

Multiplicative Improvements in Network Reliability Bounds

Timothy B. Brecht

Department of Computer Science, University of Toronto, Toronto, Ontario, M5S 1A1, Canada

Charles J. Colbourn

Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada

Multiplicative inequalities for reliability bounds are derived, by observing that certain reliability measures are positively correlated. These inequalities can be used to obtain substantial improvements on available bounds for network reliability.

1. BACKGROUND AND MOTIVATION

In the network design process, one goal is to select a network topology which is highly reliable. Although there is no universally accepted measure of reliability, the most widely used definition is a probabilistic one. The network is modelled as a *probabilistic graph* $G = (V, E)$, in which V is a set of nodes representing sites in the network, and E is a collection of undirected edges representing bidirectional point-to-point communication links. Nodes are not susceptible to failure, but edges are. Each edge e operates with some known probability p_e . In this setting, the *reliability* is the probability that the network can support some desired network operation, when edges fail independently according to the given probabilities. Three standard reliability measures arise in this way. An *all-terminal* operation requires that every pair of nodes has a path of operational edges connecting them, and *all-terminal* reliability is the probability that such an event occurs in the network. A *two-terminal* operation for specified nodes s and t requires that there be (at least) a path of operational edges connecting s with t ; *two-terminal* reliability is then the probability of this event. Finally, a *k-terminal* operation requires, for a specified set K of k *target* nodes, that every pair of target nodes has a path of operational edges connecting them.

Numerous techniques have been developed for computing these three reliability

measures [4], but each method requires exponential time for general networks. In fact, it is #P-complete to compute any of the three measures, and hence for large scale network design we must resort to approximation strategies. One main line of investigation here has studied the development of efficiently computable lower and upper bounds on the reliability measures. If these bounds can be made sufficiently tight, the bounds alone may suffice in the design process; we need only distinguish between candidate network topologies.

With this in mind, a number of general strategies have been developed for obtaining bounds. The most prevalent strategy is *subgraph counting*, pioneered by Van Slyke and Frank [11]. Here one makes the additional assumption that every edge has the same probability of operation, and then computes bounds on the number of i -edge subnetworks which support the network operation, for each i . A number of powerful combinatorial theorems lead to bounds on these subgraph counts, and thus to bounds on the reliability (see [5]). Van Slyke and Frank [11] used this method to develop the *Kruskal-Katona* bounds, which apply in each of the three reliability problems mentioned. Subsequently, Ball and Provan [1] used a theorem of Stanley to obtain much tighter bounds, which apply to all-terminal reliability, but not to the two-terminal or k -terminal problems.

The second major strategy, pioneered by Poleskii [9], employs *edge-packings* of graphs (see [2,6]). In this approach, one partitions the edges of the network into subnetworks. A lower bound is obtained by computing the probability that at least one of the subnetworks supports the desired operation; an upper bound is obtained taking subnetworks whose removal prevents network operation, and computing the probability that at least one of the subnetworks fails. While much simpler than the subgraph counting bounds, the edge-packing bounds are competitive. In fact, neither strategy leads at present to bounds which are uniformly better than the other.

We consider another general strategy here, which differs from the two basic strategies in an important respect; we assume that bounds from the two basic methods can be readily computed, and develop some statistical inequalities which tighten the resulting bounds. In a previous paper, we developed certain additive inequalities and demonstrated that they often result in a substantial tightening of reliability bounds [3]. In this paper, we establish some simple multiplicative inequalities. The inequalities are always at least as tight as the previously known additive inequalities, and are easily computed. In Section 5 we discuss computational issues and the ramifications for bounding network reliability.

2. ADDITIVE INEQUALITIES

We first introduce required definitions, and recall the additive inequalities developed in [3]. Let $G = (V, E)$ be a probabilistic graph, with each edge e having operation probability p_e . A *state* S of G is an assignment to each edge of a state, either "operational" or "failed"; the network state is operational if it supports the desired network operation, failed otherwise. A state is often represented as the subset of operational edges. The probability of a network state is the product of operation probabilities for edges operating in this state, and failure probabilities of edges failed in this state. An *event* is a Boolean proposition mapping states to $\{0, 1\}$; a particular event of interest

to us is denoted $C(K)$, the event that all nodes in a target set K can communicate via paths of operational edges.

