

Applying a Result of Frankl and Rödl to the Construction of Steiner Trees in the Hypercube

by

Zevi Miller and Dan Pritikin
Department of Mathematics and Statistics
Miami University
Oxford, Ohio 45056

Abstract

Let $Q(n)$ be the n -dimensional hypercube, and X a set of points in $Q(n)$. The Steiner problem for the hypercube is to find the smallest number $L(n,X)$ of edges in any subtree of $Q(n)$ which spans X . Let $W(k,n)$ be the set of points in $Q(n)$ having weight k , where we normalize $k+1 \leq \frac{n}{2}$. We apply a result of Frankl and Rödl on the generalized Turan problem for hypergraphs to show that

$$L(n,W(k+1,n)) \leq \binom{n}{k+1} + (1+o(1)) \left(\frac{\log(k)}{k} \right) \binom{n}{k} \text{ as } k \rightarrow \infty.$$

We also show that this is within a factor of $\log(k)$ from optimal in a certain strong sense.

1. Introduction

Let $G = (V,E)$ be an undirected graph with vertex set V and edge set E . The Steiner problem for G (first introduced in [H]) is: given a subset X of V find the minimum number $L(X,G)$ of

edges in any subtree of G which contains X among its vertices. Apart from its intrinsic interest as a graph theory problem, the Steiner problem is motivated by various layout problems in VLSI design and in communication networks.

This problem has been considered for various classes of graphs G (see the survey [Wi]). In this paper we continue a study, begun in [MP], of the Steiner problem in the case where G is the n -dimensional hypercube $Q(n)$. Among the results in the latter paper were "formulas" for $L(X, Q(n))$ when $|X| \leq 5$, and NP-completeness even in the subproblem when X is only permitted to be an arbitrary subset of the points of weight 2 in $Q(n)$.

By definition, $Q(n)$ is the graph whose point set is the set of all 2^n n -tuples over the alphabet $\{0,1\}$. Two points of $Q(n)$ form an edge if and only if the corresponding n -tuples differ in exactly one coordinate. The weight of a point is the number of 1's among its coordinates. We let $W(k,n)$ be the set of all points in $Q(n)$ of weight k . As more convenient notation, given any set $X \subset Q(n)$, we let $L(n,X) = L(X, Q(n))$.

One of the results in [MP] was the bound

$$L(n, W(k+1,n)) \leq \binom{n}{k+1} + (2+o(1)) \left(\frac{\log(k)}{k} \right) \binom{n}{k}$$

under the restriction that k is at most roughly $O(n^{1/2})$. That result made use of a result of Frankl and Rödl [FR] related to the generalized Turán problem for hypergraphs. Their result involves the construction of a set $D \subset W(k,n)$ of size at most $(1+o(1)) \left(\frac{\log(k)}{k} \right) \binom{n}{k}$ (as $k \rightarrow \infty$) which dominates $W(k+1,n)$; that is, such that for every $w \in W(k+1,n)$ there is a $v \in D$ for which vw is an edge of $Q(n)$. In this report we obtain the improved bound

$$L(n, W(k+1,n)) \leq \binom{n}{k+1} + (1+o(1)) \left(\frac{\log(k)}{k} \right) \binom{n}{k}$$

with k restricted only by $k+1 \leq \lfloor \frac{n}{2} \rfloor$. When $k+1 > \lfloor \frac{n}{2} \rfloor$ the symmetry of the hypercube then implies the dual bound obtained by replacing each occurrence of k in the second term of the right side by $k+2$. We make an analysis of the connected components structure in the subgraph of $Q(n)$ induced by $D \cup W(k+1,n)$. The results of this analysis imply that by adding to D a set E of negligibly smaller size, the resulting set $D \cup E \cup W(k+1,n)$ induces a connected subgraph of $Q(n)$. The improved bound follows.

In the last section we show that this bound is in a certain strong sense within a factor of $\log(k)$ from optimal.

2. The Upper Bound

In this section we apply a result of Frankl and Rödl [FR] on the generalized Turán problem to get an upper bound for $L(n, W(k+1, n))$.

Observe first that a trivial upper bound for $L(n, W(k+1, n))$ is obtained by noting that the subgraph of $Q(n)$ induced by $W(k+1, n) \cup W(k, n)$ is connected. It follows that $L(n, W(k+1, n)) \leq |W(k+1, n)| + |W(k, n)| - 1 = \binom{n}{k+1} + \binom{n}{k} - 1$. Of course one would expect to improve on this fairly crude bound. A better approach is to use a result of [FR] that provides a subset $Z \subset W(k, n)$ which dominates $W(k+1, n)$ and has size $|Z| \leq (1+o(1)) \left(\frac{\log(k)}{k}\right) \binom{n}{k}$ as $k \rightarrow \infty$. The hope is then that by adding not too many more points to Z we might obtain a set Z' such that the graph induced by $Z' \cup W(k+1, n)$ is connected. Indeed, the main result of this section is that such a Z' exists having an additional number of points which is negligible compared to the bound for $|Z|$ given above.

