Bounding the Size of k-Tuple Covers

Wolfgang Bein
School of Computer Science
Center for the Advanced Study of Algorithms
University of Nevada, Las Vegas
bein@egr.unlv.edu

Shahram Latifi
Department of Electrical & Computer Engineering
University of Nevada, Las Vegas
latifi@egr.unlv.edu

Linda Morales
Department of Computer Science
University of Texas at Dallas
lmorales@utdallas.edu

I. Hal Sudborough
Department of Computer Science
University of Texas at Dallas
hal@utdallas.edu

Abstract.

Suppose there are n applications and n processors. A pair cover is a set S of one-to-one mappings (assignments) of the applications to the processors such that, for every pair (A_i,A_j) of applications and every pair (p,q) of processors, there is an assignment f in S that maps (A_i,A_j) to (p,q). More generally, we consider, for all k≥1, minimum size k-tuple covers. We improve bounds given earlier in by Latifi, where the application for k-tuple covers was fault tolerance of the multidimensional star network. Let F(n,k) denote the minimum cardinality k-tuple cover for n applications and processors. We give bounds for F(n,k) that are within a small multiplicative factor of optimum.

1. Introduction

Suppose, for some n>1, there are n applications A_1, A_2, … , A_n and n processors. An *assignment* f is a one-to-one mapping of {A_1,A_2,…,A_n} to the n processors. We are interested in sets S of assignments such that, for every pair (A_i,A_j) of applications and every pair (p,q) of processors , there is an assignment f in S that maps (A_i,A_j) to (p,q). A set S of assignments that satisfies this property is called a *pair cover*. We wish to compute pair covers of minimum cardinality. More generally, we wish to compute, for all k≥1, minimum cardinality *k-tuple covers*.

Equivalently, we can define k-tuple covers as a set of permutations. Let $\pi = \pi_1, \pi_2. \dots , \pi_n$ be a permutation of length n over Z_n ={0,1,2,...,n−1}, describing an assignment of application π_i to processor i, for all i (0≤i≤n−1), and let $\lambda=(\lambda(1), \lambda(2). \dots , \lambda(k))$ be a k-tuple of distinct processors (positions) chosen from Z_n. Then, $\pi|\lambda = \pi_{\lambda(1)}, \pi_{\lambda(2)}. \dots , \pi_{\lambda(k)}$ is the sub-list consisting of the applications assigned to processors (in positions, respectively) λ(1), λ(2), … , λ(k). We say that a set S of permutations over Z_n ={0,1,2,...,n−1} is a *k-tuple cover of Z_n* if for any k-tuple z(1), z(2). … , z(k) of distinct elements from Z_n and any k-tuple λ of distinct processors (positions) chosen from Z_n there is a permutation π in S such that $\pi|\lambda$ = z(1), z(2), … , z(k).

For example, the set of all six permutations on Z_3, namely S = {012,021,120,210,201,102}, is a pair cover of minimum cardinality. The set S' = {0123, 0312, 0231, 1032, 1203, 1320, 2013, 2301, 2130, 3021, 3102, 3210} of even permutations on Z_4 is a pair cover of minimum cardinality.

A set S is said to satisfy the *P(n,k)* property if it is a set of permutations that is a k-tuple cover of Z_n. Let F(n,k) denote the minimum cardinality of any k-tuple cover for Z_n.

There are $\binom{n}{k}$ distinct ways to choose k objects from a set of n. So, there are $\binom{n}{k}$ distinct k-tuples of processors. For each k-tuple of processors there are $k!\binom{n}{k}$ different ways applications can be mapped to them. As any k-tuple of applications mapped to a chosen k-tuple of processors can be covered by a specific permutation that has those applications in the chosen positions, it follows that at most $k!\binom{n}{k}\binom{n}{k}$ permutations are needed to form a k-tuple cover. So, F(n,k) ≤ $k!\binom{n}{k}\binom{n}{k}$. On the other hand, one permutation can cover at most $\binom{n}{k}$ distinct k-tuples of positions, so any *k-tuple cover* must have at least *lowerbound(n,k)* = $k!\binom{n}{k}$ permutations. It follows that $k!\binom{n}{k} \leq F(n,k) \leq k!\binom{n}{k}\binom{n}{k}$.

We provide improved bounds for F(n,k). In particular, we show that:

1. for all n, F(n,2) ≤ 4 · *lowerbound(n,2)*,
2. for all n, F(n,3) ≤ (8n/3) · *lowerbound(n,3)*, and
3. for all n, F(n,4) ≤ ($3n^2$/2) · *lowerbound(n,4)*.

That is, we give an upper bound for the cardinality of a pair cover that is at worst within a multiplicative factor of 4 of optimum, an upper bound for the cardinality of a triple cover that is at worst within a multiplicative factor of (8n/3) of optimum, etc.

The *P(n,k) problem* is the problem of deciding, given input n, k, and r, if there is k-tuple cover of Z_n of cardinality r. For each fixed k, the problem is in the class *NP*, it is unknown whether it is *NP*-complete, or whether it is computable deterministically in polynomial time.

The problem of computing *k-tuple covers* of small cardinality, for each k≥2, has several applications [5,6]. In particular, it is useful in the design of sampling tests of individuals or disparate populations (corresponding to processors) and medicines or drugs (corresponding to applications) [6]. In the next section we describe a specific application in the area of fault tolerance in a network for parallel processing called the *star network*.

