# Computing V ertex Connectivity: New Bounds from Old Techniques 

M onika R. Henzinger ${ }^{1}$<br>Compaq Systems Research, 130 Lytton Ave., Palo Alto, California 94301<br>E-mail: monika@pa.dec.com<br>Satish R ao<br>Computer Science Division, Soda Hall, University of California, Berkeley, California 94720-1776<br>E -mail: satishr@ cs.berkeley.edu<br>and<br>H arold N. G abow<br>Department of Computer Science, University of Colorado at Boulder, Boulder, Colorado 80309<br>E-mail: hal@ cs.colorado.edu

R eceived September 16, 1996; revised September 6, 1999

The vertex connectivity $\kappa$ of a graph is the smallest number of vertices whose deletion separates the graph or makes it trivial. We present the fastest known deterministic algorithm for finding the vertex connectivity and a corresponding separator. The time for a digraph having $n$ vertices and $m$ edges is $O\left(\min \left\{\kappa^{3}+\right.\right.$ $n, \kappa n\} m$ ); for an undirected graph the term $m$ can be replaced by $\kappa n$. A randomized algorithm finds $\kappa$ with error probability $1 / 2$ in time $O(\mathrm{~nm})$. If the vertices have nonnegative weights the weighted vertex connectivity is found in time $O\left(\kappa_{1} n m \log \left(n^{2} / m\right)\right)$ where $\kappa_{1} \leq m / n$ is the unweighted vertex connectivity or in expected time $O\left(n m \log \left(n^{2} / m\right)\right)$ with error probability $1 / 2$. The main algorithm combines two previous vertex connectivity algorithms and a generalization of the preflow-push algorithm of H ao and Orlin (1994, J. Algorithms 17, 424-446) that computes edge connectivity. © 2000 A cademic Press

[^0]
## 1. INTRODUCTION

The (vertex) connectivity $\kappa$ of an undirected graph or digraph is the smallest number of vertices whose deletion separates or trivializes the graph. (M ost terms are defined precisely at the end of this section.) This is a central concept of graph theory [15]. We present efficient algorithms for computing connectivity.

M ore precisely we consider the following tasks. To compute the connectivity means to find $\kappa$ and a corresponding separator. To check $k$-connectedness means to verify that $\kappa \geq k$ or find a separator of $<k$ vertices. Lastly if each vertex has a nonnegative weight, $\kappa$ is the smallest weight of a separator or trivializor and we may wish to compute this weighted connectivity. We present algorithms for all these tasks, on undirected graphs and digraphs. Our approach is to simply combine three previous algorithms, two for computing vertex connectivity and one for network flow. We now discuss the efficiency and compare it to previous work. Throughout this paper $n$ and $m$ denote the number of vertices and edges of the given graph, respectively.

First consider (unweighted) digraphs. Even and Tarjan showed how to compute $\kappa$ in time $O\left(\kappa n^{1.5} m\right.$ ) [7]. This was improved to the following best-known bounds: Even [6] showed how to check $k$-connectedness by solving at most $k^{2}+n$ network flow problems, achieving a running time of $O((k+\sqrt{n}) k \sqrt{n} m)$. Galil extended Even's approach to compute $\kappa$ in the same time bound with $k$ replaced by $\kappa$ [8]. We find the connectivity in time $O\left(\min \left\{\kappa^{3}+n, \kappa n\right\} m\right)$ and check $k$-connectedness in the same bound with $\kappa$ replaced by $k$. This bound equals the previous best when $k$ or $\kappa$ is $\Theta(1)$ or $\Theta(\sqrt{n})$ and is superior for all other values. The largest improvement is a factor $\sqrt{n}$ when $k$ or $\kappa$ is $\Theta(n)$. Our algorithm uses $O(m)$ space, as do all other algorithms of this paper. Our high-level algorithm is essentially the same as [7] for small connectivities and [6, 8] for large connectivities.

Next consider undirected graphs. An undirected graph has the same connectivity as the digraph formed by giving each edge both directions. Hence the digraph algorithms apply. Furthermore the technique of Nag amochi and Ibaraki [17] usually allows one to work on a subgraph of $O(\mathrm{kn})$ edges rather than the whole graph. This usually allows undirected connectivity and $k$-connectedness to be computed in the digraph time bound with $m$ replaced by $k n$ or $\kappa n$. This is true for all the above algorithms. For example, we compute undirected connectivity in time $O\left(\min \left\{\kappa^{3}+\right.\right.$ $n, \kappa n\} \kappa n$ ). We also mention a recent algorithm of Henzinger: In an undirected graph with minimum degree $\delta$ it finds $\kappa$ and a corresponding separator if $\kappa \leq \delta / 2$ and verifies $\kappa>\delta / 2$ otherwise, in time $O\left(\min \{\kappa, \sqrt{n}\} n^{2}\right)[12]$.

We turn to randomized algorithms. Becker et al. give a Monte Carlo version of [7] that computes the connectivity of a digraph. For any desired probability $p \leq 1 / n$ it returns the correct answer with probability $1-p$ (i.e., the error probability is $p$ ) and runs in expected time $O(((\log 1 / p) /$ $\left.(\log (n / \kappa))) n^{1.5} m\right)$ [3]. We give an algorithm that computes the connectivity with error probability $1 / 2$ in (worst-case) time $O(n m)$. This is a strict improvement of [3] since for any desired probability $p$ repeating our algorithm $\log 1 / p$ times achieves error probability $p$ in time $O((\log 1 / p) n m)$. For undirected graphs the time to achieve error probability $1 / 2$ is $O\left(\kappa n^{2}\right)$ (replace $\kappa$ by $k$ for checking $k$-connectedness). These results use a fact about our deterministic algorithm for computing $\kappa$ : It runs in time $O\left(((n-\kappa) n)^{2}\right)$. This bound comes into play when $\kappa=n-$ $o(n)$ (e.g., the time to compute $\kappa$ is linear when $\kappa=n-O(1)$ ).

Linial et al. check undirected $k$-connectedness with a Monte Carlo algorithm that runs in time $O((M(n)+n M(k)) \log n)$ and has error probability $1 / n$ [16]. Here $M(n)$ is the time to multiply two $n \times n$ matrices and is $O\left(n^{2.38}\right)$ [5]. Our M onte Carlo algorithm (for the same error probability) is faster when $k \leq n^{.37}$ or $k \geq n^{.73}$; it is always faster if naive matrix multiplication is used to get a practical algorithm. Linial et al. also give a Las $V$ egas algorithm with expected running time $k$ times their $M$ onte Carlo bound. Cheriyan and Reif achieve similar Monte Carlo and Las V egas time bounds for digraphs [4]. Our deterministic algorithm is faster than these L as V egas algorithms if they use naive matrix multiplication.

Finally consider vertex-weighted graphs (undirected or digraphs). The naive algorithm computes the weighted connectivity by calculating $\Theta\left(n^{2}\right)$ maximum flows. We know no other results on weighted connectivity. We improve this to time $O\left(\kappa_{1} n m \log \left(n^{2} / m\right)\right)$ where $\kappa_{1}$ denotes the connectivity if vertex weights are ignored (i.e., every vertex has weight 1 ). Since $\kappa_{1} \leq m / n$ this bound is at most $O\left(m^{2} \log \left(n^{2} / m\right)\right.$ ). This improves the naive bound by a factor greater than $n$ if we use the best-known algorithm of [14] to compute a maximum flow. We also present a Monte Carlo algorithm that computes the connectivity with error probability $1 / 2$ in expected time $O\left(n m \log \left(n^{2} / m\right)\right)$.
The vertex connectivity algorithms we use transform the problem to a number of network flow problems, e.g., $\kappa n$ maximum flow problems for [7]. We show the preflow-push approach to maximum flow can reduce the number of flow problems, e.g., to $\kappa$ flow problems for [7]. The preflow-push approach is due to Goldberg and Tarjan [10]. Gallo et al. show how this approach often allows $n$ or more related maximum flow computations to be combined into one [11]. Hao and Orlin [13] use their idea to give an efficient algorithm to find the edge connectivity of a (capacitated) digraph in time $O\left(n m \log \left(n^{2} / m\right)\right.$ ). We present a flow algorithm that extends [13].

In fact [13] sketches an extension but omits some details. To clarify this we first define our flow problem. The complete definition of the VC split problem is given at the end of this section.

Consider a digraph $G=(V, E)$ with a set of terminals $T \subseteq V$. A minimum split for $T$ is a set $S \subseteq V$ that minimizes the total capacity of the entering edges subject to the constraint that both $T \cap S$ and $T-S$ are nonempty. The edge connectivity algorithm of [13] finds a minimum unrestricted cut, or in our terminology, a minimum split for $V$. We extend this algorithm to find a minimum split for $T$ when $T$ is a given vertex cover of $G$; we call this a minimum VC split. We find a minimum VC split in the same time bound as the edge connectivity algorithm of [13] (e.g., for the above unweighted connectivity algorithms this bound is $O(\mathrm{~nm})$ ).