For an event ϕ , we define $\Pr\{G;\phi\}$ to be the probability that event ϕ occurs in G ; more precisely,

$$\Pr\{G;\phi\} = \sum_{\text{state } S} \Pr\{S\}\phi(S)$$

$\Pr\{G;\psi \wedge \phi\}$ denotes the probability that events ψ and ϕ occur simultaneously; $\Pr\{G;\psi|\phi\}$ denotes the conditional probability that event ψ occurs given that event ϕ occurs. Notice that $\Pr\{G;C(V)\}$ is all-terminal reliability, $\Pr\{G;C(\{s,t\})\}$ is two-terminal reliability, and $\Pr\{G;C(K)\}$ is k -terminal reliability.

The main purpose of adopting this uniform framework for the three problems is to develop inequalities which relate the reliability measures. The main additive inequality, used in [3], can be stated as follows:

Theorem 2.1. For any probabilistic graph G and sets of nodes K_1, K_2 with $K \subseteq K_1 \cup K_2$ and $K_1 \cap K_2 \neq \emptyset$, $\Pr\{G;C(K)\} \geq \Pr\{G;C(K_1)\} + \Pr\{G;C(K_2)\} - 1$. ■

This additive inequality is easy to prove, either directly or as a consequence of certain reliability bounds assuming statistical dependence. The proof of Theorem 2.1 is by no means remarkable, but the resulting improvements in reliability bounds are! Improvements over the subgraph counting and edge-packing approaches are found for each of the three reliability problems [3].

A strengthening of Theorem 2.1 is therefore of significant interest. In the special case when every edge has operation probability equal to .5, a lemma of Kleitman [8] shows that

$$\Pr\{G;C(K_1) \wedge C(K_2)\} \geq \Pr\{G;C(K_1)\}\Pr\{G;C(K_2)\}$$

In the next section, we therefore develop a multiplicative inequality which improves on Theorem 2.1, by extending Kleitman's lemma to all possible assignments of probabilities.

3. MULTIPLICATIVE INEQUALITIES

The proof of Theorem 2.1 requires no knowledge of what the events $C(K_1)$ and $C(K_2)$ are, but requires only that together they imply $C(K)$. Hence worst-case assumptions are made about the correlation of the two events. However, in the actual context at hand, the correlation of the two events cannot realize this worst case; informally, we can observe that if all nodes in K_1 can communicate, this should not decrease the probability that all nodes in K_2 can communicate. The multiplicative inequalities arise as a result of formalizing (and proving) this observation.

We introduce a sequence of preliminary lemmas.

Lemma 3.1. $\Pr\{G;C(K')\} \geq \Pr\{G;C(K)\}$ for all $K' \subseteq K$.

Proof. For any state in which $C(K)$ holds, $C(K')$ also holds. ■

For an edge e , we define $\text{Up}(e)$ to be the event which maps a network state to 1 if e is operational in this state, 0 otherwise; then observe

Lemma 3.2. $\Pr\{G; \text{Up}(e)|C(K)\} \geq \Pr\{G; \text{Up}(e)\}.$

Proof. Consider all states for which $C(K)$ holds. Partition this collection of states into two groups, one called U with states having e operational and one called D with states having e failed. We need only show that the probability of obtaining a state of U is at least $p_e/(1 - p_e)$ times the probability of obtaining a state from D . But this follows from the observation that if $S \in D$, $S \cup \{e\} \in U$. ■

Let G be a probabilistic graph and K a target set for G . We use $G \cdot e$ and $G - e$ to denote the results of contracting and deleting the edge e , respectively. $K \cdot e$ denotes the image of K under contraction of the edge e .

Lemma 3.3. $\Pr\{G \cdot e; C(K \cdot e)\} \geq \Pr\{G - e; C(K)\}.$

Proof. $G - e$ and $G \cdot e$ have the same edge set, and thus their states are in a natural 1-1 correspondence. Observe that every state connecting K in $G - e$ also connects $K \cdot e$ in $G \cdot e$. ■

Lemma 3.4. $\Pr\{G; C(K_2)|C(K_1)\} = \Pr\{G; \text{Up}(e)|C(K_1)\} \Pr\{G \cdot e; C(K_2 \cdot e)|C(K_1 \cdot e)\} + \Pr\{G; \text{Up}(e)|C(K_1)\} \Pr\{G - e; C(K_2)|C(K_1)\},$ for any graph G with target sets K_1, K_2 and edge e .

