinct.

Suppose $k \geq 2$, and consider the statement of Theorem 1. Letting $K_0 = A_1$, we can choose K_1 containing the union of $B_1, A_2, B_2, \ldots A_k, B_k$. Letting v, v' be vertices in $V(A_1) \cap V(B_1)$ and $V(A_1) \cap V(B_k)$ respectively, we infer that v and v' are in the same connected component of G_1, so $B_1 = B_k$, contrary to distinctness. Therefore $k = 1$, the alternation is A_1, B_1 with A_1, B_1 antipodal, so there is a pair v, v^T of antipodal vertices contained in both $K_0 = A_1$ and $K_0^T = B_1$. □

Theorem 3 *Every (not necessarily antipodal) simple labelling of H_n has a pair of antipodal vertices joined by a monochromatic path.*

Proof. We proceed by induction on n. Suppose the result is true for H_n, and consider H_{n+1}. The cube H_{n+1} can be decomposed into two subcubes $0H_n$ and $1H_n$, and by inductive hypothesis $0H_n$ has a pair of antipodal vertices $0x$ and $0x^T$ joined by a monochromatic path P. Say the edges of P have label 0. If either of the edges $(0x, 1x)$, $(0x^T, 1x^T)$ has label 0, then we obtain antipodal vertices $1x, 0x^T$ or $0x, 1x^T$ joined by a monochromatic path labelled 0, by extending P. So suppose both of these edges have label 1 for the remaining case. These two edges are in components joined by edges of label 1, call these components C and D respectively. If C and D are the same component of label 1, then $0x, 1x^T$ are antipodal vertices joined by a path of label 1, and the result follows. Otherwise by Theorem 1 we have that every path from C to D must traverse an edge in the component E labelled 0 of the path P, otherwise components C and D in the boundary of E would have to be the same component. Let P^T be the path from $1x^T$ to $1x$ antipodal to P. The path P^T joins C and D, and therefore must go through an edge of E. In particular, P^T contains a vertex y^T of E, with y in P and thus also in E. But then a path labelled 0 joins the antipodal vertices y and y^T through E, completing the proof. □

From Theorem 3 we obtain the following general class satisfying the conjecture.

Corollary 1 *Even labellings of H_n induce simple labellings, so the conjecture follows for labellings induced by even labellings (that are not necessarily antipodal).*

Proof. If the labelling is not simple, then it contains a square $xyzt$ with xy, zt having label 0 and yz, tx having label 1, contrary to the assumption that every vertex with an even number of 1s in the string has all edges coming out of it with the same label. Thus the labelling is simple and the conjecture follows in this case from Theorem 3. □

4 On the Occurrences of the Special Square

We earlier proved the conjecture in the case where the square $xyztx$ with edges 0101 in that order was forbidden. We refer to an ocurrence of this square as a *special square.* We first prove a preliminary result.

Proposition 1 *Suppose the conjecture is not true in n dimensions. Then it is also not true in $n+1$ dimensions, and so on by induction.*

Proof. Suppose we have a counterexample of a labelled n-dimensional hypercube H_n. Making two copies of the labelled H_n and joining them arbitrarily but antipodally gives a counter example H_{n+1}, as an antipodal 1-edge path in H_{n+1} would induce such a path in H_n. □

Theorem 4 *Consider a counter example to the conjecture with a minimal number of special squares in an even dimension n. Let $xyztx$ be a special square in such a counter example of the conjecture, with edges labelled in turn* 0101. *Then the 1-edge components of yz and tx contain some antipodal vertices u and u^T respectively. Furthermore each component of 0-edges must contain along its boundary some special square if we consider counter examples that have a minimal number of 0-components with respect to those having a minimal number of special squares.*

Proof. If no antipodal u, u^T occur, then we can change the label of xy from 0 to 1, joining the two components. If this flip produces one or more extra special squares, we may flip again the opposite sides, from 1 to 0, and so on. As n is even, the parities of the flips from 0 to 1 and from 1 to 0 are different, so the process will eventually end, with one less special square. In the end, all special squares will join components that contain an antipodal pair of vertices u, u^T respectively.