As background, recall the Turán problem for graphs. This is to find the maximum number of edges a graph G on n points can have so that G contains no subgraph isomorphic to K_t (the complete graph on t points). This problem was solved in [Tu]. Equivalently one may ask for the smallest size of a set T of edges in K_n such that every K_t subgraph of K_n contains some edge from T .

A possible generalization of the Turán problem to hypergraphs, coming from the case $t=3$ where we are covering triangles (i.e. 3-sets) by edges (2-sets), is as follows. What is the minimum size $C(n, k)$ among all collections T of k -sets from a ground set X on n points so that every $(k+1)$ -subset of X contains at least one element from T ? The connection with hypercubes is then natural. First we view the elements of X as the coordinate positions used in describing the points of $Q(n)$. Each subset G of X then corresponds to a point of $Q(n)$, which we denote $\underline{p}(G)$, having 1's in the coordinates belonging to G and 0's in the remaining coordinates. Similarly for a collection T of subsets of X we let $\underline{p}(T)$ be the subset $\bigcup_{G \in T} \underline{p}(G)$ of the point set of $Q(n)$. Now $C(n, k)$ may be viewed as the minimum size of a set $V \subset W(k, n)$ which dominates $W(k+1, n)$, i.e. such that for any $w \in W(k+1, n)$ there is a $v \in V$ such that vw is an edge of $Q(n)$.

We now describe the construction of [FR]. The following notation will be needed. Let $X = \{1, 2, \dots, n\}$ be a ground set of n elements, and let $\binom{X}{m}$ denote the collection of m -subsets of X for any m , $1 \leq m \leq n$. Consider an integer $r < k$ to be chosen later, and assume for convenience that $n = rt$, where t is some positive integer. Now partition X as $X = X_0 \cup X_1 \cup X_2 \cup \dots \cup X_{r-1}$ where $X_i = \{c: 1 \leq c \leq n, c \equiv i \pmod{r}\}$. For any subset $G \subset X$, let $\underline{S}(G) = \{i: G \cap X_i \neq \emptyset\}$, and $\underline{s}(G) = |\underline{S}(G)|$. For any integer e , $0 \leq e \leq r-1$, define

$$V_e = \left\{ F \in \binom{X}{k}: e + \sum_{x \in F} x \equiv 0, 1, 2, \dots, \text{ or } r - \underline{s}(F) \pmod{r} \right\}.$$

Theorem 1: [FR] For any $0 \leq e \leq r-1$, the set $p(V_e)$ dominates $W(k+1,n)$.

Proof: We paraphrase the proof given in [FR].

Consider any e and any $G \in \binom{X}{k+1}$. It suffices to show that there exists a corresponding

$F \in V_e$ such that $F \subset G$. For any $g \in G$ let $y(g) = e + \sum_{x \in G - \{g\}} x$. The numbers $\{y(g) : g \in G\}$ form a

set of $s(G)$ distinct numbers modulo r . Hence by the pigeon hole principle there exists w such that $y(w) \equiv i \pmod{r}$ for some i satisfying $0 \leq i \leq r - s(G)$. The set $F = G - \{w\}$ then satisfies $F \in V_e$ and $F \subset G$, as required. ■

For any $0 \leq e \leq r-1$, let R_e be the subgraph of $Q(n)$ induced by $p(V_e) \cup W(k+1,n)$ (with V_e defined as above). Since $p(V_e)$ already dominates $W(k+1,n)$, the subgraph R_e is a good start towards constructing a subtree of $Q(n)$ which spans $W(k+1,n)$. Should R_e be connected, we would have the bound $L(n, W(k+1,n)) \leq |R_e| = \binom{n}{k+1} + |p(V_e)|$. But even when R_e turns out to be disconnected, a bound for $L(n, W(k+1,n))$ will be obtained by adding to R_e a modest number of points of $Q(n)$ so as to induce a connected subgraph. Therefore we introduce the following notation toward analyzing the structure of the connected components of R_e .