Proof. Partition the states of G in which $C(K_1)$ holds into two groups, according to whether e is operational or failed; observe:

$$\Pr\{G; C(K_2)|C(K_1)\} = \Pr\{G; (\text{Up}(e) \wedge C(K_2))|C(K_1)\} + \Pr\{G; (\overline{\text{Up}(e)} \wedge C(K_2))|C(K_1)\}$$

The event $(\text{Up}(e) \wedge C(K_2))|C(K_1)$ holds exactly when $\text{Up}(e)|C(K_1)$ holds in G and the event $C(K_2 \cdot e)|C(K_1 \cdot e)$ holds in $G \cdot e$. Moreover, these two events are statistically independent, since the first refers only to the edge e and the second refers only to the edges of $G \cdot e$, i.e., edges other than e . Simplifying the event $(\overline{\text{Up}(e)} \wedge C(K_2))|C(K_1)$ in a similar way yields the statement of the lemma. ■

Corollary 3.5. $\Pr\{G; C(K)\} = \Pr\{G; \text{Up}(e)\} \Pr\{G \cdot e; C(K \cdot e)\} + \Pr\{G; \overline{\text{Up}(e)}\} \Pr\{G - e; C(K)\}$ for any graph G , target set K , and edge e .

Proof. Take $K_2 = K$ and K_1 an arbitrary singleton set in the statement of Lemma 3.4. ■

Corollary 3.5 is the well-known factoring theorem, which has been widely used in reliability investigations (see, for example, [10]). Finally, we are in a position to state the key lemma, which formalizes the informal observation made at the outset:

Lemma 3.6. For G a connected graph, and K_1, K_2 target sets of G , $\Pr\{G; C(K_2)|C(K_1)\} \geq \Pr\{G; C(K_2)\}.$

Proof. We use induction on m , the number of edges in G . The base case is when $m = n - 1$, when G is a tree. G contains a unique minimal subtree T_1 connecting K_1 , and a unique minimal subtree T_2 connecting K_2 . Now $\Pr\{G; C(K_2)\}$ is the product of the edge operation probabilities over edges in T_2 . Next observe that every state of G in which $C(K_1)$ holds has all edges of T_1 operational; hence we can contract each

edge of T_1 in turn to form a graph $G \cdot T_1$, with the image of K_2 after the sequence of contractions being $K_2 \cdot T_1$. Then $\Pr\{G; C(K_2) | C(K_1)\} = \Pr\{G \cdot T_1; C(K_2 \cdot T_1)\}$. Now contracting T_1 may leave T_2 unaffected, or may contract it onto a subtree T'_2 . In either case, the probability is the product of edge operation probabilities over edges in the tree/subtree, and hence is at least as great as the probability of T_2 . This completes the base case.

Suppose then that the statement holds for all graphs with at most m edges, and suppose G has $m + 1$ edges. Using Corollary 3.5, select an edge e whose removal does not disconnect G and write

$$\begin{aligned} \Pr\{G; C(K_2)\} &= \Pr\{G; \overline{\text{Up}(e)}\} \Pr\{G \cdot e; C(K_2 \cdot e)\} \\ &\quad + \Pr\{G; \overline{\text{Up}(e)}\} \Pr\{G - e; C(K_2)\} \end{aligned} \quad (1)$$

Using Lemma 3.4, write

$$\begin{aligned} \Pr\{G; C(K_2) | C(K_1)\} &= \Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} \Pr\{G \cdot e; C(K_2 \cdot e) | C(K_1 \cdot e)\} \\ &\quad + \Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} \Pr\{G - e; C(K_2) | C(K_1)\} \end{aligned} \quad (2)$$

We must show that

$$\Pr\{G; C(K_2) | C(K_1)\} - \Pr\{G; C(K_2)\} \geq 0$$

Equations (1) and (2) show this is equivalent to

$$\begin{aligned} &\Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} \Pr\{G \cdot e; C(K_2 \cdot e) | C(K_1 \cdot e)\} \\ &\quad - \Pr\{G; \overline{\text{Up}(e)}\} \Pr\{G \cdot e; C(K_2 \cdot e)\} \\ &\quad + \Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} \Pr\{G - e; C(K_2) | C(K_1)\} \\ &\quad - \Pr\{G; \overline{\text{Up}(e)}\} \Pr\{G - e; C(K_2)\} \geq 0 \end{aligned}$$