Suppose a component A_1 of 0-edges does not contain a special square along its boundary. Suppose we flip A_1 to a component of 1-edges and it is no longer a counter example. Then there must be two 1-components B_1 and B_3 sharing at least one vertex with A_1 and containing antipodal vertices u, u^T respectively. Then as in Theorem 2 we obtain an antipodal cycle of components $A_1, B_1, A_2, B_2, A_3, B_3$. If A_1 is A_2 or A_3, say A_2, then the sequence $v, A_1, B_1, A_1, v^T, B_2, A_3, B_2$ shows that A_1 contains a pair of antipodal vertices v, v^T, contrary to having a counter example. If all A_i are distinct, then as in Theorem 2 by the proof of Theorem 1, we must have $B_1 = B_3$, since there is no special square in the boundary of A_1. This gives the sequence $A_1, B_1, w, A_2, B_2, A_2, w^T, B_1$, showing that A_2 contains a pair of antipodal vertices w, w^T, contrary to having a counter example. So we may flip the component A_0 to 1-edges, reducing the number of components. □

This proof indicates that a counter example for $n \geq 6$ would have to have at least four 1-edge components. Our attempt to construct such a counter example with $n = 6$ and four 1-edge components with 16 vertices each failed, so the cases with $n \geq 6$ remain open. It can be shown by similar methods that any counter example with n even must have at least two antipodal pairs of special squares.

5 1-Switch Path Joining Antipodal Vertices

Consider a labelled n-dimensional hypercube, where the labelling is not necessarily antipodal, and a path P in it joining antipodal vertices. We say that P is a k-switch path for some $k \geq 0$ if P

is the concatenation of at most $k+1$ monochromatic paths (so that the path P switches colors at most k times).

We restate the main conjecture of [1], and a new conjecture that implies it.

Conjecture 1 *Any antipodal labelling of the n-dimensional hypercube always has a 0-switch path joining antipodal vertices.*

Conjecture 2 *Any labelling of the n-dimensional hypercube always has a 1-switch path joining antipodal vertices.*

Theorem 5 *If the labelling is antipodal, then the existence of a $2k+1$-switch path joining antipodals for some $k \geq 0$ implies the existence of a $2k$-switch path joining antipodals as well. In particular, Conjecture 2 implies Conjecture 1.*

Proof. If the path starts with the monochromatic path x, y and ends with the monochromatic path z, x^T, then these two paths have different colors, so we can remove the path x, y and extend an antipodal path to obtain z, x^T, y^T, thus reducing the number of switches by one. □

Theorem 6 *The 1-switch conjecture holds for arbitrary labelings of the n-dimensional hypercube if $n \leq 5$, and via isometric paths (a path is isometric if all of its edges traverse different dimensions). Consequently, the 0-switch conjecture for antipodal labelings of the n-dimensional hypercube if $n \leq 5$, and via isometric paths.*

Proof. The 0 switch conjecture If $n \leq 2$ then a shortest antipodal path has at most 2 edges and thus at most 1 switch. If $n = 3$, then any vertex, with three incident edges, has two of these edges of the same label, so we obtain a monochromatic path of length 2 that can be extended to an antipodal isometric path of length 3 with at most one switch.

If $n = 4$, consider first the case where the label at each vertex is two 0s and two 1s. If one of these cycles of 0s or 1s is of length 6 or more then we obtain a monochromatic path of length 3 in the cycle joining two vertices at distance 3, and this path can be extended to an antipodal isometric path of length 4 with only one switch. Otherwise we have four 0-squares and four 1-squares, and we obtain an antipodal isometric path consisting of two 0-edges followed by two 1-edges, with only one switch.

In the remaining case for $n = 4$ we have a vertex that has three incident edges of the same label, say $(0000, 0001), (0000, 0010), (0000, 0100)$ are all labelled 0-edges. If we can also come out of one of $0001, 0010, 0100$ with a 0-edge, the we obtain a path of length 3 joining vertices at distance 3 with a 0-label, which can be extended to an antipodal isometric path of length 4 with only one switch. Otherwise we have a path $1010, 0010, 0011, 0001, 0101$ of 1-edges, that is, a monochromatic antipodal path.