In hindsight we discovered that H ao and Orlin sketch a similar extension of their algorithm to find a minimum cut in a bipartite digraph ([13], Theorem 9, p. 443; vertex connectivity is not discussed). They indicate that using the bipush variant of the preflow-push algorithm [2] correctly implements the extended algorithm. We do not use bipushes for our main VC split algorithm, but we do introduce a new operation called join. O ur join operation can decrease distance labels by large amounts; previous preflow-push algorithms never decrease distance labels and this is crucial for their efficiency [1, 2, 10, 11, 13, 14]. We prove that the desired time bounds hold in spite of this decrease. It is possible to implement joins slightly less efficiently so they do not decrease distances, but then they increase distances by large amounts. It seems that the phenomenon handled by joins must be addressed and analyzed in any complete algorithm for this problem. Thus, our contribution is to fill a gap in the fundamentally sound and insightful suggestion of H ao and Orlin .
We anticipate further applications of the VC split algorithm besides vertex connectivity. For instance taking the vertex cover $T=V$ shows our algorithm generalizes [13]; i.e., it can compute the edge connectivity in the same bound as [13]. The vertex cover $T=V-\{s\}$ implements Frank's digraph connectivity augmentation algorithm in the same time bound as [9].

The paper is organized as follows. Section 2 gives algorithms for vertex connectivity, assuming an efficient algorithm for finding a minimum VC split. Section 3 supplies the VC split algorithm. The rest of this section gives notation and definitions.
$\mathbf{R}_{+}$denotes the set of nonnegative real numbers. We often denote singleton sets by omitting set braces, e.g., $U \cup u$. If $U, W \subseteq V$ then a $\bar{U} W$-set is a subset of $V$ containing $W$ and disjoint from $U$. For example for $u, w \in V$, a $\bar{u} w$-set contains $w$ but not $u$. If $f$ is a function $f: S \rightarrow \mathbf{R}$ then for any set $T \subseteq S, f(T)$ denotes $\Sigma\{f(t): t \in T\}$.

In a graph with vertices $v$ and $w$, the notation $v w$ denotes an undirected edge joining $v$ and $w$ or a directed edge from $v$ to $w$; it will be clear from context which is meant. For a digraph $G, G^{R}$ denotes the reverse digraph, i.e., all edges of $G$ are reversed. The undirected version of $G$ ignores the directions of edges.
A (directed) edge $u v$ enters any $\bar{u} v$-set and leaves any $\bar{v} u$-set. For a set of vertices $W$, the in-degree $\rho(W)$ equals the total number of edges entering $W$. If the graph has a capacity function on the edges then $\rho(W)$ is the total capacity of edges entering $W$. An xy-cut is the set of edges entering some $\bar{x} y$-set. A minimum $x y$-cut is an $x y$-cut of smallest total capacity; this minimum equals the value of a maximum $x y$-flow.

Consider a digraph or an undirected graph $G=(V, E)$. For $x, y \in V$, $S \subseteq V-\{x, y\}$ is an $x, y$-separator if $G-S$ contains no path from $x$ to $y$. Define $\kappa(x, y)$ as the smallest cardinality of an $x, y$-separator if such separators exist; if not (i.e., $x=y$ or $x y \in E$ ) then set $\kappa(x, y)=n-1$. D efine the vertex connectivity by

$$
\kappa=\min \{\kappa(x, y): x, y \in V\} .
$$

A vertex-weighted graph has a function $w: V \rightarrow \mathbf{R}_{+}$. In this context define $\kappa(x, y)=\min \{w(S), \infty: S$ an $x, y$-separator $\}$. The weighted vertex connectivity $\kappa$ is again defined by the above equation. For $w$ identically 1 this is just ordinary vertex connectivity.

For an undirected graph $G=(V, E)$ the algorithm of [17] partitions $E$, in $O(m)$ time, into a sequence of forests $F_{k}, k=1, \ldots, n$. For $k=1, \ldots, n$ define the forest subgraph $F G_{k}$ by

$$
F G_{k}=\left(V, \bigcup_{i=1}^{k} F_{i}\right) .
$$

For all $k, F G_{k}$ is $k$-connected if $G$ is. In fact [17] proves the stronger property that any vertex separator of $F G_{k}$ with cardinality $<k$ is a vertex separator of $G$ (see the proof of Theorem 3.1 in [17]). Clearly $F G_{k}$ has $O(k n)$ edges.

Consider a digraph $G=(V, E)$ with a set of "terminals" $T \subseteq V$. A subset of $V$ is a split of $T$ if it is an $\bar{s} t$-set for some vertices $s, t \in T$. If each edge has a nonnegative capacity, a minimum split of $T$ is a split $S$ of $T$ that minimizes $\rho(S)$. For any $s \in T$, an $s$-split of $T$ is an $\bar{s} t$-set for some $t \in T$, and we similarly define a minimum $s$-split of $T$. Clearly a minimum split of $T$ is a minimum $s$-split of $T$ in either $G$ or $G^{R}$. The minimum $V C$ split problem is to find a minimum $s$-split of $T$, a given vertex cover of (the undirected version of) $G$.

If a function refers to a graph we include the graph as an extra argument if it is not clear, e.g., $\rho(W, G), \kappa(G)$.

## 2. VERTEX CONNECTIVITY

This section gives our algorithms for vertex connectivity. It summarizes their efficiency, assuming the results of Section 3 for the minimum $V C$ split problem.

Consider a digraph $G=(V, E)$. For any $x \in V$ define

$$
\kappa(x)=\min \{\kappa(x, y): y \in V\} .
$$

Clearly $\kappa=\min \{\kappa(x): x \in V\}$. We begin by showing that computing $\kappa(x)$ reduces to the minimum $V C$ split problem. A ssuming Theorem 3.2 this gives an algorithm to compute $\kappa(x)$ in time $O(n m)$. In this discussion fix the digraph $G$ and the vertex $x$. We use three graphs derived from $G$, illustrated in Fig. 1.

We start with the standard reduction to network flow by node splitting. Specifically define a graph $S G$ (the "split graph") having for each $v \in V$, two vertices $v_{T}, v_{S}$ and edges $v_{T} v_{S}$ and $v_{S} v_{T}$ of capacity 1 and $\infty$, respectively, and for each edge $v w \in E$ an edge $v_{S} w_{T}$ of capacity $\infty$. There is a bijection between $x, y$-separators of $G$ and finite $x_{S} y_{T}$-cuts of $S G$; specifically, separator $U \subseteq V$ corresponds to cut $\left\{u_{T} u_{S}: u \in U\right\}$. Hence, $\kappa(x, y ; G)$ equals the minimum capacity of an $x_{S} y_{T}$-cut in $S G$. Letting $y$ be arbitrary we conclude that $\kappa(x)=\kappa(x, G)$ equals the capacity of a smallest $x_{S} y_{T}$-cut for a vertex $y_{T}$ in $S G$ or $n-1$ if no such finite cut exists. A ny such cut corresponds to an $x, y$-separator of $\kappa(x)$ vertices (using the above bijection).

To find this cut using the algorithm of Section 3 it is convenient to use two other graphs. $C G$ (the "contracted graph") is derived from $S G$ by contracting $x_{S}$ with its neighbors, i.e., contract vertices $x_{S}, x_{T}$, and $z_{T}$ for $z \in V$ with $x z \in E$. Call the vertex resulting from this contraction $X_{T}$. A ssume $X_{T}$ is not the only vertex of $C G$, since otherwise $\kappa(x)=n-1$. O bserve the following:
(i) $\kappa(x)$ equals the capacity of the smallest $X_{T} y_{T}$-cut, for $y_{T}$ a vertex of $C G$.
(ii) A ny such cut corresponds to an $x, y$-separator of $\kappa(x)$ vertices.
(iii) $C=\left\{y_{T}: y_{T}\right.$ a vertex of $\left.C G\right\}$ is a vertex cover of $C G$ (this set includes $X_{T}$ ).

Property (i) holds since $C G$ is constructed by contracting infinite capacity edges incident to the source $x_{S}$ of $S G$. Property (iii) holds since $C$ is one side of a bipartition of $C G$.
Next we transform $C G$ to a unit capacity graph $U G$ : Delete all edges $v_{S} v_{T}$ and assign capacity 1 to all edges $v_{S} w_{T}$. The resulting graph $U G$ satisfies Properties (i) and (iii). For Property (i) the transformation does

not change the value of a maximum flow from $X_{T}$ to any $y_{T}$, so it does not change the value of a minimum cut. (Property (ii) need not hold since new minimum cuts may be introduced.)
Properties (i) and (iii) show that $\kappa(x)$ is the value of a minimum $X_{T}$-split of vertex cover $C$ in digraph $U G$ (recall the definition of minimum $s$-split of $T$ at the end of Section 1). Since $U G$ is a graph with unit edge capacities, Theorem 3.2 shows $\kappa(x)$ can be computed in time $O(n m)$. (Note that $n$ and $m$ for $U G$ are within a constant factor of the same parameters for $G$.) It is a simple matter to derive an $x, y$-separator of $\kappa(x)$ vertices from a minimum $X_{T}$-split: If the set of edges entering the minimum $X_{T}$-split is $\left\{u_{T} u_{S}: u \in U\right\} \cup\left\{v_{S} w_{T}: v w \in F\right\}$ (for some sets $U \subseteq V, F \subseteq E)$ the desired separator can be taken as $U \cup\{v: v w \in F$ for some vertex $w$ \}.

