Reserving Resilient Capacity for a Single Commodity with Upper Bound Constraints

G. Brightwell
Centre for Discrete and Applicable Mathematics
London School of Economics
London, U.K.

G. Oriolo
Dip. di Informatica, Sistemi e Produzione
University of "Tor Vergata"
Rome, Italy

F.B. Shepherd
Bell Laboratories
Lucent Technologies
Murray Hill, NJ, USA

Abstract

Continuing research begun in [2], we investigate problems of reserving capacity in the arcs of a network, subject to the constraint that, on the failure of any one arc, there is enough reserved capacity on the remaining arcs to support a flow of value T from a source s to a destination t. We also impose upper bounds on the amount of capacity we may reserve on the arcs: this alters the nature of the problem.

In the case where each arc has the same upper bound, we investigate the strategy of finding the minimum-cost reservation that is itself an acyclic (s, t) flow: we show that such a reservation is easy to find, always has a simple form, and has a cost at most twice that of the optimal solution.

In the case where each arc has its own upper bound, we explain why no such results can hold, but we do give an efficient algorithm for the case where we are asked for a reservation on a fixed set of arc-disjoint paths.

We consider the case where we are free to reserve on each arc as much capacity as we want but only in bundles of fixed size.

1 Introduction and Definitions

We investigate various problems of a type we call *resilience*. The theme is that we are to reserve capacity in the arcs of a network so that, if any one arc of the network fails, there is sufficient capacity reserved in the remaining arcs to allow a flow of value T from a single source s to a single destination t. Each arc has a cost, which is to be interpreted as a per-unit cost for reserving capacity in that arc.

We observed in a previous paper [2] that, while the problem as stated can be solved in polynomial time using linear programming techniques, the overhead and running times of such algorithms often makes them impractical, and there is no obvious simple combinatorial algorithm. Also, the integer version of the problem is NP-complete. We are thus led to seek simple algorithms to provide solutions that are approximately optimal.

In [2], we discussed the utility of finding the optimal reservation of a certain particularly simple type, namely one derived by reserving capacity along a set of *diverse* (arc-disjoint) paths from s to t. We showed that finding such a *diverse paths reservation* was simple in practice, and was at worst twice the cost of the overall optimum. Furthermore, the simplicity of the solution obtained may be an asset in many practical settings.

Here we discuss the more general scenario where the arcs of the network are also equipped with upper bounds on the capacity we may reserve, possibly uniform over all arcs and possibly not. Restrictions of this type are likely to be present in any practical setting and, as we shall see, they seriously alter the nature of the problem. In this case, the strategy of looking for a diverse paths reservation does not have such pleasant properties: our main aim in this paper is to highlight a slightly different strategy, namely that of finding the minimum cost reservation that is itself an acyclic (s,t) flow – i.e., the reservation vector is itself a flow from s to t, and its support does not contain any directed cycles.

We show that such a minimum cost *acyclic flow-reservation* can easily be found, and always takes a simple form, even when there is an upper bound – the same for all arcs – on the maximum capacity of an arc; furthermore its cost is at most twice that of an optimum solution.

In the case where each arc has its own upper bound, we explain why no such results can hold; in this case we give an efficient algorithm to find the optimum solution for the most simple case where we are asked for a diverse paths reservation with the set of paths pre-specified.

One of the eventual aims of this research is to use the strategies and algorithms we are developing as part of an algorithm to treat multi-commodity flow problems where there is a resilience requirement. The principle is that, if reservations with a simple structure (e.g., diverse paths reservations or acyclic flow-reservations) provide good approximate solutions for single commodities, then restricting attention to such solutions should greatly simplify the multi-commodity problem while not unduly compromising the quality of the solution obtained. Some more discussion is given in [2]

DEFINITIONS

We start with some basic definitions and facts concerning flows in networks.

Throughout, we suppose (sometimes implicitly) that we are given a directed graph (network) $D = (V, A)$ with *node set* V and *arc set* A. We shall always assume that D comes with two nodes permanently fixed as the *source* s and the *destination* t. We denote by $\kappa(D)$ the (s,t)-edge connectivity of the network, that is the minimum size of a set $X \subseteq A$ such that in $G \setminus X$ there are no (s,t) paths. We also suppose that we are given a rational number T (usually an integer) representing the required traffic flow from s to t through the network D in the case of failure. Finally, we are also given a vector (c_a) of non-negative rational (again, usually integer) per-unit costs on the arcs a of D. Even though there may be parallel arcs, we often use the notation (u,v) for an arc from u to v, when no confusion arises.

For any $S \subseteq V$, let $\delta^+_D(S)$ (or simply $\delta^+(S)$ if the context is clear) denote the set of arcs in D with tail in S and head in $V - S$. Let $\delta^-(S)$ denote the set of arcs $\delta^+(V - S)$. For a node $v \in V$, we write $\delta^+(v)$ for $\delta^+(\{v\})$, etc. We call $S \subseteq V$ an (s,t)*-set* if $s \in S$ and $t \in V - S$.

Let $\mathbb{Q}_+$ denote the set of non-negative rational numbers, so that $\mathbb{Q}_+^A$ is the set of all assignments of non-negative rationals to each member of the arc-set A, which we frequently view as a vector. For any vector $x \in \mathbb{Q}_+^A$ and $A' \subseteq A$, we denote by $x(A')$ the sum $\sum_{a \in A'} x_a$. We let $I(x)$ denote the *support* of x, that is $I(x) = \{a \in A : x_a > 0\}$. For a set S of arcs, we denote by $e(S)$ the incidence vector of S, i.e., $e(S)_i = 1$ if $i \in S$ and 0 otherwise. For a single arc a, we write e_a for $e(\{a\})$. (A path P is considered as a set of arcs, and so also has an associated incidence vector $e(P)$.) The *cost* of a vector $x \in \mathbb{Q}_+^A$ is $\text{cost}(x) = \sum_{a \in A} c_a x_a$.

Resilient Reservations

A vector $x \in \mathbb{Q}_+^A$ of capacities is *T-resilient* if:

$$x(\delta^+(S) - e) \geq T \quad \text{for each } (s,t)\text{-set } S \text{ and } e \in \delta^+(S). \tag{1}$$

We often refer to constraint (1) as the *partial T-cut constraint* associated with the pair (S, e). We also refer to capacity vectors as *reservations*.

A collection of paths from s to t is called *diverse* if no two have a common arc. Examples of T-resilient reservations are obtained by taking any collection of $m \geq 2$ diverse paths, and reserving capacity $T/(m-1)$ on each arc in any of the paths. Any reservation obtained by choosing a collection $\{P_1, \ldots, P_m\}$ of diverse paths and reserving capacity z_i on each arc of path P_i is called a *diverse paths reservation*. Diverse paths reservations have been studied in detail in [2]; they provide simple solutions to resilience problems and, in the case where we may reserve as much capacity as we want on each arc, allow easy optimization and cost at most twice as much as the optimal T-resilient reservation.