2. Fault Tolerance in Star Networks

Star graph interconnection networks are attractive alternatives to the hypercubes. Algebraic and combinatorial properties of the star graph were initially studied in [2,3,9]. The star network has superior node degree and diameter compared to the hypercube of a comparable size. More specifically, growth in node degree and diameter is sublogarithmic to network size in the star but logarithmic in a hypercube [2,3]. Various networks have been mapped to the star network [4,12,13]. A treatment of communication aspects for this network is presented in [7]. Fault tolerance of star graph has also been addressed extensively (see [9−11,14] for example).

Fault tolerance of large networks is generally measured by how much of the network structure is preserved in the presence of a given number of node or link failures. For example, parallel algorithms running on such networks utilize the topological properties of these networks. In the presence of node/link failures, the entire network is not available. Thus a natural question is: How large of a subnetwork (defined as a smaller network but with the same topolgical properties as the original one) is still available in the network in the presence of faults. In the use of star networks as the underlying interconnection topology, estimates of network reliability and fault tolerance are important in choosing algorithms and predicting their performance under failure conditions. Although, the first component failure (node or link), makes the complete star network unavailable, it is important to know how many fault-free star networks of smaller size or substars are available in the damaged structure.

Here we consider the node-failure model, and study bounds on the number of faulty nodes, which make every substar of a certain size faulty in a given star network.

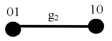

Star graph of dimension 2

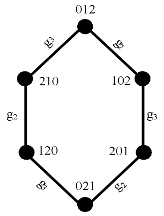

Star graph of dimension 3

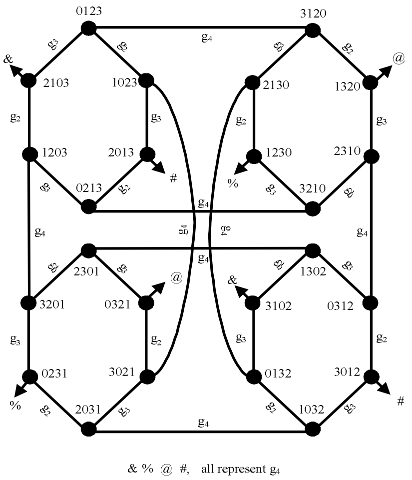

Star graph of dimension 4

Figure 1. The star graphs S_2, S_3, and S_4,

The n-dimensional star graph, S_n, is an edge- and node-symmetric graph containing $n!$ nodes and (n−1)n!/2 edges. The nodes are assigned labels, each being a distinct permutation on the set of integers Z_n = {0,..,n−1}. Two nodes are joined with a link labeled g_i or *i* (for ease of notation) if and only if the label of one can be obtained from the label of the other by swapping the first symbol and the i^{th} symbol, where

$1<i\leq n$. For instance, in a 4-star containing 24 nodes, two nodes 0123 and 3120 are neighbors and joined via an edge labeled 4. Each S_n contains n disjoint copies of S_{n-1}, but there are several ways to do the partitioning. For example, S_n can be partitioned into disjoint copies of S_{n-1} in (n−1) ways, *i.e.* by removing all i-links, for $1<i\leq n$.

For example, partitioning S_4 along the 4th link results in four copies of S_3, namely those formed by the nodes whose labels end with 0, 1, 2, or 3, respectively. If, for some i ($0\leq i\leq 3$), all nodes whose labels end with i are intact, then the copy of S_3 defined by these nodes is non-faulty. Equivalently, if for all i ($0\leq i\leq 3$), at least one of the nodes whose label ends with i is damaged, then all four of these copies of S_3 are faulty. However, in the latter case there may still be a non-faulty copy of S_3, as there are other ways to partition. For example, if the four nodes labeled with permutations 0123, 3120, 2031, and 1032 are faulty, there is still a non-faulty copy of S_3. For example, the copy of S_3 formed by the six nodes 2103, 3102, 1302, 2301, 3201, 1203 is still intact. This is one of the four disjoint copies of S_3 formed by deleting the links labeled with 3.

All copies of S_3 will be damaged if the labels for the set of faulty nodes is a set of permutations with each symbol in every position. For example, if the set of nodes in S_4 with labels 0123, 1230, 2301, and 3012 are damaged, then there is no intact copy of S_3. This is so, as any copy of S_3 in S_4 is formed by a set of permutations in which symbols in the first position and two of the other three positions are exchanged by transpositions, but the remaining position has a fixed symbol. That is, by choosing a set of permutations that covers each symbol in each position one damages every copy of S_3. For example, in the copy of S_3 formed by the six nodes 2103, 3102, 1302, 2301, 3201, 1203, symbols in position 0 are exchanged with symbols in positions 1 and 3, and the symbol in position 2 (*i.e.*, the symbol 0) is fixed. This copy of S_3 includes the node labeled 2301, which is damaged. In all, the damaged node 2301 destroys three additional copies of S_3, namely the copy of S_3 in which the symbol 2 is fixed in position 0, the copy of S_3 in which the symbol 3 is fixed in position 1, the copy of S_3 in which the symbol 1 is fixed in position 3.