Now we show how our algorithm for $\kappa(x)$ can be used to carry out a number of connectivity computations. The first group of results is summarized as follows. Define three functions

$$
T_{1}(k, n, m)=k n m, T_{2}(k, n, m)=\left(k^{3}+n\right) m, T=\min \left\{T_{1}, T_{2}\right\} .
$$

Note $T=T_{1}$ for $k \geq \sqrt{n}$ and $T=T_{2}$ for $k \leq \sqrt{n}$.
Theorem 2.1. For a digraph, the connectivity $\kappa$ can be found in time $O(T(\kappa, n, m))$ and $k$-connectedness can be checked in time $O(T(k, n, m))$. For an undirected graph the bounds are $O(T(\kappa, n, \kappa n))$ and $O(T(k, n, k n))$, respectively. The space is $O(m)$ in all cases.

Proof of Theorem 2.1 for digraphs and $T=T_{1}$. We implement the algorithm of Even and Tarjan [7] to compute $\kappa$ as follows. Let $x_{i}$, $i=1, \ldots, n$ be an arbitrary indexing of the vertices.

$$
i \leftarrow 0 ; \kappa \leftarrow n-1 ;
$$

repeat $\left\{i \leftarrow i+1 ; \kappa \leftarrow \min \left\{\kappa, \kappa\left(x_{i}, G\right), \kappa\left(x_{i}, G^{R}\right)\right\}\right\}$ until $i>\kappa$;
We refer to this as the Even-Tarjan algorithm.
To see this algorithm is correct assume the connectivity is $<n-1$ and let $S$ be a minimum vertex separator of $G$. When the algorithm halts it has $i>\kappa \geq|S|$. Some $x_{j}$ does not belong to $S$ for $j \leq|S|+1 \leq i$. Thus, $\kappa\left(x_{j}\right)=|S|$ in $G$ or $G^{R}$. Since the algorithm computes $\kappa\left(x_{j}\right)$ it halts with $\kappa=|S|$.

It is a simple matter to find a minimum vertex separator assuming each computation of $\kappa(x)$ finds a corresponding separator. To check $k$-connectedness terminate the loop when $i=k$ and test if $\kappa \geq k$.
The desired time bound with $T=T_{1}$ follows since each value $\kappa(x)$ is computed in time $O(\mathrm{~nm})$.

Proof of Theorem 2.1 for undirected graphs and $T=T_{1}$. As mentioned, the digraph algorithm applies to an undirected graph. (There is no need to consider $G^{R}$.) To check $k$-connectedness we apply the digraph algorithm to the forest subgraph $F G_{k}$ of $G$ (this subgraph, constructed by the algorithm of Nagamochi and Ibaraki [17], is defined in Section 1). Since $F G_{k}$ has $O(k n)$ edges and is found in time $O(m)$ [17] we achieve the desired time bound.

We can also compute the connectivity by working on a subgraph of $O(\kappa n)$ edges. There are two ways to achieve this. The first is to begin by finding a value $k$ with $k \leq \kappa<2 k$ (do this by checking $k$-connectedness for $k$ doubling; alternatively use the algorithm of [12]); then run our connectivity algorithm on the forest subgraph $F G_{2 k}$. We give a second method that may be superior in practice because it is one simple loop. (R eaders only interested in asymptotic results can skip this.)

Let $G=(V, E)$ be the given graph. Index the vertices of $V$ as $x_{i}$, $i=1, \ldots, n$. A lso let $x_{n+i}=x_{i}$ and $F G_{n+i}=F G_{n}=G$ for $i=1, \ldots, n$.

$$
i \leftarrow 0 ; \kappa \leftarrow n-1 ;
$$

## repeat

$$
\left\{i \leftarrow i+1 \text {; if } \kappa\left(x_{i}, F G_{i}\right)<\min \{i, \kappa\} \text { then } \kappa \leftarrow \kappa\left(x_{i}, F G_{i}\right)\right\}
$$

until $i>2 \kappa$;
To prove this algorithm is correct let $\kappa^{*}$ denote the connectivity of $G$. First observe that the iterations with $i \leq \kappa^{*}$ all have $\kappa=n-1$. This follows since $F G_{i}$ is $i$-connected so $\kappa\left(x_{i}, F G_{i}\right) \geq i$. When $i>\kappa^{*}, F G_{i}$ has connectivity $\kappa^{*}$. Thus, $\kappa \geq \kappa^{*}$ and the algorithm does not halt with $i \leq 2 \kappa^{*}$.

Now assume $\kappa^{*}<n-1$ and consider the iterations for $i=\kappa^{*}+1$, $\ldots, 2 \kappa^{*}+1$. Let $S$ be a separating set of $G$ with $\kappa^{*}$ vertices. $S$ is a separating set of each $F G_{i}$. At least one of these $\kappa^{*}+1$ iterations has $x_{i} \notin S$ (note that each iteration examines a distinct $x_{i}$ ). This iteration sets $\kappa=\kappa^{*}$. Furthermore any corresponding vertex separator of $F G_{i}$ is a vertex separator of $G$, since $i>\kappa^{*}$ (see the discussion of the forest subgraph in Section 1). Thus, the algorithm halts with the correct connectivity and separator, and $i=2 \kappa^{*}+1$. The latter bound on the number of iterations implies the desired time bound.

Proof of Theorem 2.1 for checking $k$-connectedness and $T=T_{2}$. The algorithm is an implementation of a variant of Even's algorithm for checking $k$-connectedness [6].

Given a digraph $G$ and parameter $k$, the algorithm either verifies that $\kappa \geq k$ or it finds $\kappa$ and a corresponding separator. Define $\kappa^{-}$as $\min \{\kappa, k\}$ and similarly for $\kappa^{-}(x, y)$. Thus, we seek $\kappa^{-}$and a corresponding separator if $\kappa^{-}<k$.

For $X \subseteq V$ and $y \in V-X, S \subseteq V-\{y\}$ is a weak $X, y$-separator if $G-S$ contains no path from $X$ to $y$. Equivalently a weak $X, y$-separator is a subset of $V-\{y\}$ that is an $x, y$-separator for every $x \in X$ that it does not contain. Note that $X$ is trivially a weak $X, y$-separator. Define $\kappa_{W}(X, y)$ as the smallest cardinality of a weak $X, y$-separator; also $\kappa_{W}(X)$ $=\min \{\kappa(X, y): y \in V-X\}$. If this quantity is less than $|X|$ then it equals the cardinality of an $x, y$-separator for some $x \in X$.

Choose two arbitrary sets $X, Y \subseteq V$ of cardinality $k$. (These sets can be equal. We will need the sets to be distinct when we use this algorithm to compute к.) Observe that

$$
\begin{equation*}
\kappa^{-}=\min \left\{\kappa_{W}(X)\right\} \cup\left\{\kappa_{W}\left(Y, G^{R}\right)\right\} \cup\left\{\kappa^{-}(y, x): y \in Y, x \in X\right\} . \tag{1}
\end{equation*}
$$

In proof, each quantity on the right-hand side of (1) corresponds to a separator if it is less than $k$, so the right-hand side is $\geq \kappa^{-}$. Next suppose $\kappa<k$ and let $S$ be a minimum separator of $\kappa$ vertices. Thus $V-S$ can be partitioned into two nonempty sets $A, B$ such that $G-S$ contains no path from $A$ to $B$. If $X \subseteq A \cup S$ then $\kappa_{W}(X)=|S|$ (we use the fact that $|X|=k$ ). Similarly if $Y \subseteq B \cup S$ then $\kappa_{W}\left(Y, G^{R}\right)=|S|$. In the remaining case $X$ contains a vertex of $B$ and $Y$ contains a vertex of $A$. Thus, $\kappa$ equals the minimum value in the third set of (1).
O ur algorithm works by computing the quantities on the right-hand side of (1) and corresponding separators (if such exist). We now describe how each of these quantities is computed.

Consider the first quantity $\kappa_{W}(X)$. Define a digraph $\tilde{G}$ by starting with $G$ and adding a new vertex $a$ with edges $a x, x \in X$. Fixa vertex $y \in V-X$. A subset of $V-y$ is an $a, y$-separator in ${ }_{\sim} G$ if and only if it is a weak $X, y$-separator in $G$. H ence $\kappa_{W}(X)=\kappa(a, G)$, and we can compute $\kappa_{W}(X)$ using our procedure to compute $\kappa(x)$. Similarly we compute the second quantity of (1), $\kappa_{W}\left(Y, G^{R}\right)$.