A vector $x \in \mathbb{Q}_+^A$ is an (s,t) *flow-vector* (or simply a *flow*) if $x(\delta^+(v)) = x(\delta^-(v))$ for all $v \neq s,t$ and(therefore) $x(\delta^+(s)) - x(\delta^-(s)) = x(\delta^-(t)) - x(\delta^+(t))$. In this case, the *value* of the flow is $x(\delta^+(s)) - x(\delta^-(s))$. A flow-vector of value M is also called an *M-flow*. A reservation vector $x \in \mathbb{Q}_+^A$ is a *flow-reservation* if x is an (s,t) flow vector. Diverse paths reservations are flow-reservations.

We say a vector $x \in \mathbb{Q}_+^A$ is *acyclic* if there are no directed cycles in the support $I(x)$ of x. Therefore, x is an *acyclic flow-reservation* if x is an (s,t) flow-reservation and x is acyclic.

It is clear that there is always a minimum cost T-resilient diverse paths reservation which is an acyclic flow-reservation. Conversely, we shall show (Corollary 3.3) that there always exists a minimum cost T-resilient acyclic flow-reservation which is on diverse paths (this result does not extend to the integral case).

However, in general there are flow-reservations which cost less than any acyclic one, and general reservations which cost less than any flow-reservation: we illustrate both of these possibilities with the digraph in Figure 1, with costs as indicated. An optimal T-resilient reservation gives capacity $T/3$ to all the arcs from s to u, all the arcs from s to v, and the two arcs between u and v, and capacity T to arcs (u,t) and (v,t): the cost of this reservation is $20T/3$. An optimal T-resilient flow-reservation adds capacity $T/3$ to the arcs (t,u) and (t,v) and costs $22T/3$. Finally, an optimal T-resilient acyclic flow-reservation gives capacity T to one arc from s to u, one arc from s to v, and the arcs (u,t) and (v,t); this costs $8T$.

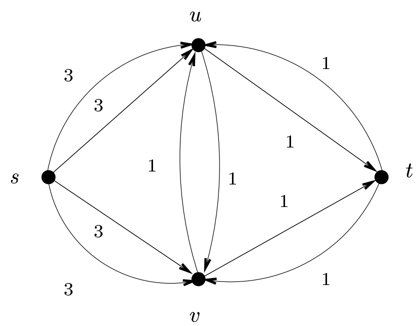

Figure 1: Per-unit costs of links as indicated.

Acyclic flow-reservations are studied in detail in Section 3; we show that they provide simple enough solutions and, in the case where we have the same upper bound on the capacity we may reserve on each arc, allow easy optimization and cost at worst twice the cost of any feasible T-resilient reservation.

Constraints on reservations

In many practical settings there are specific constraints on the reservation vector. Those we have considered in the paper are as follows.

1. Integrality: The reservation vector has to be integral.

2. Bounded Reservation: There are bounds λ_a on the capacity we may reserve on each arc a of the network. If the bounds are the same for each arc, we call the reservation *uniformly* bounded.

3. Bundled Reservation: For each arc a of the network we are free to reserve as much capacity as we want but only in bundles of μ_a units. When the size of the bundles is the same for each arc, the bundles are *homogeneous*; if all the bundles have size one this constraint reduces to integrality.

Linear programming formulations

As with many standard network flow problems, some optimization problems involving resilient reservations may be formulated as optimization problems over certain polyhedra. This has been investigated in [2], where we observe that the *resilience polyhedron*, defined as:

$$\mathcal{R}(T,D) = \left\{x \in \mathbb{Q}_+^A : x(\delta^+(S) - e) \geq T \text{ for each } (s,t)\text{-set } S \text{ and } e \in \delta^+(S)\right\} \qquad (2)$$

consists exactly of the T-resilient vectors for the network D. A consequence of this formulation is that there is a polynomial time algorithm to find a minimum cost T-resilient reservation, since the separation problem for $\mathcal{R}(T,D)$ amounts to solving at most $|A|$ maximum flow problems. Moreover, the resilience polyhedron can be described by a polynomially bounded number of variables and constraints – see [2] for further discussion.

By the same arguments, it follows that there are polynomial time algorithms to find a minimum cost T-resilient reservation if we include as additional constraints upper bounds on the capacities of the arcs and/or require that the reservation be a flow-reservation.

These linear programming algorithms are not however offered as a practical approach for finding minimum cost resilient reservation, and the remainder of the paper addresses the task of finding more direct combinatorial algorithms.

Complexity overview

In the following chart we summarize complexity results for the problem of finding minimum cost T-resilient vectors for different types of reservations and constraints. "Easy" indicates problems for which we have a *combinatorial* exact algorithm; "Polytime" indicates problems for which the only polynomial time algorithms we know arise from the linear programming formulation; "Approx" indicates NP-complete problems for which we have a combinatorial *approximation* algorithm – i.e., a combinatorial algorithm which produces a feasible output vector whose cost is within a fixed constant factor of the optimum – while "Hard" indicates problems for which no such approximation algorithm exists unless P = NP. References are given to either our earlier paper, or a result in this paper implying the assertion here.

	Reservation			
	general	flow	acyclic flow	diverse paths
Arc Constraints				
free	Polytime [2]	Polytime	Easy (3.3)	Easy [2]
integrality	Approx [2]	Approx	Easy (3.4)	Approx [2]
uniform bounds	Polytime	Polytime	Easy (3.5)	Hard (2.2)
non uniform bounds	Polytime	Polytime	?	Hard (2.2)
integrality + uniform bounds	Approx (3.6)	Approx (3.6)	Easy (3.5)	Hard (2.2)

Let us deal immediately with the various assertions made above about the problems where our reservation vector has to be a flow, but not necessarily acyclic. It was shown in [2] that it is NP-hard to solve the resilience problem when the reservation vector has to be integer. The example constructed in [2] can be augmented by adding zero-cost arcs back from the destination t to every other node, and from every other node back to s; any (integral) T-resilient reservation thus has the same cost as some (integral) T-resilient flow-reservation – it follows that the problem of finding an integral T-resilient flow-reservation is also NP-complete. Also in [2], we show that the naive algorithm that allocates capacity T to the arcs of a cheapest pair of diverse (s, t) paths gives at most twice the cost of any T-resilient reservation. As this is an integral flow-reservation, this also gives a 2-approximation algorithm for this problem.

The paper is organized as follows.