More generally, a set of nodes (or permutations labeling nodes) in S_n damages every copy of S_{n-k} if it is a k-tuple cover. That is, a copy of S_{n-k} in S_n is formed by a set of permutations in which symbols in position 0 and n−k−1 of the remaining n−1 positions are exchanged by transpositions, but k positions have fixed symbols. A set of permutations that forms a k-tuple cover includes at least one permutation from each possible copy of S_{n-k} in S_n, so if all of the nodes labeled by the permutations in a k-tuple cover are damaged, then every copy of S_{n-k} in S_n is damaged.

A careful reading of the preceding shows that a restriction of a k-tuple cover is required to damage every copy of S_{n-k} in S_n. That is, the k positions with fixed symbols are chosen from positions $\{1,2,\ldots,n-1\}$, which does not include position 0. So, let a set S of permutations over $Z_n=\{0,1,2,\ldots,n-1\}$ be a *restricted k-tuple cover of Z_n* if for any k-tuple z(1), z(2). … , z(k) of distinct elements of Z_n and any k-tuple λ of distinct processors (positions) chosen from $\{1,2,\ldots,n-1\}$ there is a permutation π in S such that $\pi|\lambda$ = z(1), z(2), … , z(k). A restricted k-tuple cover of Z_n is required to damage every copy of S_{n-k} in S_n.

A set S is said to satisfy them *R(n,k)* property if it is a set of permutations that is a restricted k-tuple cover of Z_n. Let *f(n,k)* denote the minimum cardinality of any restricted k-tuple cover for Z_n. Clearly, for all n,k, $f(n,k)\leq F(n,k)$ and $k!\binom{n}{k}\leq f(n,k)\leq k!\binom{n}{k}\binom{n-1}{k}$.

It is well known that the n! permutations on the set {1,..,n} can be expressed as a sequence of transpositions of the form (i j). Accordingly, the permutation labels in the star graph can be expressed as a sequence of an odd number (or an even number) of transpositions and traversing an edge changes the parity. It follows, of course, that the star graph is bipartite.

3. New Upper Bounds

We give here improved upper bounds for F(n,k) and f(n,k) and algorithms to compute efficient k-tuple covers and restricted k-tuple covers.

Note that for some instances of n and k, we can construct a k-tuple cover S with lowerbound(n,k) permutations, and therefore in these cases our bound for F(n,k) is optimum. For example, F(n,1)=n as the set of permutations consisting of the permutation $\pi=012\ldots n-1$ and all cyclic shifts of π covers all elements of Z_n in every position. For example, the set of faulty nodes {01234, 12340, 23401, 34012, 40123} will damage all S_4's in S_5. Furthermore, F(n,n−1) = n!, as one needs the set of all permutations over Z_n in order to cover all (n−1)-tuples of elements from Z_n in all positions.

Also, F(n,n−2) = n!/2, as the set A of all even permutations is a (n−2)-tuple cover. That is, for any (n−2)-tuple (z(1), z(2). … , z(n−2)) from Z_n and any (n−2)-tuple λ of distinct positions chosen from Z_n, let π be a permutation such that $\pi|\lambda$ = (z(1), z(2). … , z(n−2)). If π is an even permutation, *i.e.* π is in A, we are done. So, suppose π is an odd permutation. Let i and j be the two positions not included in the (n−2)-tuple λ. Transform π into π' by the transposition that exchanges the symbols in positions i and j. Then, π' is an even permutation and $\pi'|\lambda = (z(1), z(2). \ldots , z(n-2))$, as the transposition leaves the symbols in positions λ fixed. So, for all n, $F(n,1)=n$, $F(n,n-1)=n!$, and $F(n,n-2)=n!/2$ are optimum values.

We now consider $F(p,2)$, where p is prime. For each i, $(0 \le i \le p-1)$, consider the permutation $\pi(i) = 0, i, i+i \equiv 2i \pmod{p}, 3i \pmod{p}, \ldots , (p-1)i \pmod{p}$. For all i, $\pi(i)$ is a permutation, as i is a generator of the cyclic group Z_p, so the values $0, i, 2i, \ldots , (p-1)i$ are all distinct. In particular, consider the permutation $\pi(2) = 0,2,4,6,1,3,5$ over Z_7. The distance from 6 to 3 in $\pi(2)$ is 2, as the symbol 3 is the second symbol after 6 in the permutation. Similarly, the distance from 3 to 6 in this permutation is 5, as the symbol 6 is the 5th symbol after 3, where the count is done with wraparound. In general, the distance from j to k $(0 \le j \ne k \le p-1)$ in $\pi(i)$ is a, where $(k-j) \equiv a \cdot i \pmod{p}$. For example, when $p=7$, the distance from 6 to 3 in $\pi(2) = 0,2,4,6,1,3,5$ is 2, and $(3-6) \equiv 2 \cdot 2 \pmod{7}$ and the distance from 3 to 6 in $\pi(2)$ is 5, and $(6-3) \equiv 2 \cdot 5 \pmod{7}$. Observe that, for all j,k, with $j \ne k$, and m,n, with $m \ne n$, the distance from j to k in $\pi(m)$ is not the same as the distance from j to k in $\pi(n)$, as $(j-k)m - (j-k)n \equiv (j-k)(m-n)$ is not 0 (mod p). This follows from the fact that $(j-k) \ne 0 \pmod{p}$ is a generator of the cyclic group Z_p and hence $(j-k) \cdot a \equiv 0 \pmod{p}$ implies $a \equiv 0 \pmod{p}$.