If $n = 5$, the proof follows from Lemma 1 and Lemma 2 below. We may by Lemma 1 assume that 00000 and 11100 are joined by an isometric monochromatic path, say of label 0. This path can be extended to an antipodal isometrc 1-switch path unless the edges $(00000, 00010)$ and $(11100, 11110)$ have label 1 and the edges $(00010, 00011)$ and $(11110, 11111)$ have label 0. But then, if 00010 and 11110 are joined by an isometric 1-switch path, say going from label 0 to label 1, then adding the edges $(00010, 00011)$ of label 0 and $(11110, 11100)$ gives an isometric antipodal 1-switch path from 00011 to 11100 (the case going from label 1 to label 0 uses the other two edges that had to have label 1 and label 0). If 00010 and 11110 are not joined by an isometric 1-switch path, then we have three antipodal monochromatic paths in the 3-cube $***10$ by Lemma 2. But then, we could have taken the 3-cube $***00$ to contain three antipodal isometric monochromatic paths, not just the isometric monochromatic path joining 00000 and 11100, and only one of these antipodal monochromatic paths in the 3-cube $***00$ can have a corresponding antipodal pair in the 3-cube $***10$ that is not joined by an isometric 1-switch path by Lemma 2, so the proof when the corresponding antipodalpair in the 3-cube $***10$ is joined by ann isometric 1-switch path is completed as in the previous case. □

Lemma 1 *An arbitrary labelling of H_5 has an isometric monochromatic path of length 3.*

Proof. The degree of each vertex is 5, so the average degree of vertices in the graph of 0 labels or in the graph of 1 labels is at least 2.5, say of 1 labels. In particular, some connected component G of 1 labels has average degree at least 2.5. Then G must contain a vertex v of degree at least 3, and G cannot be a star rooted at v (otherwise the average degree would be less than 2). Extending in G a path of length 2 starting with v with one of the other two edges coming out of v gives an isometric path of label 1 in G, namely u_0, v, u_1, u_2. □

Lemma 2 *A labeling if H_3 that does not have an isometric 1 switch path joining some antipodal pair of vertices, has isometric monochromatic paths joining the other three antipodal pairs of vertices.*

Proof. Suppose 000 and 111 are not joined by an isometric 1 switch antipodal path. If the two edges $(000, 001)$ and $(000, 010)$ do not have the same label, then one of the two paths $000, 001, 011, 111$ and $000, 010, 011, 111$ is a 1-switch path. Therefore all edges coming out of 000 and 111 must have the same label, and the remaining hexagon $001, 011, 010, 110, 100, 101$ has the other label. This hexagon gives isometric monochromatic paths joining the three antipodal pairs other than $000, 111$. □

We say that a labelling of the n-dimensional hypercube is $(r, n-r)$-regular for some $0 \leq r \leq n$ if each vertex has r 0-edges and $n-r$ 1-edges.

Lemma 3 *Each 0-component of an $(r, n-r)$-regular labelling is either an r-dimensional subcube or a component containing two vertices at distance at least $r+1$.*

Proof. We first show that the 0-component contains two vertices at distance r. If x, y in the 0-component are at distance q less than r, then y has $q < r$ neighbors closer to x, so some z neighbor of y in the 0-component must be at distance $q+1$ from x by $(r, n-r)$-regularity. The claim thus follows by induction on q.

Suppose x, y are at distance r in the same 0-component. Then the r neighbors of x, y in the 0-componenent must be the neighbors in the r-dimensional subcube containing x, y, otherwise we have such neighbors of x at distance $r+1$ of y or viceversa. Arguing similarly for the neighbors of x and their antipodal neighbors of y at distance r in this r-dimensional subcube, we obtain all the neighbors of neighbors of x or y in the same 0-component as well, until we obtain the whole r-dimensional subcube containing x, y as the whole 0-component, or obtain distance $r+1$ for the 0-componenent. □

Theorem 7 *If a labelling is either $(n-1,1)$-regular or $(n-2,2)$-regular, then the 1-switch conjecture holds.*

Proof. For an $(n-1,1)$-regular labelling, either we have a 0-component that is a $(n-1)$-dimensional subcube, giving antipodal vertices at distance $n-1$ in the 0-subcube that can be extended to antipodal vertices in the whole cube by just one switch, or some 0-componenent already contains vertices at distance n, thus antipodal.