We start Section 2 by briefly reviewing some results from our earlier paper [2] on diverse paths reservations. These results suggest that using a diverse paths reservation is a good strategy for building resilient networks.

We then move to the scenario where reservations are bounded, that is there are upper bounds λ_a on the capacity we may reserve on each arc a of the network. We show that then finding a minimum cost bounded diverse paths reservation is NP-hard, and furthermore is not k-

approximable for any $k \geq 1$ even in the fractional case, and even if all the λ_a are equal (unless P = NP).

In Section 3 we discuss acyclic flow-reservations. Diverse paths reservations are particular examples of acyclic flow-reservations; we show that there always exists a minimum cost acyclic flow-reservation which is easy to find and is either a diverse paths reservations or, in the integral case, a convex combination of *two* diverse paths reservations. This extends to acyclic flow-reservations many of the nice practical features of diverse paths. Moreover, the optimization problem remains easy for bounded reservations if bounds are uniform, both in the fractional and integral case. In the latter case we give a very simple algorithm which outputs a bounded acyclic flow-reservation, which is a convex combination of two diverse paths reservations and whose cost is at most twice that of *any* bounded T-resilient vector. Then, in Section 4, we move to the scenario where the bounds λ_a are different. We start by giving an example which shows that in this case the ratio between the cost of any bounded resilient flow reservation (even non-acyclic) and the cost of an optimal bounded T-resilient reservation can be arbitrarily large. Moreover, the problem of finding a minimum cost bounded diverse paths reservation is still NP-hard and not k-approximable if the set of paths is not fixed. We therefore deal with the case where the set of paths is fixed and we give an $O(m \log m)$ combinatorial algorithm for both fractional and integral cases.

We close by considering, in Section 5, the case where, for each arc a of the network, capacity has to be reserved in bundles of μ_a units, even if there is no upper bound on the *total* capacity we may reserve on each arc. It is clear that the problem of finding a minimum cost T-resilient reservation with homogeneous bundles may be reduced to that of finding a minimum cost integral T-resilient reservation, however when bundles are non-homogeneous we show that the problem is NP-HARD even on a fixed set of paths.

2 Diverse Paths Reservations

2.1 Reservations with no Upper Bound Constraints

In [2] we have studied in detail the case where it is required that our T-resilient reservation vector be a diverse paths reservation, i.e., formed from a collection of diverse (s,t) paths by reserving a capacity along each path. We briefly review some results; for details see [2]. The scenario is that we are allowed to reserve as much capacity as we want on each path. We consider two cases: in the former we assume that the set of paths is pre-specified, in the latter they are not. In the former case, we may as well treat the network as consisting of a set of arcs from s to t, and so the problem may be formulated as follows.

Problem (Fixed Set of Paths without Bounds). Given a demand T and a non-decreasing sequence of costs $c_1, c_2, \ldots, c_m$, find a non-negative real vector $x = (x_1, x_2, \ldots, x_m)$ to minimize $\text{cost}(x) = \sum c_i x_i$ subject to the constraints

$$\sum_{i \neq j} x_i \geq T \quad \text{for any } j. \tag{3}$$

In [2], we give extremely simple algorithms to solve this problem, even in the integral case. In fact, we show that an optimal (fractional) T-resilient reservation is obtained at one of the solutions z^i:

$$z^i: \quad z^i_j = \begin{cases} T/(i-1) & j \le i \\ 0 & j > i, \end{cases} \tag{4}$$

for some i between 2 and m. Moreover, since the cost-sequence $\text{cost}(z^2), \ldots, \text{cost}(z^m)$ is unimodal, we may obtain the integral optimum through a simple rounding of the fractional one.

If the set of diverse paths is not given from the outset, there is a difference between the fractional and the integral case, since the latter turns out to be NP-hard. We show that a minimum cost (fractional) T-resilient reservation is obtained at one of the solutions w^i:

$$w^i = \sum_{j=1}^{i} \frac{T}{i-1} e(P^i_j),$$

for some i between 2 and $\kappa(D)$, where $\mathcal{P}^i = \{P^i_1, \ldots, P^i_i\}$ is a minimum cost collection of i diverse (s,t) paths.

In particular, since unimodality holds again, a minimum cost reservation may be found by generating the collections $\mathcal{P}^i$ from $i = 2$ up to the first value where the cost $\text{cost}(w^i)$ starts to rise. We point out that for generating the collections $\mathcal{P}^2$, $\mathcal{P}^3$, ... one may use the *successive shortest path method* – see [1] – we briefly recall this since it is the basis for algorithms we introduce later.

For any $i \ge 1$, suppose that we are inductively given a collection $\mathcal{P}^i$ of minimum cost i diverse (s,t) paths (for $i = 1$ we need a shortest (s,t) path) and let $E(\mathcal{P}^i)$ denote the set of arcs belonging to a path of the collection $\mathcal{P}^i$. At the i-th iteration, we first find a shortest (s,t) path Q_i in the *residual* network D^i – which is built from D by replacing each arc $a = (u,v) \in E(\mathcal{P}^i)$, with cost c_a, by an "artificial" arc (v,u), not present in D, with cost $-c_a$ – and then set $E(\mathcal{P}^{i+1}) = (E(\mathcal{P}^i) \setminus R) \cup F$ where R and F are respectively the set of artificial and "true" arcs of Q_i. We refer to this as the algorithm SSP.

For the integral case, we give in [2] a polynomial time $\frac{15}{14}$-approximate algorithm, and we show that this bound is the best possible (if P $\neq$ NP).

2.2 Reservations with a Uniform Bound

In the rest of this section we consider the scenario where there is one fixed $\lambda \in \mathbb{Q}^+$ such that the capacity x_a we may reserve on any arc a satisfies $x_a \le \lambda$ – we say that the reservation vector x is *uniformly bounded* by λ. Of course, if we are interested in a minimum cost T-resilient reservation, this bound is meaningful only if $\lambda \le T$.

We discuss here the extension of diverse path reservations to this scenario. We start by examining the case where we are given a fixed collection of diverse paths to use. Finding a minimum cost uniformly bounded diverse paths reservation then amounts to solving the following problem, where again each path is regarded as a single arc from s to t.

Problem (Fixed Set of Paths with Uniform Bounds). Given a demand T and a non-decreasing sequence of costs $c_1, c_2, \ldots, c_m$, find a non-negative real vector $x = (x_1, x_2, \ldots, x_m)$ to minimize $\text{cost}(x) = c_1x_1 + c_2x_2 + \cdots + c_mx_m$ subject to the constraints (3) and to $x_i \leq \lambda$ for each i.