Finally we compute each value $\kappa^{-}(y, x), y \in Y, x \in X$. A s noted above, $\kappa(y, x)$ equals the size of a minimum $y_{S} x_{T}$-cut in the split graph $S G$. Thus, we can compute $\kappa^{-}(y, x)$ by computing a maximum flow from $y_{S}$ to $x_{T}$, stopping if we verify that the maximum flow value is $\geq k$.

Now we prove the desired time bound $T=T_{2}$. It suffices to show this for digraphs. The time to compute $\kappa_{W}(X)$ and $\kappa_{W}\left(Y, G^{R}\right)$ is $O(\mathrm{~nm})$ by our discussion of $\kappa(x)$. We compute each value $\kappa^{-}(y, x)$ by the Ford-Fulkerson maximum flow algorithm, computing $\leq k$ augmenting paths. Hence, the time is $O(\mathrm{~km})$. The time to compute all $k^{2}$ of these values is $O\left(k^{3} m\right)$. This establishes the time bound.

Proof of Theorem 2.1 for computing connectivity and $T=T_{2}$. We find the connectivity $\kappa$ using this algorithm: Check $k$-connectedness for $k$
equal to successive powers of 2 until we reach a power $2^{i}>\kappa$. At that point the algorithm determines $\kappa$ and a corresponding separator. (R ecall our checking algorithm finds $\kappa$ and a corresponding separator when $k>\kappa$.)

If the graph is undirected it is easy to see that the total time is $O\left(T_{2}(\kappa, n, \kappa n)\right)$. Now consider a digraph. The desired bound $O\left(T_{2}(\kappa, n, m)\right.$ ) holds except for the computations of $\kappa_{W}(X)$ and $\kappa_{W}\left(Y, G^{R}\right)$. This takes time $O(n m)$ in each $k$-connectedness check, giving a total contribution of $O(n m \log \kappa)$. Now we reduce this to $O(\mathrm{~nm})$. Specifically we show that one execution of our VC split algorithm gives enough information to compute each desired value $\kappa_{W}(X)$ in $O(n)$ time. (Hence, the total time is $O(n m+n \log \kappa)=O(n m)$ ). The values $\kappa_{W}\left(Y, G^{R}\right)$ are computed the same way.
The procedure we have given to compute $\kappa_{W}(X)=\kappa(a, \tilde{G})$ works as follows: It constructs a contracted graph $C G$ by starting with the split graph of $G$ and contracting vertices $a_{S}, a_{T}$, and $x_{T}$ for $x \in X$ into a new vertex $X_{T}$. Then it finds a minimum $X_{T}$-split of the vertex cover $\left\{y_{T}: y_{T}\right.$ a vertex of $U G$ \} in the unit capacity graph $U G$.

Let $S G$ denote the split graph of $G$. Let $U^{\prime} G$ be the unit capacity graph formed from $S G$ in a manner similar to $U G$; i.e., delete all edges $v_{S} v_{T}$ and assign capacity 1 to all edges $v_{S} w_{T}$. (Note that we have not used a contracted graph $C G$ to form $U^{\prime} G$.) Choose an arbitrary $z \in V$. Suppose we execute our algorithm to find a minimum $z_{T}$-split of vertex cover $C=\left\{y_{T}: y \in V\right\}$ of digraph $U^{\prime} G$. We shall see that our algorithm works as follows: It orders the vertices of $C-z_{T}$ as $t_{i}, i=1, \ldots, n-1$. (The ordering is determined by the algorithm.) For each such $i$ it computes $\mu_{i}$ as the value of a minimum $S t_{i}$-cut where $S=\left\{z_{T}, t_{1}, \ldots, t_{i-1}\right\}$.

O ur connectivity algorithm saves these values $\mu_{i}$. Then for each $k$-connectedness check, we choose $X$ to be $\left\{z_{T}, t_{i}: 1 \leq i<k\right\}$. The value $\kappa_{W}(X)$ $=\kappa(a, G)$ equals $\min \left\{\mu_{i}: k, \ldots, n-1\right\}$. This follows because the contracted graph $C G$ in the computation of $\kappa(a, G)$ is the same as the split graph of $G$ with vertices $x_{T}, x \in X$ contracted.

Proof of Theorem 2.1 concluded. It is easy to combine the two approaches to connectivity into one algorithm: We execute the second algorithm to find $\kappa$ or determine that $\kappa>\sqrt{n}$. In the latter case we then switch to the first algorithm. This gives one algorithm that achieves Theorem 2.1.

We mention a related connectivity problem. In a digraph $G=(V, E)$ for any $x \in V$ define

$$
\hat{\kappa}(x)=\min \left\{\kappa(x, G), \kappa\left(x, G^{R}\right)\right\} .
$$

(If $G$ is undirected clearly $\hat{\kappa}=\kappa$.) The quantity $\hat{\kappa}(x)$ indicates how well connected $x$ is to the remaining vertices. It may be of interest to know the entire function $\hat{\kappa}$. For example to choose a "server" in a distributed network we would select a host with the highest connectivity to the "client" nodes; i.e., the server $x$ should maximize $\hat{\kappa}(x)$.
To calculate all values $\hat{\kappa}(x)$ first compute $\kappa$ and a corresponding separator $S$. Observe that any vertex $x \notin S$ has $\hat{\kappa}(x)=\kappa$. Thus, we complete the calculation by computing $\hat{\kappa}(x)$ for each $x \in S$.
Corollary 2.2. In a digraph all values $\hat{\kappa}(x), x \in V$ can be found in time $O$ (кпт).

Note that in an undirected graph we can also achieve the time bound $O\left(\kappa \kappa^{*} n^{2}\right)$ where $\kappa^{*}$ is the largest value of $\hat{\kappa}$. For this do the second step using graphs $F G_{k}$ with $k$ doubling, until all values of $\hat{\kappa}$ have been found.

In preparation for the next result we refine our time bound for very high connectivities ( $\kappa=n-o(n)$ ). Recall our algorithm for computing $\kappa(x)$. The contracted graph $C G$ contains a vertex $y_{T}$ only when $x y \notin E$. To state the main observation, for each vertex $x \in V$ define $d_{o}(x)$ to be the number of edges leaving $x$ (in the given graph $G$ ). Then the number of vertices $y_{T}$ in $C G$ is

$$
\begin{equation*}
n-d_{o}(x) \tag{2}
\end{equation*}
$$

For computing the connectivity note that $d_{o}(x) \geq \kappa$. Thus, (2) implies $C G$ has $\leq n-\kappa$ vertices $y_{T}$ and $O(n(n-\kappa))$ edges. The same bounds hold when $C G$ is constructed from $G^{R}$. Hence, the time bound $O\left(n_{T} m\right)$ of Theorem 3.2 implies that we find $\kappa(x)$ in time $O((n-\kappa) m)=O(n(n-$ $\kappa)^{2}$ ) and thus the Even-Tarjan algorithm given above to compute $\kappa$ runs in time $O\left(((n-\kappa) n)^{2}\right)$ (for digraphs or undirected graphs). For example, this time bound is linear when $\kappa=n-O(1)$. A similar bound holds for checking $k$-connectedness.

We turn to randomized algorithms. We compute the connectivity as follows. Define $\delta$ as the smallest in-degree or out-degree of a vertex. Our randomized algorithm is the E ven-Tarjan algorithm, except that we choose each vertex $x_{i}$ randomly and terminate when $i \geq 1 / \log (n / \delta)$.

To state the results define

$$
T(n, m)=n m .
$$

Theorem 2.3. The following can be done with error probability $\leq 1 / 2$ : For a digraph the connectivity $\kappa$ can be found in time $O(T(n, m))$. For an undirected graph $\kappa$ can be found in time $O(T(n, \kappa n))$ and $k$-connectedness can be checked in time $O(T(n, k n))$. For any fixed nonnegative value $f$, error
probability $\leq 1 / n^{f}$ can be achieved in the same time bounds if $\kappa$ or $k$ is $O\left(n^{e}\right)$ for any fixed value $e<1$. In all cases the space is $O(m)$.

Proof. First consider digraphs. To calculate the error probability of our algorithm let $S$ be a minimum vertex separator. To err the algorithm must choose all $i$ vertices from $S$. This occurs with probability $\leq(|S| / n)^{i} \leq$ $(\delta / n)^{i} \leq 1 / 2$.

To bound the time write $\delta=(1-\epsilon) n$ for $\epsilon$ between 0 and 1 . Using natural logarithms, $\ln (n / \delta)=-\ln (1-\epsilon)>\epsilon$. Thus, since $1 / \epsilon \geq 1$ the number of iterations is $O(1 / \epsilon)$. $\cup$ sing (2) and Theorem 3.2, each iteration computes $\kappa\left(x_{i}\right)$ in time $O((n-\delta) m)=O(\epsilon n m)$. Thus, the total time is $O(n m)$ as desired.

Now consider undirected graphs. We check $k$-connectedness using the digraph algorithm on $F G_{k}$.

To compute $\kappa$, find a value $k \in(\kappa, 2 \kappa]$ and use the digraph algorithm to compute the connectivity of $F G_{k}$. We find the desired $k$ using the algorithm of [12] in time $O\left(\kappa n^{2}\right)$ or less. This implies the desired time bound.