Using Lemmas 3.2 and 3.3, it suffices to show that

$$\begin{aligned} &\Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} [\Pr\{G \cdot e; C(K_2 \cdot e) | C(K_1 \cdot e)\} - \Pr\{G \cdot e; C(K_2 \cdot e)\}] \\ &\quad + \Pr\{G; \overline{\text{Up}(e)} | C(K_1)\} [\Pr\{G - e; C(K_2) | C(K_1)\} - \Pr\{G - e; C(K_2)\}] \geq 0 \end{aligned}$$

Now by the induction hypothesis, both terms in square brackets are nonnegative, and hence the inequality holds. ■

This key lemma underpins the multiplicative inequality:

Theorem 3.7. Let G be a graph, and K_1, K_2 be target sets of G satisfying $K_1 \cap K_2 \neq \emptyset$. Then for any $K \subset K_1 \cup K_2$, $\Pr\{G; C(K)\} \geq \Pr\{G; C(K_1)\} \Pr\{G; C(K_2)\}$.

Proof. If G is disconnected, it may happen that not all targets in K_1 (K_2) appear in the same component; in this case, the inequality is trivial. Otherwise, if all targets in K_1 and in K_2 lie in the same component, so do all targets in K ; hence we need only consider this connected component. Henceforth assume that G is connected.

Using Lemma 3.1, we need only consider the case $K = K_1 \cup K_2$. Now since $K_1 \cap K_2 \neq \emptyset$, $\Pr\{G; C(K)\} = \Pr\{G; C(K_1) \wedge C(K_2)\}$. Equivalently, $\Pr\{G; C(K)\} = \Pr\{G; C(K_1) \Pr\{G; C(K_2) | C(K_1)\}\}$. Then Lemma 3.6 completes the proof. ■

It is an easy exercise to check the Theorem 3.6 is at least as strong as Theorem 2.1, and is strictly stronger unless one of $\Pr\{G;C(K_1)\}$ or $\Pr\{G;C(K_2)\}$ is exactly one.

4. IMPROVING RELIABILITY BOUNDS

In [3], a simple strategy was developed for using Theorem 2.1 to improve reliability bounds; we extend that strategy to employ the multiplicative inequalities instead. Consider a graph $G = (V, E)$. We have a number of basic bounding strategies for obtaining two-terminal lower bounds for pairs of vertices in G ; in particular, the subgraph counting Kruskal-Katona bounds apply [11], and edge-packing bounds apply [2]. Using these methods, compute, for each pair $\{x, y\}$ of vertices in G , the best lower bound on their two-terminal reliability, L_{xy} . To put theorem 3.6 in a computationally useful form, we prove the following lemma:

Lemma 4.1. Let $SG(K)$ be a subgraph of the complete graph on vertex set V which connects all nodes in a target set K of $G = (V, E)$. Then $\Pr\{G;C(K)\} \geq \prod_{\{x,y\} \in SG(K)} L_{xy}$.

Proof. We may assume that $SG(K)$ contains no cycles, and that every degree one vertex is in K ; that is, $SG(K)$ is a *Steiner tree* for K in G . If $SG(K)$ contains vertices not in K , observe by Lemma 3.1 that $\Pr\{G;C(K)\} \geq \Pr\{G;C(S)\}$, where S is the vertex set of $SG(K)$. We therefore need only consider the case $S = K$. Let z be a degree one vertex of $SG(K)$, and let x be its neighbor in $SG(K)$. Let $K_1 = S - \{z\}$ and $K_2 = \{x, z\}$ and apply Theorem 3.6 to establish that $\Pr\{G;C(K)\} \geq L_{xz} \Pr\{G;C(K_1)\}$. Repeating this process until $|K_1| = 1$ gives the statement of the theorem. ■

This reformulation enables us to develop bounds on k -terminal reliability and all-terminal reliability using only bounds on two-terminal reliability. In the application to all-terminal reliability, observe that the best bound from Lemma 4.1 arises when one selects $SG(K)$ to be a spanning tree for which the product of its two-terminal lower bounds is maximum. Until this point, all of the computation requires only polynomial time, and hence it is important to note that spanning trees whose edge weight product is maximum are easily found; one can transform to the usual maximum weight spanning tree problem by taking logarithms of the edge weights. Hence Lemma 4.1 gives an efficient method for finding reliability bounds in the all-terminal case.