Clearly if $\lambda(m-1) < T$, then there is no feasible solution; in other words we need to use at least $\lceil \frac{T+\lambda}{\lambda} \rceil$ paths. For the remainder of this section, we set $n = \lceil \frac{T+\lambda}{\lambda} \rceil$, and assume that $n \leq m$. For $i = 2, \ldots, m$ let z^i be defined by (4) – note that z^i is feasible only when $i \geq n$ – and let $z^{(T;\lambda)}$ be defined by:

$$z_i^{(T;\lambda)} = \begin{cases} \lambda & i < n \\ T - (n-2)\lambda & i = n \\ 0 & i > n. \end{cases}$$

Notice that, if T is an integer multiple of λ, then n is an integer and $z^{(T;\lambda)} = z^n$. Otherwise $z^{T;\lambda}$ is not one of the z^i, but is a feasible vector for the present problem. The proof of the next lemma is routine, and in any case it comes later as a corollary of Lemma 3.5.

Lemma 2.1. *Given an instance of the* **Fixed Set of Paths with Uniform Bounds** *problem, let* $n = \lceil \frac{T+\lambda}{\lambda} \rceil$ *and suppose that* $n \leq m$. *Let* h *be such that* $\text{cost}(z^h) = \min_{i \in \{2,\ldots,m\}} \text{cost}(z^i)$ *(that is,* z_h *is an optimal solution for the problem without the upper bound constraint). Then an optimal solution is given by* $z^{(T;\lambda)}$ *if* $h < n$ *and by* z^h *if* $h \geq n$.

Lemma 2.1 and algorithm SSP provide then a very simple algorithm to solve the problem for a *fixed* set of paths with uniform bounds. In contrast, we next show that the version of the problem where we are free to choose the set of paths is NP-hard, and indeed is not k-approximable for any k, even in the fractional case, unless P = NP.

Denote by BDP the problem which consists of a network with two specified nodes s and t, together with non-negative integer costs on the arcs, an integer T, and a rational λ. An (optimal) solution for an instance of BDP is a minimum cost uniformly bounded T-resilient diverse paths reservation, with upper bound λ. Note that problem BDP does not require the reservation vector to be integer. Let 1-BDP denote the special case of BDP where T is fixed to 1.

Theorem 2.2. *1-BDP is NP-hard. Furthermore there is no k-approximation algorithm for 1-BDP unless P=NP.*

The proof is similar to one in [2].

Proof. Lemma 2.1 shows that the support of an optimal solution to an instance of BDP consists of a collection of diverse paths $P_1, P_2, \ldots, P_j$, where the arcs of the first $j-1$ paths reserve a common amount r of capacity, and the last path's arcs reserve capacity $T-(j-2)r$, equal to r unless $r = \lambda$.

Let 2DIV-PATHS denote the problem of determining whether a given digraph D, with four distinct nodes s_1, t_1, s_2, t_2, contains a pair of arc-disjoint paths P_1, P_2, where P_i joins s_i and t_i $(i = 1, 2)$. Fortune, Hopcroft and Wyllie [3] show that this problem is NP-complete.

Suppose that we are given an instance of 2DIV-PATHS, that is a digraph D, with four distinct nodes s_1, t_1, s_2, t_2. Construct a digraph from D by adding new nodes s, t as well as arcs (s, t), (s, s_1), (s, s_2), (t_1, t) and (t_2, t) with costs 0, 0, 1, 0 and 1 respectively. All remaining arcs have cost zero. See Figure 2. Moreover, let $T = 1$ and assume that $\frac{1}{2} < \lambda < 1$: this is our instance of 1-BDP.

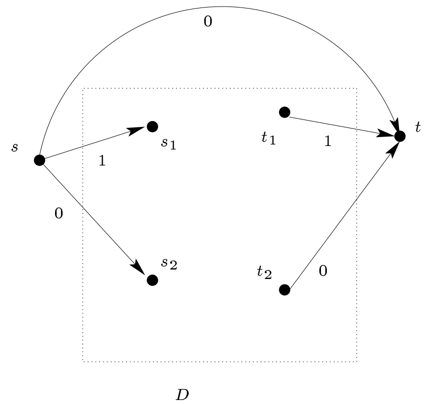

Figure 2: NP-completeness Reduction

We know that an optimal 1-resilient diverse paths reservation will have support consisting of three arc-disjoint paths; two of the paths will have reserved capacity λ and the third capacity $1 - \lambda$.

Note that if there exists a solution to 2DIV-PATHS, with P_i a path between s_i, t_i $(i = 1, 2)$, then by assigning capacity λ to the arcs (s, t), (s, s_1), (t_1, t) and the arcs of P_1, and capacity $1 - \lambda$ to (s, s_2), (t_2, t) and the arcs of P_2 we obtain a solution to 1-BDP of cost $2(1 - \lambda)$. Conversely, if 2DIV-PATHS has no solution, then either 1-BDP itself has no feasible solutions or any solution to 1-BDP uses one path of cost 0 and two paths of cost 1, from which we deduce that the reservation has cost at least 1. Thus the optimal solution to the instance of 1-BDP has cost $2(1-\lambda)$ if and only if the instance of 2DIV-PATHS has a solution, and otherwise the cost is at least 1.

For any k, setting $\lambda = 1 - \frac{1}{2(k+1)}$ implies that, unless P = NP, 1-BDP is not k-approximable. □

The hardness of this problem is extremely sensitive to the values of T and λ given. Indeed, suppose that we are given an instance of BDP where $T = (n-2)\lambda + \mu$, with $0 < \mu \le \lambda$. Then the algorithm SSP can be used to find the optimal solutions of the form $w^{(n)}, w^{(n+1)}, \ldots$. Furthermore, it is easy to see that the cost of any solution of the form $w^{(T;\lambda)}$ (analogous to $z^{(T;\lambda)}$, so having reserved capacity λ on $n-1$ paths, and μ on one) is at least $\mu\frac{n-1}{T}$ times the cost of $w^{(n)}$ (the extreme case being when all the cost of an optimal collection of n paths falls on one path, which is given capacity μ under $w^{(T;\lambda)}$ and $T/(n-1)$ under $w^{(n)}$). Hence a trivial adaptation of algorithm SSP achieves an approximation ratio of $\frac{(n-2)\lambda+\mu}{(n-1)\mu} = 1 + \frac{n-2}{n-1}\frac{\lambda-\mu}{\mu}$. It is fairly easy to adapt the construction given above to show that this is best possible for the given values of T and λ, unless P = NP.

Roughly speaking, the reason for the hardness of BDP is to be found in Lemma 2.1. If $h < n$, we get an "asymmetric" optimal solution, and then it could even be the case that any optimal diverse paths solution is attained at a collection of j (s, t)-paths that is not a minimum cost collection of j diverse (s, t)-paths.