Lemma 1. For every prime p, there is a set A of $p(p-1)$ permutations over Z_p that satisfies the *P(p,2)* property.

Proof: As there are $p-1$ permutations $\pi(1), \pi(2), \ldots , \pi(p-1)$ and $p-1$ distinct distances in a permutation of length p, it follows that, for all m,j,k, with $j \ne k$, there is a permutation in $\{ \pi(1), \pi(2), \ldots , \pi(p-1)\}$ such that the distance from j to k is m. It follows that the set A of all distinct cyclic shifts of the $p-1$ permutations $\pi(1), \pi(2), \ldots , \pi(p-1)$ is a set of $p(p-1)$ permutations such that, for all j,k $(j \ne k)$, and all pairs of distinct positions $\lambda=(a,b)$, there is a permutation π in A such that $\pi|\lambda = j, k$. □

Theorem 1. For every prime p, $F(p,2) = p(p-1)$.

Theorem 1 follows immediately from Lemma 1. Theorem 1 appeared with a sketched proof in [10]. We note that [10] also claimed that, when n is not prime, that $F(n,2) \le n \cdot (p-1)$, where p is the smallest prime greater than n. However, the proof given there is not correct, and we have not been able to prove it by other means. In particular, if true, the claim in [10] would imply that $f(6,2) \le 36$. We are unable as yet to construct a pair cover for Z_6 with 36 permutations.

We can prove the following lemma, which shows, for example, that $F(6,2) \le 42$.

Lemma 2. For $n,k \ge 1$, let A be a set of permutations over Z_{n+1} that satisfies the *P(n+1,k)* property. One can form a set B over Z_n that satisfies the *P(n,k)* property such that $|B|=|A|$.

Proof: By definition of the *P(n+1,k)* property, for any k-tuple of symbols, say $\alpha_1, \alpha_2, \ldots, \alpha_k$, chosen from Z_{n+1} and k-tuple λ of positions $i_1, i_2, \ldots , i_k$ chosen from $\{1,2,\ldots,n+1\}$ there is a permutation π in A such that $\pi|\lambda = \alpha_1, \alpha_2, \ldots, \alpha_k$. We can view π can as a list of the following form:

$$\pi = X_1\ \alpha_1\ X_2\ \alpha_2\ \ldots.\ X_k\ \alpha_k\ X_{k+1},$$

where X_i, for all i $(1 \le i \le k+1)$ is a possibly empty sub-list of π, and the symbols $\alpha_1, \alpha_2, \ldots , \alpha_k$ are in positions $i_1, i_2, \ldots , i_k$ of π. In particular, this is also true when the k-tuple of symbols is chosen from Z_n, instead of Z_{n+1}, and the k-tuple λ of positions is chosen from the smaller set $\{1, 2, \ldots , n\}$. Notice that, in this latter case, the sub-list X_{k+1} is not empty and the symbol ***n*** which is in Z_{n+1}, but not in Z_n, is not one of the α_i and hence must be in one of the X_j $(1 \le j \le k+1)$. Consequently, one can create a permutation π' over the symbols Z_n by replacing the symbol ***n*** in π with the last symbol of π, with the understanding that, if ***n*** is the last symbol of π, then this operation simply deletes it. This operation both removes the last symbol of π and the symbol ***n***, so π' is a permutation over Z_n. Furthermore, all symbols other than the last one and ***n*** remain fixed in the positions they were in π by this operation, *e.g.* $\alpha_1, \alpha_2, \ldots, \alpha_k$ shown above are unaffected by this operation. That is, for any k-tuple of symbols $\alpha_1, \alpha_2, \ldots, \alpha_k$ from Z_n and any k-tuple λ of positions from $\{1,2,\ldots,n\}$, if $\pi|\lambda = \alpha_1, \alpha_2, \ldots, \alpha_k$, then $\pi'|\lambda = \alpha_1, \alpha_2, \ldots, \alpha_k$. That is, the set $B = \{ \pi' \mid \pi \text{ is in } A \}$ satisfies the *P(n,k)* property and $|B| = |A|$. □

We can restate Lemma 2 in terms of $F(n,k)$:

Theorem 2. For all n,k, $F(n,k) \le F(n+1,k)$.

Example 1. Let A be the set of 42 permutations of Z_7 that satisfy the *P(7,2)* property obtained by the construction of Lemma 1. We can obtain 42 permutations of Z_6 that satisfy the *P(6,2)* property by the above construction. For example, the permutations 0123456, 0246135, 0362514 in set A, would be transformed into the permutations 012345, 024513, 034251 by the operation described.

Using Theorems 1 and 2 we can show that our upper bound for F(n,2) is never greater than 4 times the optimum.