For an $(n-2,2)$-regular labelling, either we have a 0-component with two vertices at distance $n-1$, which can be extended with just one switch to antipodal vertices at distance n, or with have four $(n-2)$-dimensional 0-cubes, in which case two aditional 1-edges give a 1-switch antipodal path. □

Note that the Lemma, the Theorem and their proofs also hold if we assume degrees at least r instead of exactly r for the 0-component in the Lemma, and degrees at least $n-2$ for the 0-components in the Theorem, so that the 1-components are either vertices, paths, or cycles.

We may extend the problem by considering labellings with $r \geq 2$ possible labels (instead of just $r = 2$ with labels 0 and 1). In that case, we may conjecture the existence of $(r-1)$-switch antipodal paths. Notice that if such antipodal paths exist, there necessarily exist monochromatic paths joining vertices at distance $\lceil n/r \rceil$.

The following bound shall be improved later.

Theorem 8 *For any n, r, there exist monochromatic paths joining vertices at distance $\lceil n/(2r) \rceil$.*

Proof. Let $N = 2^n$ be the number of vertices. Then one of the r labels occurs in $Nn/(2r)$ edges, say label i. Then some i-component has K vertices and $Kn/(2r)$ i-edges. If dimension j is used s_j times by this component, then $\sum_{1 \leq j \leq n} s_j \geq Kn/(2r)$ and $\sum_{1 \leq j \leq n} (s_j/K) \geq n/(2r)$. If we fix a vertex x in this i-component and select a random vertex y in the same i-component, then with probability at least s_j/K dimension j will be flipped from x to y, since there are at least s_j vertices on either side of dimension j in this i-component. Summing this probability over all j we infer an expected number of flipped dimensions at least $n/(2r)$ from x to y, in particular some y achieves $\lceil n/(2r) \rceil$ flips. □

Notice that if we partition the dimensions almost evenly among the r labels, the bound above on distance between vertices in monochromatic paths can only be improved from $\lceil n/(2r) \rceil$ to $\lceil n/r \rceil$. We carry through this improvement in detail.

Let G be a connected induced subgraph of the n-dimensional hypercube H_n. Let h be the maximum Hamming distance between vertices of G in H_n. Let $d = 2e/v$ be the average degree of vertices in G, where v and e are the number of vertices and edges in G.

Theorem 9 *If G is a connected induced subgraph of H_n, then $h \geq d$, with equality only if G is a subcube H_d.*

Note that if G is a subcube H_d, then $h = d$. We first prove the first statement of the theorem.

We say that $x \leq y$ for $x, y \in V(H_n)$ if $x_i \leq y_i$ for $1 \leq i \leq n$, and we say that $x < y$ if $x \leq y$ and $x \neq y$. We say that G as above is an *ideal* if whenever $x \leq y$ and $y \in V(G)$, we have $x \in V(G)$.

Lemma 4 *We may assume w.l.o.g. that G is an ideal.*

Proof. If G is not an ideal, then for some $1 \le i \le n$, there exist $x \in V(G)$ and $y \in V(H_n) \setminus V(G)$ such that $x_i = 1, y_i = 0$, and $x_j = y_j$ for $j \ne i$, $1 \le j \le n$. Replacing every such x by the corresponding y, we obtain a new G' such that $h' \le h$ and $d' \ge d$. Repeating this for all $1 \le i \le n$, we obtain an ideal G^n such that if the theorem holds for G^n then it also holds for G, since $h' \ge d'$ implies $h \ge h' \ge d' \ge d$. □

Let $|x|$ be the number of x_i equal to 1. A maximal vertex in an ideal G is a vertex $x \in V(G)$ such that $y \notin V(G)$ for all $y > x$, $y \in H_n$.

Lemma 5 *We may assume w.l.o.g. that if x is a maximal vertex in an ideal G, then $|x| > d/2$.*

Proof. If $|x| \le d/2$, then removing x from G to obtain G' gives $d' = 2e'/v' \ge (2e - d)/(v - 1) = (dv - d)/(v - 1) = d$ and $h' \le h$, so if $h' \ge d'$ then $h \ge h' \ge d' \ge d$. □

Lemma 6 *We may assume w.l.o.g. that if x is a vertex in an ideal G with $x_1 = 0$ and $x_i = 1$, then y with $y_1 = 1$, $y_i = 0$ and $y_j = x_j$ for $j \ne 1, i$, $1 \le j \le n$ is such that y is also in G.*

Proof. If the condition in the lemma does not hold for some x, y, i then replacing all such x by the corresponding y gives G' satisfying $d' \ge d$ and $h' \le h$, so if $h' \ge d'$ then $h \ge h' \ge d' \ge d$. □

Let $z = \max(x, y)$ for $x, y \in V(H_n)$ be given by $z_i = \max(x_i, y_i)$.