The conclusion is that upper bounds on the capacity make diverse paths reservations hard to handle in general (this, of course, extends to the case when the bounds λ_a are different). In the next section we deal with acyclic flow-reservations, which have many of the nice properties of diverse paths reservations, yet do allow easy optimization in the bounded case.

3 Acyclic Flow-Reservations

In this section we show that seeking an acyclic flow-reservation is a good strategy for building resilient networks when we are given a uniform upper bound on the capacity we may reserve on each arc. We start by discussing acyclic flow-reservations in networks with no upper bound constraints; later we move to the uniformly bounded scenario.

3.1 Acyclic Flows with no Upper Bound Constraints

We recall that a reservation vector x is an *acyclic flow-reservation* if x is an (s, t) flow-vector and there are no cycles in the support $I(x)$ of x. As we pointed out in Section 1, this is, in general, a severe restriction since there may exist non-acyclic flow-reservations which cost less than any acyclic one.

We consider the problem of finding a minimum cost T-resilient acyclic flow, and we start with the problem of finding an optimal one of value M, that is an M-flow.

Lemma 3.1. *Let x be an acyclic flow of value M in a network D. Then x is T-resilient if and only if $x_a \leq M - T$ for every arc a.*

In particular, a minimum cost T-resilient acyclic M-flow is a minimum cost flow of value M in the network with a bound of $\lambda = M - T$ on the capacity of each arc.

Proof. Suppose that x is an acyclic flow of value M. If every arc has reservation $x_a \leq M - T$ then the failure of one arc removes at most $M - T$ units of flow, leaving total flow of at least $M - (M - T) = T$.

Conversely, consider an ordering $s = v_1, v_2, \ldots, v_n = t$ of V such that all arcs in the support of x go from a lower numbered node to a higher. If there is some arc $a = (v_i, v_j)$ with $x_a > M - T$, let $S = \{v_1, \ldots, v_i\}$. If a fails, this removes more than $M - T$ units of capacity from the cut induced by S, so the maximum (s, t)-flow in $D - a$ is $M - x_a < T$.

By cancelling zero-cost cycles, we see that there is some M-flow x, of minimum cost subject to $x_a \leq M - T$ for each a, which is acyclic. The second part of the result now follows from the first. □

We also see that there is *some* T-resilient acyclic M-flow, provided that the (s,t)-connectivity $\kappa(D)$ is at least $\lceil \frac{M}{M-T} \rceil$, or equivalently $M \geq \frac{\kappa(D)}{\kappa(D)-1}T$. From now on in this section, we set $n = \lceil \frac{M}{M-T} \rceil$, which again should be thought of as the minimum number of paths in a T-resilient reservation, and we assume that $n \leq \kappa(D)$.

Lemma 3.1 implies that a minimum cost T-resilient acyclic M-flow can be found efficiently by solving a minimum cost M-flow problem on the network with capacities $M - T$ on each arc. Even better, this problem can be solved by the algorithm SSP; for any i between 2 and $\kappa(D)$, let $\mathcal{P}^i = \{P_1^i, \dots, P_i^i\}$ be a minimum cost collection of i diverse (s,t) paths. Since all the arcs have capacities $M - T$, at the end of the $(n-1)$st iteration the algorithm SSP has a current solution $\bar{w}$ sending $M - T$ units of flow on each path of the family $\mathcal{P}^{n-1}$. Finally, in the nth iteration the algorithm finds an (s,t) path Q_{n-1}, that is of minimum cost in the residual network, and updates $\bar{w}$ to the final value $w^{(M)}$ by "adding" $M - (n-1)(M-T)$ units of flow on Q_{n-1}.

It is routine to check that the solution $w^{(M)}$ may be expressed as follows. Recall that, by definition, $\frac{n}{n-1}T \leq M < \frac{n-1}{n-2}T$ and, for any i between 2 and $\kappa(D)$, diverse path reservation vectors w^i are defined as:

$$w^i = \sum_{j=1}^{i} \frac{T}{i-1} e(P_j^i). \tag{5}$$

Then the solution $w^{(M)}$ is given by:

$$w^{(M)} = \begin{cases} w^2 + (M-2T)e(P_1) & \text{if } n = 2 \\ \alpha w^{n-1} + (1-\alpha) w^n & \text{if } n \geq 3 \end{cases} \tag{6}$$

where α, $0 \leq \alpha < 1$, is such that $M = \alpha \frac{n-1}{n-2}T + (1-\alpha)\frac{n}{n-1}T$. Therefore we have proved the following.

Theorem 3.2. *Let D be a network equipped with non-negative costs on its arcs, and let M and T be such that $n = \lceil \frac{M}{M-T} \rceil \leq \kappa(D)$. Then $w^{(M)}$ is a minimum cost T-resilient acyclic M-flow.*

In other words, if $\frac{M}{M-T}$ is integral, there is a minimum cost M-flow which is a diverse paths reservation. Else, there exists a minimum cost M-flow which is a convex combination either of a pair of diverse paths reservations (when $M < 2T$), or of one such reservation and a reservation using a single shortest path (when $M > 2T$).

It follows that the minimum of the cost $w^{(M)}$ over *all* possible values of M is attained at some w^i (observe that $w^{(M)}$ dominates w^2 for $M > 2T$). Hence we have the following corollary.

Corollary 3.3. *Let D be a network equipped with non-negative costs on its arcs and let T be an integer. There is some minimum cost T-resilient acyclic flow-reservation that is a diverse paths reservation, and this reservation can be found using the successive shortest paths algorithm.*

As we have mentioned, the problem of finding a minimum cost *integral* T-resilient diverse paths reservation (or flow-reservation) is NP-hard. In contrast, we show that finding a minimum cost integral T-resilient acyclic flow-reservation is quite easy.

Let $c(M) = \text{cost}(w^{(M)})$, that is $c(M)$ is equal to the cost of a minimum cost T-resilient acyclic M-flow – as we have seen, $c(M)$ is also the cost of a minimum cost M-flow in the network where each arc has capacity $M - T$.

Lemma 3.4. *The function $c(M)$ is convex in the interval $[\frac{\kappa(D)}{\kappa(D)-1}T, \infty)$ where it is defined.*

Proof. Let $M_1 < M_2 < M_3$ be any three points in the interval $[\frac{\kappa(D)}{\kappa(D)-1}T, \infty)$, and take $\alpha \in (0, 1)$ such that $M_2 = \alpha M_1 + (1 - \alpha)M_3$.