In the case of two-terminal reliability, one might expect no improvement since the basic bounds used are two-terminal. However, let us remark that one can derive a simple multiplicative triangle inequality (which is similar in spirit to the additive triangle inequality of [3]):

Lemma 4.2. For any three vertices x, y, z of G , $\Pr\{G;C(\{x, z\})\} \geq L_{xy}L_{yz}$.

Proof. By Lemma 4.1 $\Pr\{G;C(\{x, y, z\})\} \geq L_{xy}L_{yz}$; applying Lemma 3.1 completes the proof. ■

Improved bounds on two-terminal reliability can therefore also be efficiently obtained, by repeatedly applying Lemma 4.2 to improve the values of the $\{L_{xy}\}$ until no further improvement is possible.

Finally, note that for k -terminal reliability, the best bound arises from Lemma 4.1

by taking a Steiner tree for K whose product of edge weights is maximum; however, finding such Steiner trees is NP-hard [7]. Nevertheless, to obtain reasonable bounds, it suffices to find a Steiner tree near the maximum. The details here parallel the additive case quite closely, and so are omitted; see [3].

Theorem 3.7 also leads to *upper* bounds on reliability measures as follows. We can write

$$\Pr\{G;C(K_1)\} \leq \frac{\Pr\{G;C(K)\}}{\Pr\{G;C(K_2)\}}$$

for $K_1 \cup K_2 \subset K$ and $K_1 \cap K_2 \neq \emptyset$. Although none of these quantities are known exactly, we can overestimate $\Pr\{G;C(K)\}$ by using an efficiently computable upper bound, and underestimate $\Pr\{G;C(K_2)\}$ by using the lower bound just computed. This leads to an upper bound on $\Pr\{G;C(K_1)\}$. In order to efficiently compute the upper bound on $\Pr\{G;C(K)\}$, we can always take K to contain all vertices, and employ the Ball-Provan upper bound on all-terminal reliability. While this approach is somewhat less direct than that taken for the lower bounds, it is of interest that the multiplicative inequalities lead to improved upper bounds, while the additive inequalities do not.

5. SOME COMPUTATIONAL RESULTS

The main goal of deriving methods for improving bounds is to achieve a substantial tightening of the bound while retaining the existence of efficient algorithms for determining the bound. We have seen that the multiplicative inequalities can be employed in an efficient manner, but it remains to see that some substantial improvement results. We content ourselves with some small examples on a skeleton of the 1979 Arpanet, a network whose analysis is of some practical importance (primarily as a concrete

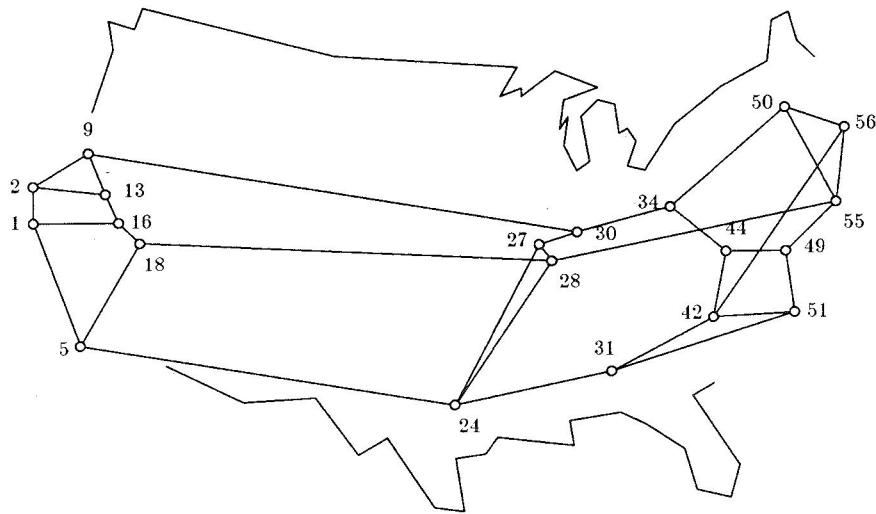


FIG. 1. A skeleton of the 1979 Arpanet.

TABLE I. Arpa skeleton: two-terminal example.

p	Mincost	Additive improvement	Multiplicative improvement
.3	.016350	.016350	.019268
.5	.127960	.127960	.153597
.7	.470107	.470107	.539538
.9	.938300	.961890	.962369
.95	.989620	.994491	.994500
.98	.999205	.999615	.999615
.99	.999895	.999950	.999950

example of bidirectional point-to-point communication networks in existence). This network is depicted in Figure 1.