Lemma 7 *h is the maximum over all pairs x, y of maximal vertices in the ideal G of $|\max(x, y)|$.*

Proof. The maximum distance between vertices of G cannot be more than the number of 1s in either of the two vertices, thus at most $|\max(x, y)|$ for some maximal vertices x, y. Conversely, if whenever $x_i = y_i = 1$ we change y_i to $y'_i = 0$, then the distance between x and y' is $|x| + |y'| = |\max(x, y')| = |\max(x, y)|$. □

Lemma 8 *We may assume w.l.o.g. that if x, y are maximal vertices, then $x_i = y_i = 1$ for some i.*

Proof. If not, then $|x|, |y| > d/2$ by Lemma 5, so $h \ge |\max(x, y)| = |x| + |y| > d$ and the theorem follows. □

Lemma 9 *We may assume w.l.o.g that $x_1 = 1$ for every maximal vertex in an ideal G.*

Proof. If $x_1 = 0$, then switching $x_1 = 0, x_i = 1$ with $y_1 = 1, y_i = 0, y_j = x_j$ for $j \ne 1, i$ gives y also in G for any such i by Lemma 6, and adding z with $z_1 = 1, z_j = x_j$ for $j \ne 1$ gives G' with $d' = 2e'/v' > (2e + d)/(v + 1) = (dv + d)/(v + 1) = d$ by Lemma 5, and $h' = h$ since for all maximal vertices u, we have $|\max(z, u)| = |\max(t, u)|$, where t differs from z in just some i with $z_i = u_i = 1$ and $t_i = 0$ by Lemma 8. So if $h' \ge d'$ then $h = h' \ge d' \ge d$. □

Now by Lemma 9, G is the union of two identical graphs G' with $x_1 = 0$ and G'' with $x_1 = 1$, and assuming inductively that the theorem holds for G' in H_{n-1} we have $h = h' + 1 \ge d' + 1 = d$. This completes the proof of the first statement of Theorem 9. For the second statement, note that for each of the modifications applied in the lemmas, either $d' > d$, in which case $h \ge h' \ge d' > d$, or the resulting graph is not a subcube if the original graph was not a subcube. Finally, the result follows by induction on the two identical G' and G''.

Corollary 2 *If the edges of H_n are colored with r different colors $0 \leq i < r$, and h_i is the maximum Hamming distance in components of the subgraph induced by color i, then $\sum_{0\leq i<r} h_i \geq n$. In particular, for some $0 \leq i < r$, $h_i \geq \lceil n/r \rceil$.*

Proof. If the number of edges of color i is $\alpha_i n 2^{n-1}$ with $\sum_{0\leq i<r} \alpha_i = 1$, then the average degree for color i is $\alpha_i n$, and in particular some component of color i has average degree $d_i \geq \alpha_i n$, so by the theorem we have $h_i \geq d_i$ and

$$\sum_{0\leq i<r} h_i \geq \sum_{0\leq i<r} d_i \geq \sum_{0\leq i<r} \alpha_i n = n.$$

□

In particular, in the case with just two colors, we have $h_0 + h_1 \geq n$. It has been conjectured for this case that there always exists a pair of antipodal vertices joined by a path of color 0 followed by a path of color 1, which implies, if the coloring gives opposite colors to antipodal edges, that there exists a monochromatic path joining antipodal vertices. This can be verified if the number of edges of color 1 is strictly less than 2^n, since in this case $d_1 < 2$ (because the total number of edges is $n2^{n-1}$), so $d_0 > n - 2$ and therefore $h_0 \geq n - 1$, giving a path of color 0 of endpoint distance n-1 that can be extended by an edge of color 0 or 1 to a path joining antipodal vertices.

References

[1] M. Devos and S. Norine, Edge-antipodal colorings of cubes, Open Problem Garden, `http://garden.irmacs.sfu.ca/?q=op/edge antipodal colorings of cubes`