Consider the flow $w = \alpha w^{(M_1)} + (1-\alpha)w^{(M_3)}$: this has value M_2, and each arc carries flow at most $\alpha(M_1-T)+(1-\alpha)(M_3-T) = M_2-T$, so that $c(M_2) \leq \text{cost}(w) = \alpha c(M_1)+(1-\alpha)c(M_3)$, as required. □

Now, observe that, when T is an integer and the allocations must be integers, M is necessarily an integer too and the optimal vectors $w^{(M)}$ are integral. So one of these is optimal for the integral problem. So to find the minimum cost integral T-resilient acyclic M-flow, we need only to minimize $c(M)$ over integers M. By Lemma 3.4, we know that, if the fractional optimum is a vector $w^h = w^{(M^*)}$, where $M^* = \frac{hT}{h-1}$, then the optimal integer solution is either $w^{(\lfloor M^* \rfloor)}$ or $w^{(\lfloor M^* \rfloor)}$. Running the successive shortest paths algorithm SSP for the fractional problem thus provides all the information needed to solve the integral problem as well.

3.2 Reservations with Uniform Bounds

We now go back to the scenario where the capacity we may reserve on each arc is uniformly bounded by $\lambda \in \mathbb{Q}^+$.

As we discussed in Section 1, a minimum cost (uniformly) bounded T-resilient reservation can be found in polynomial time by linear programming while, if we add integrality, the problem is NP-complete. On the other hand, in Section 2.2 we proved that, if we restrict to diverse paths reservations, the problem is not k-approximable for any $k \geq 1$, even in the fractional case. We now show that, if we go instead for acyclic flow-reservations, finding a minimum cost uniformly bounded T-resilient reservation is easy, even in the integral case, and that there always exists such a reservation whose cost is at most twice the cost of *any* uniformly bounded T-resilient reservation.

Indeed, Lemma 3.1 essentially reduces the problem of finding a minimum cost T-resilient acyclic flow, uniformly bounded by λ, to the problem of finding a minimum cost T-resilient acyclic flow of some value at most $T+\lambda$.

This latter problem is easy. Let h be such that $\text{cost}(w^h) = \min_{i\in\{2,\dots,\kappa(D)\}} \text{cost}(w^i)$, where the w^i are defined by (5), that is, w^h is an optimal diverse paths reservation in the absence of the upper bound constraints. Note that w^h is feasible for the bounded problem if $T \leq \lambda(h-1)$. The next lemma follows easily from Theorem 3.2 and Lemma 3.4.

Lemma 3.5. *If c is a vector of non-negative costs, a minimum cost T-resilient acyclic flow of value at most $T+\lambda$ is given by*

$$w^{(T,\lambda,-)} = \begin{cases} w^{(T+\lambda)} & \text{if } T \geq \lambda(h-1) \\ w^h & \text{if } T \leq \lambda(h-1). \end{cases}$$

We point out that the integral version of the previous lemma follows as before. Indeed, if $T \geq \lambda(h-1)$ then $w^{(T,\lambda,-)} = w^{(T+\lambda)}$, which is integral, and otherwise the flow value of the optimal integral solution is either $\lfloor \frac{hT}{h-1} \rfloor$ or $\lceil \frac{hT}{h-1} \rceil$; in both cases by Theorem 3.2 we are done.

In any case, the optimal solution $w^{(T,\lambda,-)}$ may easily be found by means of the successive shortest paths algorithm. Also it is either a diverse paths reservation or a convex combination of two T-resilient diverse paths reservations.

Finally, as the next proposition shows, a simple 2-approximate algorithm for uniformly bounded T-resilience in a network D – in the fractional or integral case – is to find and output a minimum cost T-resilient acyclic $(T+\lambda)$-flow.

Proposition 3.6. *If x is a T-resilient vector bounded by λ, then* $\text{cost}(x) \geq \frac{1}{2}\text{cost}(w^{(T+\lambda)})$.

Proof. As usual, we may assume that $\lambda \leq T$. Let x be a minimum cost T-resilient vector bounded by λ. Define x' such that, for each arc a, $x'_a = \min(2x_a, \lambda)$; it follows that $x' \geq x$.

Claim. For each (s,t)-set S, $x'(\delta^+(S)) \geq T+\lambda$.

Let S be any (s,t)-set. Suppose that there exists $e \in \delta^+(S)$ such that $x_e \geq \frac{\lambda}{2}$, it follows that $x'_e = \lambda$. Then $x'(\delta^+(S)) = x'_e + \sum_{a\in\delta^+(S)-e} x'_a \geq \lambda + \sum_{a\in\delta^+(S)-e} x_a \geq \lambda + T$.

If, on the other hand, $x_e < \frac{\lambda}{2}$ for each $e \in \delta^+(S)$, then $x'(\delta^+(S)) = 2x(\delta^+(S)) > 2T > T+\lambda$. *(End of Claim.)*

Since $x'(\delta^+(S)) \geq T+\lambda$ for each (s,t)-set S, then there exists a flow-vector x'' of value $T+\lambda$ such that $x'' \leq x'$; moreover, since $x''_a \leq x'_a \leq \lambda$, x'' is a T-resilient vector bounded by λ. It follows that $\text{cost}(w^{(T+\lambda)}) \leq \text{cost}(x'') \leq \text{cost}(x') \leq 2\text{cost}(x)$. □

We close this section with some remarks. First, it is easy to check that, if λ is an integer, then $w^{(T+\lambda)}$ is integral too. Moreover, if $\lambda = T$, then $w^{(T+\lambda)} = w^{(2T)}$ and we have the analogous result of [2] for the case with no upper bound constraints. There we prove that the cost of the solution $w^{(2T)}$ is at most twice the cost of any T-resilient solution.

In fact, we may show that the constant $\frac{1}{2}$ in Proposition 3.6 is best possible using the following example from [2]. The network D has three nodes s, u and t, and $\lambda = T$. There are two arcs of cost 1 from s to u, and r arcs of large cost c from u to t. The minimum cost 1-resilient reservation involves giving capacity 1 to the arcs between s and u, and capacity $\frac{1}{r-1}$ to the arcs between u and t, at a cost of $2 + cr/(r-1)$. However, the minimum cost diverse paths reservation uses only two of the arcs from u to t, at a total cost of $2 + 2c$. For large c and r, the ratio between these two costs can be made arbitrarily close to 2: the optimal diverse paths reservation is almost twice as expensive as the optimal 1-resilient reservation.

On the other hand, one would expect that the normal situation is even better than suggested by Proposition 3.6; in fact, we may output the vector $w^{(T,\lambda,-)}$ (see Lemma 3.5), that is a minimum cost T-resilient acyclic flow with value at most $T+\lambda$, whose cost is typically lower than that of $w^{(T+\lambda)}$.

It follows that in many practical settings $w^{(T,\lambda,-)}$ is not very far from the overall optimal solution in cost, and we recall that this vector is either a diverse paths reservation or a convex combination of two diverse paths reservations. So, given its simplicity, it is likely that this is a good solution to adopt.