For a given r , $0 \leq e \leq r-1$, and a given $G \subset X$ with $|G| = k$, we define the congruence vector of G as the vector $\underline{C}(G) = (c_0, c_1, \dots, c_{r-1})$, where $c_i = |G \cap X_i|$ is the number of elements of G congruent to $i \pmod{r}$. For a congruence vector D we let the block of D be $\{H \subset X : C(H) = D\}$, and the term block will refer to any collection B of k -subsets from X such that B is the block of D for some congruence vector D . Again, such a B may also be viewed as a subset of $W(k,n)$. Abusing our notation somewhat, we will refer to a block B by the congruence vector D defining B when there is no chance of confusion. As an example with $k=4$ and $r=3$ and n large enough, we have the subset (of X) $G = \{3,6,4,8\}$ with $C(G) = (2,1,1)$ and associated block $B = (2,1,1)$ (consisting of all $H \subset X$ satisfying $C(H) = (2,1,1)$). For a block B we let $s(B)$ be the number of nonzero entries in the congruence vector defining B (so that $s(B) = 3$ for the block of the preceding example). Note that each of the sets V_e defined above partitions as a disjoint union of blocks since if $G \in V_e$ then any $H \in \binom{X}{k}$ satisfying $C(H) = C(G)$ must also belong to V_e . As k is arbitrary in these definitions, we can speak about blocks which are subsets of $W(t,n)$ for any t . Finally for any subset $M \subset W(k,n)$ we let $\underline{Nb}(M) = \{u \in W(k+1,n) : uv \in E(Q(n)) \text{ for some } v \in M\}$.

Lemma 1: For any block B , the subgraph of $Q(n)$ induced by $p(B) \cup \underline{Nb}(p(B))$ is connected.

Proof: Let $G, H \in B$. We will construct a sequence of "exchanges" leading from G to H . Corresponding to this sequence is a path in $Q(n)$ from $p(G)$ to $p(H)$ whose points alternate between $p(B)$ and $Nb(p(B))$.

Let $G \setminus H = \{x_1, x_2, \dots, x_c\}$ and $H \setminus G = \{y_1, y_2, \dots, y_c\}$ where $x_i \equiv y_i \pmod{r}$ for all $i \leq c$. Now let $G_0 = G$, and inductively let $G_i = G_{i-1} \cup \{y_{(i+1)/2}\}$ for i odd and $G_i = G_{i-1} \setminus \{x_{i/2}\}$ for i even, $1 \leq i \leq 2c$. An easy induction shows that $G_{2c} = H$, while $G_i \in B$ for i even, and $p(G_i)p(G_{i+1})$ is an edge of $Q(n)$ for all i . Hence $p(G_0)p(G_1)\dots p(G_{2c})$ is a path in $Q(n)$ from $p(G)$ to $p(H)$ alternating between $p(B)$ and $Nb(p(B))$, as desired. ■

The preceding lemma motivates the following approach to clarifying the connected components structure in the graphs R_e . Let $B = (c_0, c_1, \dots, c_{r-1})$ and $B' = (d_0, d_1, \dots, d_{r-1})$ be two blocks in the disjoint block decomposition of V_e . Also let f_i be the r -dimensional vector having a 1 in coordinate i and 0's everywhere else, and let $e_{it} = f_i - f_t$. We will say that B and B' are related if they satisfy

$$B = B' + e_{it} \quad (*)$$

for some $0 \leq i, t \leq r-1$, where $+$ denotes vector addition. We then define the block graph $BG(e)$ as the graph whose vertices are the blocks in the disjoint block decomposition of V_e , and whose edges are unordered block pairs BB' for which B and B' are related (in the above sense). When S is a set of points in a graph G , we denote by $\langle S \rangle$ the subgraph of G induced by S .

Lemma 2: Let $F \subset BG(e)$. Then

$\langle p(F) \cup Nb(p(F)) \rangle$ is a connected subgraph of $Q(n) \iff \langle F \rangle$ is a connected subgraph of $BG(e)$.

Proof: \Leftarrow : Suppose BB' is an edge of $\langle F \rangle$ with B and B' satisfying $(*)$ for some i and t . Then the block $B'' = B' + f_i$ (note that $B'' \subset W(k+1, n)$) satisfies

$$p(B'') \subset [Nb(p(B')) \cap Nb(p(B))].$$

It follows by lemma 1 that $p(B \cup B') \cup Nb(p(B \cup B'))$ induces a connected subgraph of $\langle F \rangle$. The connectedness of F and the preceding observation then imply by a simple induction that $\langle p(F) \cup Nb(p(F)) \rangle$ is connected, as required.