For the last part of the theorem first consider digraphs. We change the termination test to $i \geq f /(1-e)$. It is easy to see this achieves the desired error probability. The time bound follows since each iteration uses time $O(n m)$. U ndirected graphs are similar.

To make this paper self-contained we sketch a procedure that for a given undirected graph computes a value $k \in(\kappa, 4 \kappa)$ in time $O\left(\kappa n^{2}\right)$. This procedure is conceptually simple and can replace that of [12] in the above algorithm to compute undirected connectivity.

O ur algorithm is based on the following fact. Let $x$ be a vertex of degree exactly $k$ in $F G_{k}$. (The existence of such a vertex when $k \leq \kappa$ is proved in [17]. $x$ can be chosen as the last vertex scanned by the algorithm.) Then $\kappa\left(x, F G_{k}\right)$ equals $k$ if $k \leq \kappa$ and is $<k$ if $k \geq 2 \kappa$.
To prove the fact define $N$ as the set of all neighbors of $x$. If $k \leq \kappa$ then $F G_{k}$ is $k$-connected. Thus, $N$ is a separator (or trivializor) that shows $\kappa\left(x, F G_{k}\right)=k$.

Suppose $k \geq 2 \kappa$. Define $S$ as a vertex separator of $\kappa$ vertices. The desired conclusion $\kappa\left(x, F G_{k}\right)<k$ is clear if $x \notin S\left(\kappa\left(x, F G_{k}\right)=\kappa\right)$ or if $F G_{k}$ contains a vertex of degree $<k$. Suppose neither of these is the case. Partition $V-S$ into nonempty sets $A$ and $B$ such that $G-S$ contains no path from $A$ to $B$. Without loss of generality $|A \cap N| \leq k / 2$. Consider the set $T=(S-x) \cup(A \cap N)$. N ote that $A-N$ contains some vertex $a$. (If $A-N=\varnothing$ then $|A| \leq k / 2$. This implies any vertex of $A$ has degree $\leq(k / 2-1)+|S| \leq k-1<k$, contradicting our initial assumption.) Thus, $T$ is an $a, x$-separator. Furthermore $|T| \leq(\kappa-1)+k / 2 \leq k-1$. Thus, $T$ shows $\kappa\left(x, F G_{k}\right)<k$ as desired.

The algorithm to compute $k \in(\kappa, 4 \kappa)$ calculates $\kappa\left(x, F G_{k}\right)$ for $k$ doubling until $\kappa\left(x, F G_{k}\right)<k$. The above fact shows the algorithm halts with $k$ larger than $\kappa$ but no larger than the power of 2 in $[2 \kappa, 4 \kappa)$. The time is $O\left(\kappa n^{2}\right)$ since each calculation of $\kappa(x)$ uses time $O\left(k n^{2}\right)$.

The final topic of this section is computing weighted vertex connectivity. Let $w: V \rightarrow \mathbf{R}_{+}$be a vertex weight function. Our algorithm first finds a minimum unweighted separator $S$ (using one of the previous algorithms). Then it sets

$$
\kappa=\min \left\{w(S), \kappa(x, G), \kappa\left(x, G^{R}\right): x \in S\right\} .
$$

(All references to $\kappa$ refer to weighted connectivity.) This is correct since if $S$ is not a minimum weight separator then some $x \in S$ is not contained in some minimum weight separator, so the weighted connectivity is $\kappa(x, G)$ or $\kappa\left(x, G^{R}\right)$.
To calculate the values $\kappa(x, G)$ and $\kappa\left(x, G^{R}\right)$ we proceed similar to the unweighted case: D efine the split graph $S G$ as before, except that each edge $v_{T} v_{S}$ has capacity $w(v)$. D efine the contracted graph $C G$ as before. A pplying Theorem 3.3 to $C G$ shows $\kappa(x, G)$ can be found in time $O\left(n m \log \left(n^{2} / m\right)\right)$.

Theorem 2.4. For any digraph or undirected graph with nonnegative vertex weights, the weighted connectivity can be found in time $O\left(\kappa_{1} n m \log \left(n^{2} / m\right)\right)$, where $\kappa_{1}$ is the unweighted vertex connectivity. The space is $O(m)$.

A randomized algorithm also works. Define

$$
\delta=\min \{w(\{x: x v \in E\}), w(\{x: v x \in E\}): v \in V\} .
$$

We use the Even-Tarjan algorithm, except that we choose each vertex $x_{i}$ randomly with any $v \in V$ having probability $w(v) / w(V)$, and we terminate when $i \geq 1 / \log (w(V) / \delta)$.
Theorem 2.5. For any digraph or undirected graph with nonnegative vertex weights, the weighted connectivity can be found with error probability $\leq 1 / 2$ in expected time $O\left(n m \log \left(n^{2} / m\right)\right)$. The space is $O(m)$.

Proof. To calculate the error probability let $S$ be a minimum weight vertex separator. To err the algorithm must choose all $i$ vertices from $S$. This occurs with probability $\leq(\kappa / w(V))^{i}$. Note that $\delta \geq \kappa$, since all quantities in the definition of $\delta$ are the weights of vertex separators or trivializors. Thus, the error probability is $\leq(\delta / w(V))^{i} \leq 1 / 2$.

To bound the expected time write $\delta=(1-\epsilon) w(V)$ for $\epsilon$ between 0 and 1. The same calculation as Theorem 2.3 shows there are $O(1 / \epsilon)$ iterations. Thus, it suffices to show the expected time for an iteration is $O\left(\epsilon m n \log \left(n^{2} / m\right)\right.$ ).

Equation (2) gives the size of our vertex cover of $C G$ (recall Property (iii) of $C G$ ). Hence, we can apply Theorem 3.3 with $n_{T}=n-d_{o}(x)$. It shows that any quantity $\kappa(x, G)$ is computed in time $O((n-$ $\left.d_{o}(x)\right) m \log \left(n^{2} / m\right)$ ). The expected value of $d_{o}(x)$ is

$$
\begin{aligned}
& \sum\left\{w(x) d_{o}(x) / w(V): x \in V\right\} \\
& \quad=(1 / w(V)) \sum\{w(\{x: x v \in E\}): v \in V\} \geq n \delta / w(V),
\end{aligned}
$$

since any vertex $v$ has $w(\{x: x v \in E\}) \geq \delta$. Hence, $n-d_{o}(x)$ has expectation $\leq n(1-\delta / w(V))=n \epsilon$. So the expected time to compute $\kappa(x, G)$ is $O\left(\epsilon n m \log \left(n^{2} / m\right)\right.$ ). Similarly for computing $\kappa\left(x, G^{R}\right)$. Thus, the expected time for an iteration is as desired.

## 3. MINIMUM VERTEX COVER SPLITS

This section presents an algorithm to find a minimum vertex cover split. Recall that we are given a digraph with vertex cover $T$ and vertex $s \in T$; we seek an $\bar{s} t$-set that has the smallest possible in-degree subject to the constraint that $t \in T$. Our algorithm runs in time $O(n m)$ on a unit capacity multigraph and time $O\left(n m \log \left(n^{2} / m\right)\right)$ on a capacitated graph. M ore precisely write $n_{T}=|T|$. The time bounds are actually $O\left(n_{T} m\right)$ and $O\left(n_{T} m \log \left(2+n_{T}^{2} / m\right)\right.$ ), respectively (see Theorems 3.2 and 3.3).

We solve the special case of the problem where the undirected version of $G$ is bipartite and the vertex cover $T$ is one set of the bipartition. The general problem reduces to this special case as follows. Let an instance of the general problem be specified by digraph $G=(V, E)$, vertex cover $T$, and $s \in T$. Replace each edge $v w \in E$ that has $v, w \in T$ by two edges $v x, x w$. Here $x$ is a new vertex (on the two new edges and no others) and the two new edges have the same capacity as the original. The new graph is bipartite (no edge joins two vertices of $V-T$ since $T$ is a vertex cover of $G$ ). For any $t \in T$ a minimum $s t$-cut has the same value in both graphs (since a maximum st-flow has the same value). The parameters $n_{T}$ and $m$ change by at most a constant factor. Hence it suffices to consider the bipartite case. (We note that a version of our split algorithm works directly on the general problem. We do not present this version because the Relabel routine has more cases. We also note that the graphs generated by the reductions of Section 2 are already bipartite.)