4 Resilient Reservation with Differing Upper Bounds

We now consider the scenario where there are bounds on the capacity we may reserve on each arc of the network which are, in general, not the same for each arc. Hence any (resilient) vectors must satisfy the additional constraint:

$$x_a \leq \lambda_a \quad \text{for each } a \in A. \tag{7}$$

We call a vector satisfying constraints (7) *bounded.*

As the following example shows, it is hard to approximate a minimum cost *non-uniformly* bounded reservation by a (general) flow reservation. Consider a network D with three nodes s, u and t. There are two arcs with upper bound $u_a = T$ and cost 1 from s to u, three arcs with upper bound $u_a = T/2$ and cost 1 from u to t and one more arc with upper bound $u_a = T/2$ and large cost c from u to t. Then the minimum cost bounded T-resilient reservation involves giving maximum capacity to all the arcs of cost 1, at a total cost of $7T/2$. However, the only T-resilient flow-reservation uses all arcs to full capacity, at a total cost of $(7+c)T/2$. By increasing c, the ratio between these two costs can be made arbitrarily large.

Hence, when the bounds are different, any strategy that finds a flow-reservation (which includes an acyclic flow-reservation or a diverse paths reservation) cannot be guaranteed to be cost-efficient.

Moreover, optimization problems seem to be challenging when the bounds are different. First, we know from Section 2 that the problem of finding a minimum cost bounded T-resilient diverse paths reservation is hard. Also the reduction we used to deal with uniformly bounded acyclic flow reservations (based on Lemma 3.1) does not extend to this case, even in the particular case of diverse paths reservation on a fixed set of paths. On the positive side, linear programming is a method for finding in polynomial time a minimum cost (fractional) bounded T-resilient *general* flow reservation.

The problem of finding a strategy for general bounded resilience with some or all of the good properties that we have been dealing with in other scenarios (simplicity of the solution type sought, good approximation properties, good optimization properties) therefore remains open. However, we are able to treat the case where we require a diverse paths reservation with the paths pre-specified.

Hence, in the rest of this section we suppose that we are given a set of paths $P_1, \dots, P_m$, considered as single arcs from s to t, and for each arc P_i an upper bound λ_i on the capacity we may reserve (in general, λ_i is equal to the smallest upper bound among those of the arcs of a path P_i). The problem of finding a minimum cost resilient reservation is as follows.

Problem (Fixed Bounded Problem). Given a demand T, a vector $\lambda \in \mathbb{Q}_+^m$ and a non-decreasing sequence of costs $c_1, c_2, \dots, c_m$, find a non-negative real vector $x = (x_1, x_2, \dots, x_m)$ to minimize $\mathrm{cost}(x) = c_1x_1 + c_2x_2 + \dots + c_mx_m$ subject to the partial cut constraints (3) and to $x_i \le \lambda_i$ for each i.

First, notice that the set of bounded T-resilient vectors is non-empty if and only if $\sum_j \lambda_j \ge T + \max_j\{\lambda_j\}$. We assume from now on that this holds.

As before, to find a minimum cost bounded T-resilient vector, we start with the problem of finding a minimum cost T-resilient bounded M-flow (that is, a T-resilient vector x with $x_i \le \lambda_i$ for each i, and value $\sum_i x_i = M$). This is equivalent to finding a minimum cost M-flow satisfying $x_i \le \min(\lambda_i, M - T)$ for each i. Observe that, if $\sum_{i=1}^m \min\{M - T, \lambda_i\} < M$ then no T-resilient vector has value M. Otherwise, let n be the smallest index such that $\sum_{i=1}^n \min\{M - T, \lambda_i\} \ge M$. The following result is immediate.

Lemma 4.1. *The vector $x^{(M)}$ is a minimum cost T-resilient M-flow, where:*

$$x_i^{(M)} = \begin{cases} \min\{M - T, \lambda_i\} & i < n \\ M - \sum_{i=1}^{n-1} x_i^{(M)} & i = n. \end{cases}$$

This gives the structure of a bounded minimum cost T-resilient M-flow $x^{(M)}$; from now on let $c(M) = \mathrm{cost}(x^{(M)})$. Let M^* be the minimum rational such that $\sum_{i=1}^m \min\{M^* - T, \lambda_i\} \ge M^*$, so that M^* is the minimum value of M for which there exists a bounded T-resilient M-flow.

The proof of the following lemma is almost the same as that of Lemma 3.4, and we skip it.

Lemma 4.2. *The function $c(M)$ is convex on the interval $[M^*, \infty)$.*

Now let j be the smallest index such that $\sum_{i=1..j} \lambda_i \ge T + \max_{i=1..j} \lambda_i$. Let $\Delta = \max_{i=1..j-1} \lambda_i$. Then, from Lemma 4.1, the vector $x^{(T+\Delta)}$:

$$x_i^{(T+\Delta)} = \begin{cases} \lambda_i & i < j \\ T + \Delta - \sum_{i=1}^{j-1} \lambda_i & i = j. \end{cases}$$

is a minimum cost T-resilient vector with value $T + \Delta$. Also for any $M > T + \Delta$, then $c(M) \ge c(T + \Delta)$. In fact, in this case $x_i^{(M)} \ge x_i^{(T+\Delta)}$ for any $1 \le i \le j$.

Therefore, we now know that to find a minimum cost T-resilient vector amounts to minimizing the convex function $c(M)$ over the interval $[M^*, T+\Delta]$.

As usual, $c(M)$ is piecewise linear and, as we show in the following, it has at most $2m$ "turnpoints" in the interval $[M^*, T+\Delta]$, so the problem can be solved by generating the successive turnpoints and evaluating $c(M)$ until the function starts to rise.

We scan the interval from right to left, i.e., we decrease M. The first turnpoint is $M = T+\Delta$; we show how to generate the next turnpoints.

We denote by $\mathcal{R}(T+\Delta)$ the subset of indices i for which $x_i^{(T+\Delta)} = \Delta$, by definition $|\mathcal{R}(T+\Delta)| \geq 1$. Assume without loss of generality that $0 < x_j^{(T+\Delta)} < \min(\lambda_j, \Delta)$ (the index j has been defined above).