\Rightarrow : Let B, B' be distinct vertices of F , and let $v \in p(B)$, $v' \in p(B')$. Let $v = v_0, v_1, v_2, \dots, v_t = v'$ be a path joining v and v' in $\langle p(F) \cup Nb(p(F)) \rangle$. Observe that $v_i \in p(F)$ for i even and $v_i \in W(k+1, n)$ for i odd, and that t is even. For i even, let B_i be the block of V_e to which v_i belongs. Then each pair $B_i B_{i+2}$ are related (if not identical). Hence B and B' are joined in $\langle F \rangle$ by the walk $B = B_0, B_2, \dots, B_t = B'$. ■

In view of lemma 2, we can determine the connected components structure of R_e by instead determining the connected components structure of $BG(e)$. To facilitate this, we will use the following terminology. We will view a block $B = (c_0, c_1, \dots, c_{r-1})$ as an ordered set of r piles of

"chips". We will refer to these piles as "pile 0", "pile 1", ..., "pile r-1". Alternatively, view a block as a configuration of $\sum_{i=0}^{r-1} c_i = k$ indistinguishable objects into r distinguishable cells. We make notational use of the cyclic nature of the subscripts in $(c_0, c_1, \dots, c_{r-1})$ as follows. When a coordinate entry in an r-tuple is expressed involving a subscript i, that entire entry is understood to denote c_i , the number of chips in pile i. For example, $(0, 0, \dots, 0, x_i+1, 0, \dots, 0)$ denotes the r-tuple with only one nonzero entry, namely x_i+1 , as its coordinate c_i . Further, when c_i is expressed using a term involving a subscript i as just discussed, we may display the entries c_{i-1}, c_{i-2} , etc. as appearing to the left of c_i , even when i-1 or i-2, etc. are "greater than i" when taken as nonnegative residues modulo r. For example, $(0, 1, 0, 0, x_2, 5)$ denotes the same block more usually denoted by $(0, 0, x_2, 5, 0, 1)$, namely, the block having x_2 chips in pile 2, 5 chips in pile 3, 1 chip in pile 5, and 0 chips in the other piles.

A block $(c_0, c_1, \dots, c_{r-1})$ (or equivalently, the corresponding piles of chips) will be called legal if it belongs to V_e . When e is fixed by context, the sum of B is $\text{sum}(B) = e + \sum_{i=0}^{r-1} c_i$, thought

of as an integer modulo r. The relation (*) shows that traversing an edge $B'B$ of $BG(e)$ is equivalent, in the language of chips, to transferring a single chip in configuration B' from pile t to pile i, obtaining as a result the configuration B. We call the transfer of a single chip from one pile of a block to another (possibly empty) pile a move. Depending upon which two piles are involved, a move may or may not result in a block which lies in V_e . A move will be called legal iff both its starting and ending blocks (B' and B) are legal (i.e., they lie in V_e). By definition, then, a move is legal iff B and B' are legal, i.e. $\text{sum}(B)$ (resp. $\text{sum}(B')$) $\equiv 0, 1, 2, \dots$ or $r-s(B)$ (resp. $s(B')$) mod r. We will refer to the set $\{B \in BG(e) : s(B)=i\}$ as level i. For a given i, $2 \leq i \leq r$, we let $\omega(i) = r-i+1$. Thus for fixed e, $\omega(i)$ is the "first" disallowed sum value for any block of V_e lying in level i, the others being $\omega(i)+1, \omega(i)+2, \dots$ through r-1. When $i = 1$, any sum value is allowed for a block at level i, i.e. all blocks in level 1 are legal.

Lemma 3: Let $r \geq 6$. Then for any e, the vertices of $BG(e)$ lying in level 1 are in the same connected component of $BG(e)$.

Proof: Let B be a vertex in $BG(e)$ all of whose chips are in pile i. It suffices to find a sequence of legal moves ending with the vertex all of whose chips are in pile i-1. (Notice that all blocks lying in level 1 are legal.) As notation in this proof, we augment the r-tuple $(c_0, c_1, \dots, c_{r-1})$ denoting a block B by $(c_0, c_1, \dots, c_{r-1}; z)$, where $z = \text{sum}(B)$. Also a "long" sequence of consecutive 0's among the coordinates will be written in exponent form; e.g. $(0^t, c_t, \dots, c_{t+m}, 0^{r-t-m-1}; z)$ will denote a block whose first t and last r-t-m-1 piles are empty.

We move chips from pile i to pile $i-1$ one at a time (reducing the sum by $1 \pmod r$ each time) until we reach a block with sum $\omega(2)+1$ or until pile i is exhausted, whichever comes first. In the latter case we would be done, so suppose we arrive at a block B' with sum $\omega(2)+1$ lying in level 2, and represent B' by $(0^{i-1}, x_{i-1}, x_i, 0^{r-i+1}; \omega(2)+1)$. (Possibly $x_{i-1} = 0$ if the original block B already had sum $\omega(2)+1$, and then $B'=B$.)

We now make the move $B' \rightarrow (0^{i-3}, 1, 0, x_{i-1}, x_i - 1, 0^{r-i-1}; \omega(2) - 2)$, which is legal since $\omega(3) = \omega(2) - 1$. Now again move one chip at a time from pile i to pile $i-1$, until either pile i is exhausted or the sum reaches $\omega(2)+1$, whichever comes first.