Theorem 3. For all $n>1$, $F(n,2) \le 4\cdot \text{lowerbound}(n,2) = 4\cdot n\cdot(n-1)$

Proof. By Bertrand's postulate [1] for any n there is a prime p such that

$$n \le p \le 2n-2.$$

By Lemmas 1 and 2 we have that

$$\begin{aligned} n(n-1) \le F(n,2) \le F(p,2) &= p(p-1) \\ \le (2n-2)(2n-3) &= 4n^2-10n+6 \\ &\le 4\cdot n\cdot(n-1), \end{aligned}$$

for all $n>1$, implying the result. □

Lemma 3. Let A and B be sets of permutations over Z_n that satisfy the *P(n,k)* property and the *R(n,k−1)* property, respectively. Let $|A|=m$ and $|B|=r$. One can construct a set H of $m+nr$ permutations over Z_{n+1} that satisfies the *R(n+1,k)* property

Proof: From A form the set D of permutations over Z_{n+1} defined by $D = \{ \boldsymbol{n}, \pi_0, \pi_1, \dots, \pi_{n-1} \mid \pi_0, \pi_1, \dots, \pi_{n-1} \text{ is in } A \}$. That is, for each of the m permutations π of A form a permutation π' of D by putting a new symbol ***n*** in position 0 and putting the elements in the list π in order in positions 1, 2, …, n−1. As adding the element ***n*** to position 0 shifts the elements of each permutation of A one place to the right, it follows that D is a restricted k-tuple cover for the set of symbols in Z_n, but not all symbols in Z_{n+1}. Furthermore, $|D|=m$.

In order to transform D into a set H of permutations that satisfies the *R(n+1,k)* property, we add to the set D a set E of permutations that allow for any k-tuple of positions chosen from {0, … , n−1} the choice of any k-tuple of symbols formed by the new symbol *n* and any choice of k−1 symbols from Z_n. Notice that no permutation in D allows such a k-tuple, as the new symbol *n* occurs only in position 0 in every permutation of D. Form the set E, where

$$\begin{aligned} E = \{\ & \pi_0, \boldsymbol{n}, \pi_1, \dots, \pi_{n-1} \\ & \pi_0, \pi_1, \boldsymbol{n}, \dots, \pi_{n-1}, \dots, \\ & \pi_0, \pi_1, \dots, \boldsymbol{n}, \pi_{n-1}, \\ & \pi_0, \pi_1, \dots, \pi_{n-1}, \boldsymbol{n} \\ & \quad \mid \pi_0, \pi_1, \dots, \pi_{n-1} \text{ is in } B\ \}. \end{aligned}$$

That is, E is formed from each permutation π of B by inserting the new symbol ***n*** immediately after the i^{th} symbol, for all i ($0 \le i \le n-1$). As $|B|=r$, $|E|=nr$. Let $H = D \cup E$. It follows that H satisfies the *R(n+1,k)* property and $|H| = m+nr$. □

Theorem 4. For $n,k \ge 1$ it holds that $f(n+1,k) \le F(n,k) + n\, f(n,k-1)$.

Theorem 4 follows directly from Lemma 3.

Example 2. As $F(5,2) = f(5,2) = 20$ and $f(5,3)=60$ it follows from Lemma 4 that $F(6,3) \le 60+5\cdot 20=160$. From Lemma 1 we have $F(7,2)=42$ and recall $F(7,1)=7$. Therefore, by Theorem 4, $f(8,2) \le 42+7\cdot 7 = 91$.

For a permutation $\pi = \pi_0, \pi_1, \dots, \pi_n$ over Z_{n+1}, and any j ($0 \le j \le n$), let *swap(π, n, j)* be the permutation π', where π' is obtained by exchanging the elements j and n wherever they are in π. It should be noted that, if j=n, then swap(π, n, j) = π.

Lemma 4. Let A and B be sets of permutations over Z_{n+1} and Z_n, respectively, that satisfy the *R(n+1,k)* property and the *P(n,k−1)* property, respectively. Let $|A|=m$ and $|B|=r$. One can construct from A and B a set C of $m+(n+1)r$ permutations that satisfies the *P(n+1,k)* property.

Proof: By definition, the set A satisfies the *R(n+1,k)* property, but not necessarily the *P(n+1,k)* property, (where position 0 is also included). A set S of permutations over Z_{n+1} is said to satisfy the *PZ(n+1,k)* property if, for any k-tuple $\lambda=(\lambda(1), \lambda(2), \dots, \lambda(k))$ of distinct positions chosen from {0,1,…,n}, *such that λ(1) = 0*, and any k-tuple of symbols z(1), z(2), … , z(k) chosen from Z_{n+1}, there is some permutation π in S such that $\pi|\lambda = z(1), z(2), \dots, z(k)$. It follows that, if D satisfies the *PZ(n+1,k)* property, then $C = A \cup D$ satisfies the *P(n+1,k)* property.

Now let the set $D = \bigcup_{0\le t\le n}\{\text{swap}(\pi, n, t) \mid \pi = \mathbf{n}, \pi_0, \dots, \pi_{n-1} \text{ such that } \pi_0, \dots, \pi_{n-1} \text{ is in B}\}$. Clearly, D has $(n+1)r$ elements, as there are n+1 choices for t and $|B|=r$. We prove that D satisfies the *PZ(n+1,k)* property. Let $\lambda=(\lambda(1), \lambda(2), \dots, \lambda(k))$ be any k-tuple of positions, *where λ(1) = 0*, and let z(1), z(2), … , z(k) be any k-tuple of symbols chosen from Z_{n+1}.

If z(1) = n, consider the vector of positions $\alpha = (\lambda(2)-1, \dots, \lambda(k)-1)$ from {1,…, n} and the vector of symbols z(2), … , z(k) from Z_n. As B satisfies the *P(n,k−1)* property, there is a permutation $\beta = \pi_0, \dots, \pi_{n-1}$ in B such that $\beta|\alpha = z(2), \dots, z(k)$. Let $\pi = \mathbf{n}, \pi_0, \dots, \pi_{n-1}$. It follows that π is in D and $\pi|\lambda = z(1), z(2), \dots, z(k)$.