We briefly review the preflow-push algorithm (see [10] for a complete treatment). Fix a digraph $G=(V, E)$ with distinguished vertices $s$ (the source) and $t$ (the sink) and capacity function $c: E \rightarrow \mathbf{R}_{+}$. Without loss of generality assume an edge $e$ belongs to $E$ if and only if its reverse edge,
denoted $e^{R}$, belongs to $E$. A flow is function $f: E \rightarrow \mathbf{R}_{+}$satisfying the capacity constraint (for each edge $e \in E, f(e) \leq c(e)$ ) and the conservation constraint (each vertex $\neq s, t$ has excess 0 , where the excess at vertex $v$ is $e(v)=f(\{w v: w v \in E\})-f(\{v w: v w \in E\}))$. A preflow obeys the same conditions except that instead of conservation we require that each vertex $v \neq s$ has $e(v) \geq 0$. The value of a flow or preflow is $e(t)$; a maximum flow ( preflow) is one with maximum value. The maximum value of a preflow is the same as that of a flow, because any preflow can be converted to a flow of the same value.
Fix a preflow $f$. A vertex $v \neq s, t$ is overflowing if $e(v)>0$. An edge $e$ has residual capacity $r(e)=c(e)-f(e)+f\left(e^{R}\right)$. An edge with $r(e)>0$ is a residual edge; an edge with $r(e)=0$ is saturated. To push $\delta$ units of flow along $e$ means to increase $f(e)$ by $\delta$ and then reduce both $f(e)$ and $f\left(e^{R}\right)$ by $\min \left\{f(e), f\left(e^{R}\right)\right\}$. Pushing up to $\min \{e(v), r(v w)\}$ units of flow along edge $v w$ keeps $f$ a valid preflow. A push of $r(v w)$ units is a saturating push (it makes $v w$ saturated).
A distance function is a mapping $d: V \rightarrow \mathbf{N}$ such that any residual edge $v w$ not incident to $s$ satisfies $d(v) \leq d(w)+1$.

The preflow-push algorithm of [10] computes a maximum preflow. To do this it maintains a preflow and corresponding distance function, repeatedly pushing flow on residual edges $v w$ that have $d(v)=d(w)+1$ and increasing distance labels until there are no overflowing vertices. At that point the preflow is maximum and the algorithm halts.
The algorithm of [13] finds a minimum unrestricted cut (i.e., an $x y$-cut of minimum capacity where $x$ and $y$ are arbitrary). It maintains a preflow from source $S$ to sink $t$. Here $S$ is a contraction of vertices of $V$ (we use $S$ to denote both the subset of $V$ and its contraction); $t$ is a vertex not in $S$. Initially $S$ consists of one vertex $s \in V$. The algorithm repeatedly finds a maximum preflow from $S$ to $t$ (following [10]) and then transfers $t$ to $S$ and chooses a new sink vertex $t^{\prime} \notin S$ that has minimum distance $d\left(t^{\prime}\right)$. At certain points the algorithm makes a set of vertices "dormant." The dormant vertices are temporarily ignored-the algorithm pushes flow and increases distance labels only on vertices that are "awake" (i.e., vertices not belonging to $S$ or a dormant set). When no awake vertices remain (because they have all been chosen as sinks and transferred to $S$ ) the algorithm awakens the last dormant set and repeats this process on them. A s mentioned, for each sink $t$ the algorithm computes a maximum preflow from $S$ to $t$; its value equals that of a minimum $S t$-cut. The minimum unrestricted cut of the given graph $G$ has value equal to the smallest of these $S t$-cuts on $G$ or $G^{R}$.

Our algorithm extends the Hao -Orlin minimum unrestricted cut algorithm. The main change is a new operation Join which adjoins an awake vertex to a previously created dormant set.

### 3.1. Algorithm Statement

W e present the algorithm using the same organization as [13]: We divide the algorithm into a number of subroutines. We state each routine and then briefly comment on how it works and how it differs from [13]. The formal analysis of our algorithm does not rely on these comments.

W e use the following notation. The algorithm maintains a partition of $V$ into sets $S$ (the source set), $D_{i}, i=1, \ldots, \gamma$ (the dormant sets), and $W$ (the awake set). Let $\mathscr{P}$ denote this partition. ( $\mathscr{P}$ changes over time.) Also define $D_{0} \equiv S$. The function $d$ is the distance function. For a set $U \subseteq V$, $d_{\text {min }}(U)=\min \{d(v): v \in U\}$ and $d_{\text {max }}(U)=\max \{d(v): v \in U\}$.

R ecall that we are given $G, T, s$ where $G$ is a digraph whose undirected version is bipartite, $T$ is one set of the bipartition, and $s \in T$. We seek an $\bar{s}$-set that contains a vertex of $T$ and has the smallest possible in-degree. The following main routine returns with $W^{*}$ equal to the desired set.

```
procedure Min_VC_Split(G,T,s);
```

1. Initialize; /* initialize flow, distances $d(v)$, etc. */

## repeat

3. \{New_Sink; / * choose new sink, etc. */
4. while a vertex of $W$ is overflowing do
5. \{choose an overflowing vertex $v \in W$;
6. if there is a residual edge $v w$ with $w \in W$ and $d(v)=$ $d(w)+1$ then
7. push $\min \{e(v), r(v w)\}$ units of flow along $v w$
8. else Relabel $(v)$ \}
9. if $\rho(W, G)<\rho\left(W^{*}, G\right)$ then $\left.W^{*} \leftarrow W\right\}$
10. until $T \subseteq S \cup t$;

This is essentially the main routine of [13]. The Initialize and New_Sink routines collectively initialize the data structures for the first iteration. Here is Initialize; New_Sink, which selects the next sink $t$ of the preflow, is given later.
procedure Initialize;

1. $\quad W^{*} \leftarrow$ any $\bar{s}$-set separating $T$, e.g., $\{v\}$ for some $v \in T-s$;
2. $S \leftarrow \varnothing ; \gamma \leftarrow 0$; / * $S$ is the same set as $D_{0} * /$
3. $W \leftarrow V$; $t \leftarrow s$;
4. set the flow in each edge of $G$ to 0 ;
5. for $v \in V$ do if $v \in T$ then $d(v) \leftarrow 0$ else $d(v) \leftarrow 1$;

We shall see that in the first iteration New_Sink completes the initialization by setting $S$ to $\{s\}$ and $t$ to the first sink and saturating the edges directed from $s$.
The Relabel routine either increases the distance label $d(v)$ or adjusts the dormant sets. The Create_Dormant routine creates a new dormant set as in [13]. The Join routine adjoins $v$ to an existing dormant set.
procedure Relabel $(v) ; / * v \in W * /$

1. if $v$ is the only vertex in $W$ at distance $d(v)$ and $d(v)<d_{\text {max }}(W)$

## then

2. Create_Dormant $(\{u \in W: d(u) \geq d(v)\})$
3. else if no residual edge goes from $v$ to $W$ then
4. if $v \in T$ then Create_Dormant $(\{v\})$ else Join $(v)$
5. else $d(v) \leftarrow \min \{d(w)+1$ : $v w$ a residual edge, $w \in W\}$;
procedure Create_Dormant $(U)$; /* $U \subseteq W * /$
6. $W \leftarrow W-U ; \gamma \leftarrow \gamma+1 ; D_{\gamma} \leftarrow U$;
procedure $\operatorname{Join}(v) ; / * v \in W-T * /$
7. $W \leftarrow W-v ; D_{\gamma} \leftarrow D_{\gamma} \cup v ; d(v) \leftarrow d_{\text {min }}\left(D_{\gamma} \cap T\right)+1$;

In line 1 of Relabel if $v$ is the only awake vertex at distance $d(v)$ and $d(v)=d_{\max }(W)$ then the condition of line 3 holds; i.e., no residual edge goes from $v$ to $W$. (This follows from the $d$-V alidity property below, which says $d$ is a valid distance function.) In Join the resetting of $d(v)$ preserves the validity of $d$. This resetting can change $d(v)$ arbitrarily; in particular, $d(v)$ can decrease. This contrasts with most preflow-push algorithms where any vertex $v$ has $d(v)$ nondecreasing (e.g., [1, 2, 10, 11, 13, 14]). Finally note that Join can add $v \notin T$ to $S=D_{0}$.

New_Sink transfers the previous sink to $S$ and saturates all edges directed from it. Then it chooses an awake vertex of $T$ as the next sink. This may necessitate waking the last dormant set $D_{\gamma}$.
procedure New_Sink;

1. $W \leftarrow W-t ; S \leftarrow S \cup t$;
2. saturate each residual edge $t v$ directed from $t$;
3. if $W \cap T=\varnothing$ then
4. $\quad$ for $v \in W$ do Join $(v)$;
5. $\left.W \leftarrow D_{\gamma} ; \gamma \leftarrow \gamma-1\right\}$
6. for each $v \in W-T$ with $d(v)=d_{\text {min }}(W)$ do $d(v) \leftarrow d(v)+2$;
7. choose $t$ as a vertex of $W$ with $d(t)=d_{\text {min }}(W)$;

The purpose of the Join operations in line 4 is to update distance labels. Now consider the loop of line 6 that increases distance values by 2 . It applies when the awake vertex $t^{\prime} \in T$ with smallest distance value has $d\left(t^{\prime}\right)=d_{\text {min }}(W)+1$. Thus, the loop ensures that the new sink $t^{\prime}$ has the smallest distance value. This is needed in the following two situations.