Assume first that pile i is exhausted. At that stage we are at a block $C = (0^{i-3}, 1, 0, k-1, 0^{r-i}; z)$, where the $k-1$ entry indicates the number of chips in pile $i-1$, and where z is neither $\omega(2)$ nor $\omega(2)-1 \pmod r$. Thus $C \subset V_e$ since $s(C) = 2$ while $\text{sum}(C)$ is not the (single) disallowed value $\omega(2)$. The move $C \rightarrow (0^{i-1}, k, 0^{r-i})$ is then legal since the resulting block lies in V_e , and we are done.

So suppose that pile i was not exhausted, and we have reached a block $D = (0^{i-3}, 1, 0, y_{i-1}, y_i, 0^{r-i-1}; \omega(2) + 1)$. Now make the moves $D \rightarrow (0^{i-1}, y_{i-1}+1, y_i, 0^{r-i-1}; \omega(2) + 3) \rightarrow (0^{i-1}, y_{i-1}+2, y_i - 1, 0^{r-i-1}; \omega(2) + 2) \rightarrow (0^{i-1}, y_{i-1}+3, y_i - 2, 0^{r-i-1}; \omega(2) + 1) = E$, where if either $y_i - 1$ or $y_i - 2$ is 0 then we stop and are done with the proof. In the final block of the sequence we must have $y_i - 2 < x_i$. Substituting E for B' and repeating the process starting from B' , we are done by induction. ■

Lemma 4: Every block $B \in \text{BG}(e)$ satisfying $s(B) \leq r-1$ is joined by a path in $\text{BG}(e)$ to some block B' in $\text{BG}(e)$ having exactly two nonempty piles, these piles being consecutive.

Proof: Call a block in $\text{BG}(e)$ having exactly two nonempty piles, these being consecutive, a 2-block. Observe that for any $B \in \text{BG}(e)$ we can assume $s(B) \geq 2$, since when $s(B) = 1$ we can just move a chip from the nonempty pile to either of the two piles next to it, obtaining a 2-block. (At least one of these moves must result in a legal block.)

The basic plan is to make moves which "compress" the piles of chips until we reach a 2-block. That is, for any block $B = (c_0, c_1, c_2, \dots, c_{r-1})$ let $m(B) = \min\{i: c_i > 0\}$ and $M(B) = \max\{i: c_i > 0\}$. Then our moves will be such that m is nondecreasing, M is nonincreasing, with at least one out of every two consecutive moves managing to include either the transfer of a chip from pile M to pile $M-1$, or from pile m to pile $m+1$, with no chip ever transferred to either of piles m or M .

Suppose first that $\lceil \frac{r}{2} \rceil + 1 < s(B) \leq r - 1$. Let $X(B) = \min\{i: c_i > 0, c_{i+1} > 0, i \neq r-1\}$. Note that $X(B)$ exists because of $\lceil \frac{r}{2} \rceil + 1 < s(B)$ and the pigeon hole principle. Consider the following procedure.

Procedure Compress

Input: A block $B \in BG(e)$ satisfying $\lceil \frac{r}{2} \rceil + 1 < s(B) \leq r - 1$.

Output: A block $B' \in BG(e)$ satisfying $s(B') \leq \lceil \frac{r}{2} \rceil + 1$ joined to B by a path in $BG(e)$.

begin

$B' \leftarrow B$

(1) If $s(B') \leq \lceil \frac{r}{2} \rceil + 1$, then return B' .

Otherwise:

(2) If $\text{sum}(B') \neq 0$, then move one chip from pile $M(B')$ to pile $M(B') - 1$.

(3) Otherwise (so $\text{sum}(B') = 0$)

(a) If $\omega(s(B)) = 2$, then move one chip from pile $X(B)$ to pile $X(B)+1$.

(b) Otherwise (so $\omega(s(B)) \geq 3$)

Move one chip from pile $m(B)$ to pile $m(B)+1$.

Let C be the block resulting from executing whichever of (2) or (3) is applicable.

$B' \leftarrow C$

Repeat step (1).

end

We claim that this procedure is valid; that is, the block B' returned has the properties specified in the description of the output.

Observe first that if all the moves made in this procedure were legal, then we would be done. For then, each execution of step (2) or (3)b moves a chip from the extreme right to the left or from the extreme left to the right. On applying (3)a, the resulting block C satisfies $\text{sum}(C) > 0$, so when we repeat step (1) we will either be done (with the desired block) or (2) will be executed. Hence a single cycle of the procedure results in either a block of the desired kind, or a nondecrease in m along with a nonincrease in M and a move of one chip from pile m to the right, or from pile M to the left. Since the total number of chips is finite, we must reach a block B' satisfying $s(B') \leq \lceil \frac{r}{2} \rceil + 1$. This B' is in the same component as the original B since all moves were legal.