Suppose, for all i, $z(i) \ne n$. Let z(1)=p. Consider the vector of positions $\alpha = (\lambda(2)-1, \dots, \lambda(k)-1)$ from {0,…, n−1} and the vector of symbols z(2), … , z(k) from Z_n. As B satisfies the *P(n,k−1)* property, there is a permutation $\beta = \pi_0, \dots, \pi_{n-1}$ in B such that $\beta|\alpha = z(2), \dots, z(k)$. Let $\pi = \text{swap}(< \mathbf{n}, \pi_0, \dots, \pi_{n-1} >, n, p)$. It follows that π is in D and $\pi|\lambda = z(1), z(2), \dots, z(k)$.

Suppose z(1)=p≠n and z(i)=n, for some i>1. Form the new vector of symbols w(2), … , w(k) from Z_n, where, for all j≠i, w(j)=z(j), and w(i)=p, and create the vector α = (λ(2)−1, … , λ(k)−1) from {0,…, n−1}. As B satisfies the *P(n,k−1)* property, there is a permutation $\beta = \pi_0, \dots, \pi_{n-1}$ in B such that β|α = w(2), … , w(k). Let $\pi = \text{swap}(< \mathbf{n}, \pi_0, \dots, \pi_{n-1} >, n, p)$. As the swap operation puts p in position 0 and **n** in position λ(i) in π, it follows that π is in D and π|λ = z(1), z(2),s … , z(k).

So, C satisfies the *PZ(n+1,k)* property and has (n+1)r elements. Hence, C = A ∪ D has m+(n+1)r elements and satisfies the *P(n+1,k)* property. □

Theorem 5. For $n, k \geq 1$ it holds that $F(n+1,k) \leq f(n+1,k) + (n+1)\cdot F(n,k-1)$

Theorem 5 follows directly from Lemma 4.
We note that the result of Theorem 5 can be improved for certain values of n and k.

Corollary 1. Let A, B, C be sets of permutations over Z_n satisfying the *P(n,k)* property, the *R(n,k−1)* and the *P(n,k−1)* property, respectively. Let |A|=m, |B|=r, and |C|=q. One can construct a set I of permutations over Z_{n+1} satisfying the *P(n+1,k)* property such that |I|=m+n(r+q).

Proof. In the proof of Lemma 3, we formed a set D of permutations over Z_{n+1} defined by D = { ***n***,π_0, π_1. … , π_{n-1} | π_0, π_1. … , π_{n-1} is in A } and constructed a set E, where E = { π_0, ***n***, π_1, … , π_{n-1}; π_0, π_1, ***n***, … , π_{n-1}; … , π_0, π_1, … , ***n***, π_{n-1}; π_0, π_1, … , π_{n-1}, ***n*** | π_0, π_1. … , π_{n-1} is in B }. It follows that H=D ∪ E satisfies the *R(n+1,k)* property and |H|=m+n·r. In Lemma 4, we joined to the set H a set J that satisfies the *P(n+1,k₀)* property. Let

$$J = \bigcup_{0 \leq t \leq n} \{\text{swap}(\pi, n, t) \mid \pi = \mathbf{n}, \pi_0, \dots, \pi_{n-1} \text{ such that } \pi_0, \dots, \pi_{n-1} \text{ is in C}\}.$$

Note that it is not necessary to include t=n in J as for any k-tuple of symbols z(1), z(2), … , z(k) from Z_{n+1} and any k-tuple of positions λ=(λ(1), λ(2). … , λ(k)) where λ(1)=0 and z(1)=n, there is a permutation π in D such that π|λ is z(1), z(2), … , z(k). So, in this case, |J|=n·q. To conclude, let I = H ∪ J. It follows that |I| = m+n(r+q). □

Theorem 6. For $n, k \geq 1$ it holds that F(n+1,k) ≤ F(n,k) + n·(F(n,k−1)+f(n,k−1)).

Theorem 6 follows directly from Corollary 1.

Theorem 7. For $n, k \geq 1$ it holds that F(n+1,k) ≤ F(n,k) + 2n·F(n,k−1).

Theorem 7 follows directly from Theorem 6, using the inequality f(n,k−1) ≤ F(n,k−1).

4. Results Based on Upper Bounds

Theorem 6 gives a reasonably good upper bound for F(n,3). In fact, using the weaker statement of Theorem 7, we show that the upper bound obtained for F(n,3) is never greater than (8n/3) times the optimum, given by lowerbound(n,3).

Theorem 8. For all n≥4, F(n,3) ≤ (8n/3) · *lowerbound*(n,3) = (8n/3)·n·(n−1)·(n−2) = (8/3)(n^4−3n^3+2n^2).

Proof (by induction on n).