In both situations, New_Sink is entered with the old sink $t$ being the only vertex at distance $d(t)$. The first situation is when some other vertex of $T$ is awake. We shall see this implies there is such a vertex $t^{\prime}$ with $d\left(t^{\prime}\right)=d(t)+2$. This quantity equals $d_{\text {min }}(W)+1$ after line 1 deletes $t$ from $W$. Hence, the loop ensures $d\left(t^{\prime}\right)=d_{\text {min }}(W)$ as desired. The second situation is when no other vertex of $T$ is awake and line 5 wakens $D_{\gamma}$. If $D_{\gamma}$ was created (in Relabel) when $v$ was the only awake vertex at distance $d(v)$ for some $v \notin T$, then a vertex $t^{\prime} \in T \cap D_{\gamma}$ with smallest distance value has $d\left(t^{\prime}\right)=d(v)+1$. Thus, again the loop ensures $d\left(t^{\prime}\right)=$ $d_{\text {min }}(W)$. This concludes the statement of our algorithm.

### 3.2. Algorithm Analysis

We prove that our algorithm has properties similar to the ones shown in [13]. We state these nine properties below and discuss how they are used in the overall argument. Then we present the proofs of the properties.

The first two properties, besides being crucial overall, will ensure that the algorithm is well-defined and halts. Recall the definition of $\mathscr{P}$ from the start of Section 3.1.

Parity property. $d(v)$ is even if $v \in T$ and odd if $v \notin T$.
$d$-Consecutiveness property. If $U \in \mathscr{P}-S$ then $d(U)$ is a set of consecutive integers.
The next paragraph shows the above two properties imply that a dormant set always contains a vertex of $T$. This fact has two important consequences. First, it makes the algorithm well-defined-in Join, the assignment to $d(v)$ needs $D_{\gamma} \cap T \neq \varnothing$. Second, it guarantees that line 7 of New_Sink always chooses a sink $t$ that belongs to $T$. (This follows by examining lines 3-7 of New_Sink and using the fact, $d$-Consecutiveness and Parity.)

We now prove that a dormant set always contains a vertex of $T$. The dormant set $D_{0}$ contains $s \in T$ (after the first execution of New_Sink). A dormant set created in line 4 of Relabel (by Create_Dormant) clearly contains a vertex of $T$. Consider a dormant set created in line 2 of Relabel. The new dormant set has vertices with at least two distinct distance values. Hence the desired property follows from $d$-Consecutiveness and Parity.
The next two properties are important for correctness of the algorithm.
$d$-Validity property. A residual edge $v w$ with $v$ and $w$ in the same set of $\mathscr{P}-S$ has $d(v) \leq d(w)+1$.

Dormancy property. No residual edge enters $W$ or goes from $D_{i}$ to $D_{j}$, $0 \leq i<j$.

The next property implies the algorithm is correct.
Optimality property. Whenever Min_vc_Split updates $W^{*}, W$ is an $\overline{S \cap T} t$-set of minimum in-degree $\rho(W, G)$.

To see this implies correctness, let $U$ be a minimum $s$-split of $T$. Let $t$ be the first vertex of $U$ chosen as the sink by new_Sink. Thus, when $t$ is the sink $U$ is an $\overline{S \cap T} t$-set. (Note that $U$ need not be an $\bar{S} t$-set; i.e., $U$ may contain vertices of $S-T$.) Now the Optimality property implies Min_VC_Split finds a set $W$ of the same in-degree as $U$.

The last four properties are used in the efficiency analysis.
Increasing distance property. A value $d(v)$ can decrease only in an operation Join $(v)$. An operation Relabel $(v)$ either increases $d(v)$ or makes $v$ dormant.

The main purpose of the next property is to bound the total increase in $d(v)$ in operations other than Join $(v)$. For any vertex $v \in V$ define the following measure of increase:
$I(v)=$ (the total amount that $d(v)$ has increased in all operations other

$$
\text { than Join }(v))+2 \times(\text { the number of operations Join }(v)) \text {. }
$$

A vertex $v \in T$ is never in a join operation. Hence, at any moment $I(v)=d(v)$ for any $v \in T$.
$d$-Bound property. Each value $d(v)$ is less than $2 n_{T}$. Each value $I(v)$ is at most $2 n_{T}$.

Transfer bound property. New_Sink awakens fewer than $n_{T}$ dormant sets. Create_Dormant creates fewer than $n_{T}$ dormant sets. A given vertex becomes dormant fewer than $n_{T}$ times.

Saturation bound property. Any edge and its reverse are collectively saturated at most $2 n_{T}$ times.

Now we prove the properties. The arguments are similar to [13] except for the $d$-Bound. For the first property and others note that the algorithm changes distance values in four places: Initialize (line 5), Relabel (line 5), Join, and New_Sink (line 6).

Proof of Parity. This property is first established by Initialize. It is preserved when line 5 of Relabel changes a distance value because $G$ is bipartite. It is obviously preserved when Join or line 6 of New_Sink changes a distance value.

Proof of $d$-Consecutiveness. The argument is by induction on the number of steps of the algorithm. $d$-Consecutiveness is first established by Initialize.
Consider the Relabel routine. The call to Create_Dormant in line 2 obviously preserves $d$-Consecutiveness (for both $W$ and the new dormant set). If this call is not made then either $d(v)=d_{\max }(W)$ or there is another vertex at distance $d(v)$. In either case $d(W-v)$ is a set of consecutive integers. This implies the $d$-Consecutiveness of $W$ is preserved in all other cases of Relabel (i.e., in the calls to Create_Dormant or Join in line 4 and when line 5 redefines $d(v)$ ). In line $4 d$-Consecutiveness is obvious for the new dormant set $\{v\}$ formed by Create_Dormant. Finally any call to Join preserves $d$-Consecutiveness for $D_{\gamma}$ when it redefines $d(v)$.

Next consider the New_Sink routine. $d$-Consecutiveness is preserved when line 1 deletes $t$ from $W$ since $d(t)=d_{\min }(W)$. The calls to Join (line 4) preserve $d$-Consecutiveness as noted above. Line 6 preserves $d$-Consecutiveness of $W$ by $d$-Consecutiveness prior to this line and Parity.

Proof of $d$-Validity and Dormancy. These two properties are proved together by induction on the number of steps of the algorithm. Both properties are established when Initialize sets the distance values in line 5.

Consider a push operation in line 7 of Min_VC_Split. The only residual edge that can be created by the push is $w v . d$-V alidity is preserved since $d(w)<d(v)$ by line 6 . Dormancy is preserved since $w \in W$ by line 6 .

Consider the Relabel routine. We begin by checking Dormancy is preserved. Consider the call to Create_Dormant in line 2, and suppose it makes $D_{\gamma}$ the new dormant set. Take any residual edge $u w$ with $u \in D_{\gamma}$ and $w$ belonging to $W$ before the call. Then $d(w) \geq d(v)$, by $d$-V alidity if $u \neq v$ and by line 6 of Min_VC_Split if $u=v$. Thus, $w \in D_{\gamma}$ and Dormancy is preserved. Line 4 of Relabel preserves Dormancy by the test of line 3.

N ext we check that $d$-V alidity is preserved in Relabel. The first case to consider is an operation Join $(v)$. Take any vertex $w \in D_{\gamma}$. A residual edge joining $v$ and $w$ is directed from $v$ to $w$ (by Dormancy before Join $(v)$ ). Also $w \in T$ (since $v \notin T$ ). Hence, the new value of $d(v)$ satisfies the validity inequality.
The second case to consider is when line 5 of Relabel redefines $d(v)$. The new value obviously preserves $d$-V alidity for edges directed from $v$.

For edges directed to $v$ observe that this operation increases $d(v)$ (by $d$-Validity prior to line 5 and the definition of $v$ in line 6 of Min_VC_Split).

Finally consider the New_Sink routine. Dormancy is preserved when line 2 saturates the edges directed from $t$. When $W \cap T=\varnothing$ a Join operation preserves $d$-V alidity, by the argument above. A lso it is obvious that in line 5 redefining $W$ to $D_{\gamma}$ preserves Dormancy. Lastly we show line 6 preserves $d$-V alidity. Immediately before line 6 increases a value $d(v)$, any vertex $w \in W \cap T$ has $d(w) \geq d(v)+1$ by Parity. This implies $d$ V alidity.

Proof of Optimality. When Min_VC_Split updates $W^{*}$ all vertices with negative excess are in $S \cap T$. Hence, we have a preflow from $S \cap T$ to $t$. The set $W$ contains $t$ but no vertex of $S \cap T$. No vertex of $W$ is overflowing, by definition. Every edge of $G$ entering $W$ is saturated, by D ormancy. This implies we have a maximum value preflow from $S \cap T$ to $t$ and $W$ is an $\overline{S \cap T} t$-set of minimum in-degree.

Proof of Increasing Distance. Line 5 of Relabel increases a distance value $d(v)$ (as shown in the proof of $d$ - V alidity and Dormancy). This implies the second part of the Increasing Distance property. Line 6 of New_Sink also increases distance values.