Next we show that starting with a legal block B , any single move in this procedure (as detailed in steps (2) or (3)) results in a legal block C . It would then follow by induction that all moves executed are legal, and the validity of Procedure Compress would follow from the preceding paragraph.

Suppose first that (2) is executed. Then $s(C) \leq s(B) + 1$ while $\text{sum}(C) = \text{sum}(B) - 1$. Hence it follows from $0 < \text{sum}(B) \leq \omega(s(B)) - 1$ and $\omega(t+1) = \omega(t) - 1$ for any t , that $0 \leq \text{sum}(C) \leq \omega(s(C)) - 1$. Thus C is legal.

Now suppose (3) is executed. If $\omega(s(B)) = 2$, then after executing (3)a we have $\omega(s(C)) \geq \omega(s(B)) = 2$ (since a chip transfer was made between nonempty contiguous piles), and $\text{sum}(C) = \text{sum}(B) + 1 = 1$. Hence C is legal. Otherwise, after executing (3)b we get $\text{sum}(C) = 1$ and $\omega(s(C)) \geq \omega(s(B)) - 1 \geq 2$. Hence again C is legal. The validity of Procedure Compress is thus proved.

We have shown so far that there is a path in $BG(e)$ from any block B satisfying $\lceil \frac{r}{2} \rceil + 1 < s(B) \leq r - 1$ to some block B' satisfying $s(B') \leq \lceil \frac{r}{2} \rceil + 1$. The proof of the lemma is now completed by observing that for any such B' there is a path in $BG(e)$ joining it to some 2-block of $BG(e)$. The proof of this observation consists in just applying a trivially modified version of Compress, which we call Compress', to B' . Compress' is the same as Compress, except that the phrase "If $s(B') \leq \lceil \frac{r}{2} \rceil + 1$ " in step (1) is replaced by "If B' is a 2-block". The validity of B' is proved in the same way, and in fact consideration of step (3)a can be omitted since the condition $s(B) \leq \lceil \frac{r}{2} \rceil + 1$ implies that $\omega(s(B)) \geq 3$. ■

Theorem 2: For any $r \geq 6$, and any $0 \leq e \leq r-1$, the graph $BG(e)$ has the following connected components:

- (1) A "large" component $\underline{L}(e) = \{B \in BG(e) : s(B) \leq r-1\} \cup \{B \in BG(e) : s(B) = r, \text{ and } B \text{ has at least one pile containing exactly one chip}\}$.
- (2) A singleton component $\{B\}$ for each block $B \in BG(e)$ satisfying $s(B) = r$ and having at least two chips in each pile. We denote the collection of such components $\underline{I}(e)$.

Proof: Let B be an arbitrary block of $BG(e)$ satisfying $s(B) \leq r-1$. By lemma 4 we know that B is in the same connected component of $BG(e)$ as some 2-block B' , the latter having all its chips in, say, piles i and $i+1$. Now proceed as in the proof of lemma 3 to show that there is a path in $BG(e)$ from B' to a block B'' in level 1. (Move one chip at a time from pile $i+1$ to pile i until the sum reaches $\omega(2) + 1$, etc.) It follows by lemma 3 that all vertices B in $BG(e)$ satisfying $s(B) \leq r-1$ are contained in a single connected component.

Now take a vertex B of $BG(e)$ satisfying $s(B) = r$. Then $\text{sum}(B)=0$, since 0 is the only allowable sum value for blocks at level r .

Suppose first that every pile of B has at least two chips. Then any single chip move from B results in a block C for which $s(C)$ is still r , but $\text{sum}(C) \neq \text{sum}(B) = 0$. Hence C is not legal. Thus $\{B\}$ must be a singleton connected component of $BG(e)$, proving statement (2).

On the other hand suppose B has some pile i containing exactly one chip. Then moving this chip to (the nonempty) pile $i+1$ gives a block C satisfying $s(C) = r-1$ and $\text{sum}(C) = \text{sum}(B)+1 = 1$. Hence C is legal, and B belongs to the same "large" connected component of $BG(e)$ as C , proving statement (1). This completes the proof of the theorem. ■

Corollary 1.1: Suppose $r \geq 6$ and $k \leq 2r - 1$. Then the graph $BG(e)$ is connected for every e , $0 \leq e \leq r-1$.

Proof: Let B be a block of $BG(e)$ at level r . Then $k \leq 2r - 1$ implies that some pile of B has exactly one chip. Thus the "large" component of $BG(e)$ is all of $BG(e)$. ■

We now proceed to our upper bound for $L(n, W(k+1, n))$.