For the basis case, observe that F(4,3) = 4! = 4·3·2. For the inductive step, assume that F(n,3) ≤ (8n/3) ·n· (n−1)·(n−2). By Theorem 7, with k=3, we have F(n+1,3) ≤ F(n,3) + 2n·F(n,2). Substituting, using the inequality F(n,2) ≤ 4 ·n· (n−1) from Theorem 3, and the upper bound for F(n,3), given by the inductive hypothesis, into the previous inequality, we obtain:

$$\begin{aligned} F(n+1,3) &\leq (8n/3)\cdot n\cdot (n-1)\cdot(n-2) + 8n\cdot n\cdot (n-1) \\ &= (8/3)(n^4-3n^3+2n^2) + 8\cdot(n^3-n^2) \\ &= (8/3)(n^4-n^2) \\ &= (8/3)\cdot n^2\cdot(n^2-1) \\ &= (8n/3)\cdot n\cdot(n+1)\cdot(n-1) \\ &= (8n/3)\cdot(n+1)\cdot n\cdot(n-1) \\ &= (8n/3)\cdot \mathit{lowerbound}(n+1,3). \quad \square \end{aligned}$$

In the following we have computed estimates of f(n,k), see Table 1.

n

	2	3	4	5	6	7	8
1	2	3	4	5	6	7	8
2		6	12	20	42	42	91
3			24	60	160	472	934
4				120	360	1320	4984
5					720	2520	11760
6						5040	20160
7							40320

Table 1. Estimates of f(n,k) computed using Theorem 4, where the columns are for n and the rows are for k.

Observe that values obtained for f(n,3) are better than those described by the upper bound given for F(n,3) in Theorem 8. For example, we have f(6,3) ≤ 160, f(7,3) ≤ 472, and f(8,3) ≤ 934, as given in Table 1. The computation showing that f(6,3) ≤ 160 was given earlier as an example (Example 2) of the use of Theorem 4. To obtain f(7,3) ≤ 472, we use Theorem 4 again. However, to do this one must first use Theorem 5, which gives a bound for the needed value F(6,3) using the computed bound for f(6,3). Specifically, as F(5,2) = 20, it follows from Theorem 5 that $F(6,3) \le f(6,3) + 6\cdot F(5,2) \le 160 + 6\cdot 20 = 280$. (Actually, using details of this particular case, we get $F(6,3) \le 220$.) So, by Theorem 4, we obtain $f(7,3) \le F(6,3) + 6\cdot f(6,2) \le 220 + 6\cdot 42 = 472$.

Similarly, the computation of f(8,3) is done by Theorem 4, but again we must first obtain a value for F(7,3) from f(7,3) using Theorem 5. Specifically, by Theorem 5, $F(7,3) \le f(7,3) + 7\cdot F(6,2) = 472 + 7\cdot 42 = 766$. (Actually, using details of this particular case, we get $F(7,3) \le 640$.) So, by Theorem 4, $f(8,3) \le F(7,3) + 7\cdot f(7,2) \le 640 + 7\cdot 42 = 934$.

Observe that computing a better value for a single entry, such as f(6,2), will improve our upper bounds almost everywhere. That is, as we use a recursive technique to compute values of F(n,k), any improvement for a particular value, say of n=a and k=b, will improve values of F(n,k), for all n≥a and k≥b as well.

Theorem 9. For all n≥5,

$$
\begin{aligned}
F(n,4) &\le (3n^2/2)\cdot \text{lower-bound}(n,4) \\
&= (3n^2/2)\cdot n\cdot(n-1)\cdot(n-2)\cdot(n-3) \\
&= (3/2)(n^6-6n^5+11n^4-6n^3).
\end{aligned}
$$

Proof (by induction on n).

For the basis case, observe that F(5,4)=5!=120. Then for the inductive step, assume that $F(n,4) \le (3n^2/2)\cdot n\cdot(n-1)\cdot(n-2)\cdot(n-3)$. By Theorem 7, with k=4, we have $F(n+1,4) \le F(n,4) + 2n\cdot F(n,3)$. Substituting, using the inequality $F(n,3) \le (8/3)(n^4-3n^3+2n^2)$ from Theorem 8, and the upper bound for F(n,4) from the inductive hypothesis, into the previous inequality, we obtain:

$$
\begin{aligned}
F(n+1,4) &\le (3n^2/2)\cdot n\cdot(n-1)\cdot(n-2)\cdot(n-3) + \\
&\qquad 2n\cdot(8/3)(n^4-3n^3+2n^2) \\
&= (3/2)n^6 - (11/3)n^5 + (1/2)n^4 + (5/3)n^3 \\
&= (3/2)(n^6 - (22/9)n^5 + (1/3)n^4 + (10/9)n^3) \\
&< (3/2)(n^6 - (8/3)n^4 + (2/3)n^3 + 2n^2 + (4/3)n) \\
&= (3(n+1)^2/2)\cdot(n+1)\cdot n\cdot(n-1)\cdot(n-2) \\
&= \text{lowerbound}(n+1,4)
\end{aligned}
$$

where the strict inequality is true, for all n≥5.

□

So, the values we compute for f(n,4), for all n≥5, are at worst $(3n^2/2)$ times optimum. Actually, we can show that the values actually computed for F(n,4), using Theoorem 4, are much better than indicated by the closed form expression $(3n^2/2)\cdot$*lowerbound*(n,4). This can be seen by the values shown in Table 1.

It should also be noted that the techniques used in the proofs of Theorems 8 and 9 can also be used to obtain closed form upper bounds for F(n,k), for k>4. These upper bounds are significant improvements on the combinatorial upper bound given earlier in the introduction. The combinatorial upper bound given earlier stated, $F(n,k) \le k!\binom{n}{k}\binom{n}{k}$. Specifically, for k=3, the term $\binom{n}{k}$, i.e. $\binom{n}{3}$, is cubic. So, the combinatorial upper bound for F(n,3) in the introduction, namely $F(n,3) = 3!\binom{n}{3}\binom{n}{3}$, is an $\Omega(n^3)$ factor times the lower bound for F(n,3), namely $3!\binom{n}{3}$.