Then, it follows from Lemma 4.1 that, for ε small enough but positive, the optimal solution $x^{T+\Delta-\varepsilon}$ is defined as follows:

$$x_i^{(T+\Delta-\varepsilon)} = \begin{cases} x_i^{(T+\Delta)} - \varepsilon = \Delta - \varepsilon & i \in \mathcal{R}(T+\Delta) \\ x_i^{(T+\Delta)} = \lambda_i & i \notin \mathcal{R}(T+\Delta), i < j \\ x_i^{(T+\Delta)} + \varepsilon(|\mathcal{R}(T+\Delta)| - 1) & i = j \\ x_i^{(T+\Delta)} = 0 & i > j. \end{cases}$$

Therefore, we keep on increasing ε (reducing M) until it reaches a value ε^* for which:

(i) either $\Delta - \varepsilon^* = \lambda_i$ for some $i < j$ not in the set $\mathcal{R}(T+\Delta)$

(ii) or $x_j^{(T+\Delta-\varepsilon^*)} = \min(\Delta - \varepsilon^*, \lambda_j)$.

The second turnpoint is therefore $T+\Delta-\varepsilon^*$. We set $\mathcal{R}(T+\Delta-\varepsilon^*) = \{i : x_i^{(T+\Delta-\varepsilon^*)} = \Delta-\varepsilon^*\}$ and, if $x_j^{(T+\Delta-\varepsilon^*)} = \min(\Delta - \varepsilon^*, \lambda_j)$, then $j = j+1$. Then we iterate.

Observe that, if we scan the interval $[M^*, T+\Delta]$ from right to left, at each turnpoint either there is an index entering the set $r(M)$ or an index entering the support $I(M)$ of $x^{(M)}$. Also if $M_1 > M_2$, then $r(M_2) \subseteq r(M_1)$ and $I(M_2) \subseteq I(M_1)$. Hence there are at most $2m$ turnpoints. Finally, since $c(M)$ changes linearly between turnpoints, the problem can be solved by generating them as explained above and evaluating $c(M)$ until the function starts to rise (because of convexity). It is easily checked then this approach results in an $O(m \log m)$ algorithm for the problem.

This extends to the integral case, since from convexity, the usual rounding arguments apply. We do not go into details.

5 Bundled resilience

In this section we adopt the viewpoint that on each arc a of the network we are free to reserve as much capacity as we want but only in bundles of μ_a units, and each bundle comes with cost c_a. As usual, our objective is to minimize the total cost of the capacity reservations we make, subject to supporting a given target amount T of traffic from s to t, even if one of the arcs in the network fails.

Here there is a dramatic difference between the case (*homogeneous bundles*) where the size of the bundles is the same for all the arcs in the network, that is $\mu_a = \mu$ for each arc a, and the case (*non-homogeneous bundles*) where there are different sizes. In fact, as we show, in the former case any resilience problem reduces to an integral resilience problem (and hence it is easy for diverse paths reservation and flow-reservation) while in the latter even the problem of finding a minimum cost resilient reservation on a fixed set of paths is NP-complete.

HOMOGENEOUS BUNDLES

In this case a resilient reservation is a vector $x \in \mathbb{Z}_+^m$ such that for any $s-t$ cut S and any $e \in \delta^+(S)$ then $\sum_{f \in \delta^+(S) \setminus e} \mu x_f \geq T$. It follows that, without loss of generality, we may assume that $T = h\mu$, where h is a positive integer. Then, trivially, any T-resilient reservation problem with homogeneous bundles reduces to the corresponding integer T'-resilient reservation problem (i.e., without bundles) with $T' = \frac{T}{\mu}$. If there are upper bounds on the capacities, these can be transformed in the same way.

NON-HOMOGENEOUS BUNDLES

Generally speaking, resilience problems with bundles of different sizes are hard to solve. In fact, finding a minimum cost reservation is hard even if we are looking for a diverse paths reservation on a *given* set of paths. Denote this problem by BFP.

Any instance of BFP consists of integers m and T, and integer m-vectors c and μ. The problem may be formulated as follows.

$$
\begin{array}{ll}
\min\ c_1x_1 + c_2x_2 + \cdots + c_mx_m & \\
\sum_{i \neq j}\ \mu_i x_i \geq T & \text{for any } j \\
x_i \geq 0 & \text{an integer for any } i.
\end{array}
$$

The previous program is very similar to the maximization version of the INTEGER KNAPSACK problem. An instance of this problem (which has been proved to be NP-complete both in the minimization and maximization versions [4]) consists of integers p and B, and positive size-p integral vectors w and a. The problem may be formulated as follows:

$$
\begin{array}{ll}
\min\ w_1y_1 + w_2y_2 + \cdots + w_py_p & \\
\sum_i a_i y_i \geq B & \\
y_i \geq 0 & \text{an integer for any } i,
\end{array}
$$

where w and a are vectors of integers and B is an integer too.

It is routine to show that any instance of INTEGER KNAPSACK, may be reduced to BFP in polynomial time: namely set $T = B$; $m = p + 1$; $c_i = w_i$ for $i \leq p$ and $c_{p+1} = 0$; $\mu_i = a_i$ for $i \leq p$ and $\mu_{p+1} = B$. Then if x is a minimum cost B-resilient vector, and y is such that $y_i = x_i$ for $i \leq p$, then y is an optimal solution to INTEGER KNAPSACK.

6 Conclusions

We have shown that, in the presence of a uniform upper bound on all arc capacities, a sensible strategy for the resilience problem is to look for an optimal solution that is an acyclic flow-reservation. This is a slightly more flexible approach than insisting on a diverse-paths reservation, which is crucial for the problems we have considered here.

We do not know of any similarly effective approach in the case where the upper bounds can be different. This seems to us to be an important problem: in a practical setting with more than one commodity, if for instance the flows are routed one commodity at a time, then upper bounds on the arc-capacities are the critical constraints.

The problem underlying most of this research is the lack of a practical "combinatorial" algorithm to solve general resilience problems exactly, even without upper bound constraints.

7 Acknowledgements

Much of this work was carried out while the second author held a Postdoctoral Fellowship at the Centre for Discrete and Applicable Mathematics, LSE, funded by the European Union as part of the TMR programme on Discrete Optimisation (DONET), FMRX-CT98-0202. The authors gratefully acknowledge the support of DONET for this research. The authors are also indebted to the anonymous referees, who read the paper carefully and made a number of valuable suggestions, including one that substantially simplified our algorithm to solve the Fixed Bounded Problem in Section 4.

References

[1] A.K. Ahuja, T.L. Magnanti, J.B. Orlin, *Network Flows – Theory, Algorithms, and Applications*, Prentice Hall, Englewood Cliffs, New Jersey, 1993.

[2] G. Brightwell, G. Oriolo, B. Shepherd, Reserving resilient capacity in a network, *SIAM J. on Disc. Math.* **14** (2001) 524–539.

[3] S. Fortune, J. Hopcroft, J. Wyllie, The directed subgraph homeomorphism problem, *Theoretical Computer Science* **10** (1980) 111-121.

[4] G.S. Lueker, Two NP-complete problems in non negative integer programming, *Report No. 178, Computer Science Laboratory, Princeton University*, 1975.