Theorem 3:

$$\begin{aligned} L(n, W(k+1, n)) &\leq \binom{n}{k+1} + \binom{n}{k} \frac{1}{\lfloor k/\log(k) \rfloor} \frac{\log(k)}{\log(k) - 1} + O\left(\frac{1}{\lfloor k/\log(k) \rfloor} \binom{k - \lfloor k/\log(k) \rfloor - 1}{\lfloor k/\log(k) \rfloor - 1}\right) \\ &= \binom{n}{k+1} + (1+o(1)) \left(\frac{\log(k)}{k}\right) \binom{n}{k}. \end{aligned}$$

Proof: Recall the partition of $BG(e)$ into the "large component", which we called $L(e)$, and the set of isolated blocks, which we called $I(e)$. Our first step is to show that there is a subset C_e of $Q(n)$ such that $|C_e| = |V_e| + O(|I(e)|)$ and $C_e \cup W(k+1, n)$ induces a connected subgraph of $Q(n)$.

Let $B \in I(e)$. Suppose we remove two chips of B from pile 0, and place one of them at some pile $i \neq 0$ and the other at some pile $j \neq 0$, where $i+j = r$. Then the resulting block B_1 belongs to $L(e) \cup I(e)$, and there must be vertices $x_0 \in p(B)$ and $x_1 \in p(B_1)$ joined by a path of length 4 whose vertices have weight $k, k-1$, or $k-2$. Repeat this procedure on B_1 in place of B (where the i and j in this step need not be the same as they were in the previous step - they need only sum to r), obtaining a block $B_2 \notin \{B, B_1\}$ and a vertex $x_2 \in p(B_2)$ at distance 4 from x_1 . Continue this process until we have found a sequence of blocks $B = B_0, B_1, B_2, \dots, B_m, B_{m+1}$ and a sequence of vertices $x_0, x_1, x_2, \dots, x_m, x_{m+1}$, where $x_i \in p(B_i)$, such that

- (i) $B_i \in I(e)$ for $0 \leq i \leq m$, $B_{m+1} \in L(e)$,
- (ii) For each i , $0 \leq i \leq m$, there is a path of length 4 joining x_i and x_{i+1} where vertices of the path have weights $k, k-1$, or $k-2$.

We note that there must exist an m for which $B_{m+1} \in L(e)$ since the successive removal of chips from pile 0 must eventually lead to a block (which we called B_{m+1}) having either one chip or no chips in pile 0. This block must belong to $L(e)$ by theorem 1. We also get a path in $Q(n)$ from x_0 to the set $p(L(e))$, every fourth point of which belongs to $p(I(e))$ (starting with x_0). Call this path $P(B)$.

We now construct the set C_e as follows. Choose some $B \in I(e)$, and let $D^{(1)} = P(B)$. Now suppose inductively that we have constructed a set $D^{(i)} \subset Q(n)$. If $D^{(i)}$ contains a point of $p(B)$ for

every $B \in I(e)$, then let $D_e = D^{(i)}$. Otherwise, let $B' \in I(e)$ be such that $p(B') \cap D^{(i)} = \emptyset$. Let $Q(B')$ be a subpath of $P(B')$ starting at a point in B' and ending at the first point of $P(B')$ which lies either in $D^{(i)}$ or in $L(e)$. Now let $D^{(i+1)} = D^{(i)} \cup Q(B')$. This procedure eventually produces a set $D_e = D^{(t)}$, where t is the smallest integer such that $D^{(t)}$ contains a point of $p(B)$ for every $B \in I(e)$. Now let $C_e = D_e \cup p(V_e)$.

We now observe that the graph induced by $C_e \cup W(k+1, n)$ is connected. This follows from the following facts.

(a) The graph induced by $p(L(e)) \cup \text{Nb}(p(L(e)))$ is connected by lemma 2 and theorem 2.

(b) For each $B \in I(e)$ the subgraph induced by $p(B) \cup \text{Nb}(p(B))$ is connected by lemma 1.

(c) Each of the subgraphs mentioned in (b) is joined to $p(L(e))$ in $Q(n)$ by a path in D_e .

Note also that $|D_e| = O(|I(e)|)$, where the constant factor implicit in the big O is 4. The set C_e with the desired properties has thus been constructed.

The reader is referred to Figure 1 for an illustration of how $C_e \cup W(k+1, n)$ is connected. In that figure, rectangles denote blocks. While we made no use of blocks within $W(k+1, n)$, the set $W(k+1, n)$ does partition into blocks, and for each block B in $W(k, n)$, $\text{Nb}(p(B))$ is a union of blocks. For each block B in $W(k, n)$, the edges between $p(B)$ and the blocks comprising $\text{Nb}(p(B))$ are indicated in Figure 1 by thickened edges. Recall by lemma 1 that each $p(B) \cup \text{Nb}(p(B))$ induces a connected subgraph. By theorem 2 and lemma 2, the subgraph induced by $p(L(e)) \cup \text{Nb}(p(L(e)))$ is connected. In Figure 1, dots denote points of D_e . Observe that each isolated block B' in $I(e)$ is joined to the large component $L(e)$ by the path $P(B')$, these paths indicated by the thin directed edges.