As the upper bound given for F(n,3) in Theorem 8 is an O(n) factor times the lower bound for F(n,3), a significant improvement has been described. Similar statements about the relative sizes of the upper bounds for F(n,k) can be made for k≥4.

5. Conclusions and Work in Progress

We have given new upper bounds on the size of k-tuple covers which are significant improvements on earlier bounds. Our work is initially motivated by computing fault tolerance bounds for the multi-dimensional star network [10]. As k-tuple covers have a direct application to this problem, our results also contribute significantly to a better understanding of this aspect of the use of star networks. The results reported in Table 1 are significant for practical applications of star networks.

Observe that, for each fixed natural number k, an algorithm polynomial in n exists to verify whether a given set S of permutations is a k-tuple cover for Z_n, when |S| is polynomial in n. Such an algorithm need only check that, for every k-tuple z(1), z(2). … , z(k) of distinct elements from Z_n and any k-tuple λ of distinct positions chosen from Z_n there is a permutation π in S such that $\pi|\lambda$ = z(1), z(2), … , z(k). As there are $\binom{n}{k}$ different k-tuples of distinct positions and $k!\binom{n}{k}$ different k-tuples of distinct elements from Z_n, and these are polynomial in n for any fixed k, verifying that such a permutation exists in S is possible in polynomial time.

Consequently, the problem of deciding, for a given natural number s, whether there is a k-tuple cover of size s for Z_n, can be done non-deterministically in time polynomial in n. That is, one can guess a set of permutations S, where |S| =s, then use the deterministic polynomial verifier described above to determine whether the set S is a k-tuple cover. So, this problem is in the class *NP*, but it is unknown whether it is in the corresponding deterministic class *P*, or whether it is *NP-complete*.

We conjecture that the upper bounds for f(n,k) and F(n,k) described here can in many cases be improved. We are currently working on improvements. We also conjecture that the combinatorial lower bound, namely F(n,k) ≥ k! $\binom{n}{k}$, can be improved in many cases. We have been able to show that F(6,2) is, in fact, larger than 3! $\binom{6}{2}$ = 6·5 = 30, but it is, as yet, unclear whether the technique used can be extended to describe better lower bounds for infinitely many cases.

References

[1] M. Aigner and G. Ziegler, G. *Proofs from THE BOOK*. Berlin, New York: Springer (2003). ISBN 3-540-40460-0.

[2] S.B. Akers, D. Harel, and B. Krishnamurthy, "*The star graph: an attractive alternative to the n-cube*," Proc. *Int'l Conf. Parallel Processing,* 1987, pp. 393-400.

[3] S.B. Akers and B. Krishnamurthy, "A Group-Theoretic Model for Symmetric Interconnection Networks", *IEEE Transactions on Computers,* Volume 38, Issue 4 (April 1989), Pages: 555 – 566.

[4] N. Bagherzadeh, M. Dowd and N. Nassif, "Embedding an Arbitrary Binary Tree into the Star Graph", *IEEE Transactions on Computers,* 45(4): 475-481, 1996.

[5] G. E. Box, J. S. Hunter, J.S., W. G. Hunter, "Statistics for Experimenters: Design, Innovation, and Discovery", 2nd Edition, Wiley, 2005, ISBN: 0471718130

[6] S. R. Dalal, C. L. Mallows "Factor Covering Designs for Testing Software", *Technometrics*, 40(3),1998, pp. 234–243.

[7] S. Fujita, "A Fault-Tolerant Broadcast Scheme in the Star Graph under the Single-Port, Half-Duplex Communication Model", *IEEE Transactions on Computers*, v.48 n.10, p.1123-1126, October 1999.

[8] N. Imani, H. Sarbazi-Azad, and S.G. Akl, "On some combinatorial properties of star graph" *ISPAN 2005, Proceedings of the 8th International Symposium on Parallel Architectures, Algorithms and Networks,* pp. 58-65, 2005.

[9] S. Latifi, "On the fault-diameter of the star graph", *Information Processing Letters*, Volume 46 , Issue 3 (June 1993), pp. 143 – 150, 1993.

[10] S. Latifi, "A study of fault tolerance in star graphs", *Information Processing Letters*, Vol. 102 , pp. 196 – 200, 2007.

[11] T.K. Li, J.M. Tan, L. Hsu, "Hyper Hamiltonian laceablility on edge fault star graph", Information Sciences, vol. 165, issues 1-2 (September 2004), pp. 59-71, 2004.

[12] Z. Miller, D. Pritikin, I. H. Sudborough, "Near Embeddings of Hypercubes into Cayley Graphs on the Symmetric Group", *IEEE Trans. Comput.* 43, 1, pp. 13-22, 1994.

[13] L. Morales, I. H. Sudborough, "Comparing Star and Pancake Networks." *Lecture Notes in Computer Science, vol. 2566. (2002)*, pp.18-36, Springer-Verlag, New York.

[14] S. Sur, P. K. Srimani, "A Fault-Tolerant Routing Algorithm for Star Graph Interconnection Network, " *Int'l Conf. Parallel Processing,* (3) 1991: 267-270.