We first give a two-terminal example, using nodes 5 and 56. The best basic bound here is the mincost bound of [2]. We report the mincost bound, the results of using additive improvement [3], and the result of using the multiplicative improvements, assuming that each edge has operation probability p (see Table I). Whereas the additive inequalities only afforded improvements for high p , the multiplicative inequalities yield improvements throughout the range. The improvements seen here are substantial except for p near zero or one.

In k -terminal problems, two main strategies are competitors to the multiplicative inequalities. The first is the Ball-Provan all-terminal lower bound, which provides a lower bound on k -terminal reliability for every set of target nodes. The second is the additive method used earlier. We used a heuristic method for finding Steiner trees [12] to implement the additive and the multiplicative bounds for a 4-terminal problem in the Arpa skeleton, using nodes 5, 9, 34, and 55. Results for selected values of p are given in Table II. Except for p values very near 1, the multiplicative inequalities yield the best available bound on this 4-terminal problem. It is important to remark that this good behavior is observed despite the fact that heuristics rather than exact methods are used to find the Steiner tree used.

In the all-terminal problem, the best available bound typically is the Ball-Provan

TABLE II. Arpa skeleton; 4-terminal problem.

p	Ball-Provan	Additive improvement	Multiplicative improvement
.3	.000018	0	.000578
.5	.009154	0	.023955
.7	.154748	0	.284424
.9	.861718	.918666	.921048
.95	.982022	.987908	.987961
.98	.999227	.999141	.999142
.99	.999936	.999889	.999889

bound [1]. In our example, the Ball-Provan bound outperforms even the multiplicative improvements at present, since they are using two-terminal lower bounds. For the Arpa skeleton with $p = .98$, the Ball-Provan bound yields .999227, while the multiplicative bound is only .998247. It appears that better two-terminal basic bounds are required if one wants to outperform the Ball-Provan bounds with this strategy.

Computational results for upper bounds are also easily obtained. No improvements are obtained in the all-terminal case, but worthwhile improvements result in the two-terminal and k -terminal problems.

VI. CONCLUDING REMARKS

The multiplicative inequalities developed here improve not only on the additive inequalities used earlier, but provide significant improvements on the basic bounding strategies which were previously available. While the observations used are straightforward, the observed tightening of reliability bounds is substantial. It appears that the strategy here is successful largely because it considers the local structure of the network, whereas the basic bounding strategies consider only "global" information.

ACKNOWLEDGMENT

Research of the second author is supported by NSERC Canada under Grant A0579.

References

- [1] M. O. Ball and J. S. Provan, Bounds on the reliability polynomial for shellable independence systems. *SIAM J. Alg. Disc. Meth.* **3** (1982) 166–181.
- [2] T. B. Brecht and C. J. Colbourn, Lower bounds for two-terminal network reliability. *Discrete Appl. Math.*, **21** (1988) 185–198.
- [3] T. B. Brecht and C. J. Colbourn, Improving reliability bounds in computer networks. *Networks* **16** (1986) 369–380.
- [4] C. J. Colbourn, *The Combinatorics of Network Reliability*, Oxford University Press, New York (1987).
- [5] C. J. Colbourn, The reliability polynomial. *Ars Combinatoria* **21A** (1986) 31–58.
- [6] C. J. Colbourn, Edge-packings of graphs and network reliability. *Discrete Math.*, **72** (1988) 49–61.
- [7] M. R. Garey and D. S. Johnson, *Computers and Intractability*. Freeman, San Francisco (1979).
- [8] D. J. Kleitman, Families of non-disjoint subsets. *J. Comb. Theory* **1** (1966) 153–155.
- [9] V. P. Poleskii, A lower boundary for the reliability of information networks. *Prof. Inf. Trans.* **7** (1971) 165–171.
- [10] A. Satyanarayana and M. K. Chang, Network reliability and the factoring theorem. *Networks* **13** (1983) 107–120.
- [11] R. M. Van Slyke and H. Frank, Network reliability analysis I. *Networks* **1** (1972) 279–290.
- [12] R. T. Wong, A dual ascent approach for Steiner tree problems on a directed graph. *Math. Programming* **28** (1984) 271–287.

Received November 1986

Accepted June 1988