We can now estimate $L(n, W(k+1, n))$. First we have

$$\sum_{e=0}^{r-1} |C_e| = \sum_{e=0}^{r-1} |V_e| + O\left(\sum_{e=0}^{r-1} |I(e)|\right)$$

so it follows that for some t we have

$$|C_t| \leq \frac{1}{r} \sum_{e=0}^{r-1} |V_e| + O\left(\frac{1}{r} \sum_{e=0}^{r-1} |I(e)|\right).$$

Now the expression $\frac{1}{r} \sum_{e=0}^{r-1} |V_e|$ is estimated in [FR], and with the choice $r = \lfloor k/\log(k) \rfloor$, is shown to

be bounded above by $\binom{n}{k} \frac{1}{\lfloor k/\log(k) \rfloor \log(k) - 1}$. Consider now the second term

$O\left(\frac{1}{r} \sum_{e=0}^{r-1} |I(e)|\right)$. We observe that the union $\bigcup_{0 \leq e \leq r-1} I(e)$ has size equal to the number of ways of

distributing k indistinguishable chips among r distinguishable piles, where there are at least two chips in each pile. This number is $\binom{k-r-1}{r-1}$. Since this union is disjoint, it follows that

$$O\left(\frac{1}{r} \sum_{e=0}^{r-1} |I(e)|\right) \leq \frac{1}{r} \binom{k-r-1}{r-1}.$$

Plugging in the choice $r = \lfloor k/\log(k) \rfloor$, the first inequality of the theorem follows.

Consider now the equality in the theorem. Using standard bounds on binomial coefficients we have $\binom{k - \lfloor k/\log(k) \rfloor - 1}{\lfloor k/\log(k) \rfloor - 1} \leq (e \log(k))^{k/\log(k)}$ and $\binom{n}{k} \geq \left(\frac{n}{k}\right)^k$. Now taking logarithms and simplifying, it follows easily that the ratio of the right side of the first inequality to the right side of the second approaches 0 as $k \rightarrow \infty$.

This completes the proof of the theorem. ■

3. Optimality

We now show that the upper bound for $L(n, W(k+1, n))$ given in Theorem 3 is asymptotically optimal to within a factor of $\log(k)$ in a certain strong sense. Let

$M(k) = \min\{ |H \setminus W(k+1, n)| : H \text{ is a connected subgraph of } Q(n) \text{ spanning } W(k+1, n) \}$, so that $M(k) = L(n, W(k+1, n)) - \binom{n}{k+1}$. We will show that $(1+o(1)) \left(\frac{\log(k)}{k}\right) \binom{n}{k}$ is optimal for $M(k)$

to within a factor of $\log(k)$ as $k \rightarrow \infty$. It is interesting to observe that the evidently nearly optimal construction used in Theorem 3 makes use only of points having weights $k+1, k, k-1,$ and $k-2$, a fairly restricted list of weights.

Let H be a connected subgraph of $Q(n)$ spanning $W(k+1, n)$. Each vertex of $W(k+1, n)$ must have at least one neighbor in H . For $x \in H$, let $N(x)$ be the number of neighbors of x which lie in $W(k+1, n)$. Note that $N(x) \leq n - k$ (recalling the assumption $k+1 \leq \lfloor \frac{n}{2} \rfloor$). Then

$$\binom{n}{k+1} \leq \sum_{x \in H \setminus W(k+1, n)} N(x) \leq (n - k) |H \setminus W(k+1, n)|,$$

and it follows that

$$|H \setminus W(k+1, n)| \geq \frac{1}{n - k} \binom{n}{k+1} = \frac{1}{k+1} \binom{n}{k}.$$

Therefore we have $M(k) \geq \frac{1}{k+1} \binom{n}{k}$, showing that the result of Theorem 3 is within a factor of $\log(k)$ from optimality for $M(k)$ as $k \rightarrow \infty$.

References

[FR] P. Frankl and V. Rödl, Lower bounds for Turan's Problem, Graphs and Combinatorics 1 (1985) 213-216.

- [H] S.L. Hakimi, Steiner's problem in graphs and its implications, Networks 1 (1971) 113-133.
- [MP] Z. Miller and M. Perkel, The Steiner problem in the hypercube, Networks 22 (1992) 1-19.
- [Tu] P. Turan, Eine Extremalaufgabe aus der Graphentheorie, Mat. Fiz. Lapok 48 (1941) 436-452.
- [Wi] P. Winter, Steiner Problem in networks: a survey, Networks 17 (1987) 129-167.