We prove the $d$-Bound using the following lemma.
Lemma 3.1. Whenever a vertex $v$ is awake,

$$
\begin{equation*}
I(v)<2|S \cap T|+d(v)-d(t) \tag{3}
\end{equation*}
$$

Immediately before an operation Join(v),

$$
\begin{equation*}
I(v)<2|(S \cup W) \cap T| . \tag{4}
\end{equation*}
$$

Let us show that the lemma implies the $d$-Bound. Inequality (3) implies the first part of the $d$-Bound: First consider a vertex $v \in T$. When $v$ is chosen as the sink $d(v)$ has achieved its maximum value and (3) gives $d(v)=I(v)<2|S \cap T| \leq 2\left(n_{T}-1\right)$. Next consider a vertex $v \notin T$. The algorithm always defines $d(v)$ as one more than the distance of a vertex in $T$, so we always have $d(v)<2 n_{T}$.

Inequality (4) essentially implies the second part of the $d$-Bound. For the second part we need only consider vertices $v \notin T$. First consider a vertex $v \notin T$ that becomes dormant in an operation Join $(v)$ and is never subsequently reawakened. B efore this join inequality (4) shows $I(v)<2 n_{T}$. The join increases $I(v)$ by 2 , so its final value is $\leq 2 n_{T}$ (since $I(v)$ is even). Thus, the second part of the $d$-Bound holds for $v$. To extend this to
all vertices $v \notin T$ we make the convention that when $T \subseteq S \cup t$ in line 10 of Min_VC_Split, the algorithm actually continues to execute New_Sink, line 4 of which performs Join $(v)$ for each $v$ remaining in $W$. With this convention the previous argument applies to every $v \notin T$. We conclude that Lemma 3.1 implies the $d$-Bound.

Before proving Lemma 3.1 we give a useful inequality. Consider the moment when an execution of create_Dormant $(U)$ is completed. Let $v \in U \cap T$ have minimum distance; i.e., $d(v)=d_{\text {min }}(U \cap T)$. Then

$$
\begin{equation*}
d(v) \leq d(t)+2|W \cap T| . \tag{5}
\end{equation*}
$$

This follows from $d$-Consecutiveness and Parity. Similar reasoning shows that (5) holds (with strict inequality) immediately before an execution of Join $(v)$ in Relabel $(v)$.

Proof of Lemma 3.1. The argument is by induction on the number of steps of the algorithm. At the start of the algorithm (after the first execution of New_Sink) for any vertex $v, I(v)=d(t)=0, S=\{s\}$, and $d(v) \geq 0$, so (3) holds.

Consider the Relabel routine. When line 5 of Relabel increases $d(v)$, obviously (3) is preserved. The only remaining case to check in Relabel is when Join( $v$ ) is executed. We wish to show (4) holds right before the join. Inequality (5) shows right before the join,

$$
0 \leq 2|W \cap T|+d(t)-d(v) .
$$

Also (3) holds right before the join since $v$ is awake. Adding these inequalities gives (4).
N ow consider the New_Sink routine. The following three cases exhaust all the possibilities.

Case 1: $v$ is awake throughout New_Sink. We first claim that no set gets awakened in New_Sink. To prove this it suffices to show $W \cap T \neq \varnothing$ in line 3 . This holds if $v \in T$ since then $v \in W \cap T$ by Case 1 . It holds if $v \notin T$ since $\operatorname{Join}(v)$ is not executed in Case 1.

Now we show that New_Sink does not decrease the quantity $2|S \cap T|$ $-d(t)$ on the right-hand side of (3). Line 1 increases $2|S \cap T|$ by 2 . When line 7 chooses a new value for $t$ the quantity $d(t)$ increases by at most 2 from its previous value, by the claim, $d$-Consecutiveness, and Parity. The net effect is that $2|S \cap T|-d(t)$ does not decrease.

Finally if line 6 increases $d(v)$ both sides of (3) increase by 2 . We conclude that (3) is preserved in Case 1.
Case 2: $v$ is awake when New_Sink starts and becomes dormant in New_Sink. If $v \in T$ it becomes dormant in line 1 and there is nothing
to prove. Suppose $v \notin T$, so $v$ becomes dormant in line 4 when $\operatorname{Join}(v)$ is executed and $v$ is reawakened in line 5 . Immediately before New_Sink starts, $d(v)=d(t)+1$ (by the test of line 3). Thus, (3) gives

$$
\begin{equation*}
I(v)<2|S \cap T|+1 \tag{6}
\end{equation*}
$$

Inequality (6) becomes $I(v)<2|S \cap T|-1$ when line 1 adds $t$ to $S$. This gives the desired inequality (2) right before the operation Join $(v)$ (note $W \cap T=\varnothing$ at this moment). The join increases $I(v)$ by 2 , so (6) holds once again. The join also sets $d(v)$ so that when $v$ awakens in line 5 and the next sink $t$ is chosen in line 7, we have $d(v)=d(t)+1$. Thus (6) implies that (3) holds at the end of New_Sink, as desired.

Case 3: $v$ is dormant when New_Sink starts. If $v$ remains dormant throughout the execution of New_Sink there is nothing to prove. So assume $v$ is awakened in line 5. We examine the two possibilities for how $v$ became dormant.

The first possibility is that an operation Create_Dormant $(U)$ made $v$ dormant. Right after the execution of Create_Dormant $(U)$ inequality (5) holds, so

$$
0 \leq 2|W \cap T|+d(t)-d_{\min }(U \cap T) .
$$

Inequality (3) holds right before Create_Dormant( $U$ ). It still holds right after Create_Dormant $(U)$ since no term changes. Adding (3) to the above inequality gives

$$
\begin{equation*}
I(v)<2|(S \cup W) \cap T|+d(v)-d_{\min }(U \cap T) . \tag{7}
\end{equation*}
$$

By the time $v$ gets awakened in New_Sink, all vertices in $W \cap T$ in (7) have been transferred to $S$. Furthermore some vertex at distance $d_{\text {min }}(U \cap$ $T$ ) becomes the sink. H ence, inequality (7) implies that (3) holds at the end of New_Sink, as desired.

The second possibility is that an operation Join $(v)$ in Relabel made $v$ dormant. Right before Join $(v)$ inequality (4) holds. By the time $v$ gets awakened in New_Sink, all vertices in $W \cap T$ in this inequality have been transferred to $S$. Also the Join $(v)$ operation increases $I(v)$ by 2 . Thus, the inequality becomes $I(v)<2|S \cap T|+2$. Since $I(v)$ is even this implies $I(v)<2|S \cap T|+1$. The join operation sets $d(v)$ so $d(v)=$ $d(t)+1$ when the next sink $t$ is chosen. This implies that (3) holds at the end of New_Sink, as desired.

Proof of Transfer Bound. There are $n_{T}-1$ executions of New_Sink, since each one chooses a new vertex of $T$ as sink. This gives the first part of the Transfer Bound. There are $<n_{T}-1$ executions of

Create_Dormant, since each one creates a dormant set which eventually gets awakened in line 5 of New_Sink. This gives the second part. A given vertex becomes dormant $<n_{T}$ times, since each time its new dormant set either equals $D_{0}$ or was created by Create_Dormant.

Proof of Saturation Bound. Consider the saturating pushes from $v$ to $w$ or $w$ to $v$ for two fixed vertices $v \notin T$ and $w \in T$. We claim that during the interval of time after one such push and before the next, the quantity $I(v)+d(w)$ increases by at least 2 . W hen the algorithm ends this quantity is $<4 n_{T}$ ( $d$-Bound). Thus, the claim implies there are $\leq 2 n_{T}$ saturating pushes total.

Now we prove the claim. The claim is obvious if Join(v) occurs in the interval, so assume it does not. Thus, neither $d(v)$ nor $d(w)$ decreases (I ncreasing Distance). Without loss of generality suppose the first saturating push is from $v$ to $w$. The next push along $v w$ or $w v$ goes from $w$ to $v$, which implies $d(w)$ has increased by $\geq 2$ (line 6 of Min_TC_Split).

Now we estimate the time for the algorithm. We use the inequality $n \leq m$, which holds since a vertex not on any edge can be deleted.

First we give some necessary data structures. We use the current edge data structure to scan the edges incident to a vertex [10]. We use the additional rule that every operation $\operatorname{Join}(v)$ resets $v$ 's current edge pointer to the beginning of its adjacency list.
These rules ensure that whenever $v$ 's current edge pointer reaches the end of $v$ 's adjacency list, every edge $v w$ has been scanned with $d(v)$ equal to its current value. Hence, as in [10] each time $v$ 's entire adjacency list has been scanned, Relabel $(v)$ is executed. The latter occurs $O\left(n_{T}\right)$ times by the Increasing Distance property, $d$-Bound, and Transfer Bound. A vertex $v \notin T$ may also have $O\left(n_{T}\right)$ partial scans of its adjacency list, each ending in the operation $\operatorname{Join}(v)$. We conclude that any given edge is scanned $O\left(n_{T}\right)$ times. Thus, the total time for scanning edges (line 6 of Min_VC_Split and lines 3 and 5 of Relabel) is $O\left(n_{T} m\right)$.

We maintain several lists of awake vertices. To choose an overflowing vertex (in line 4-5 of Min_VC_Split) we maintain a list of the overflowing vertices of $W$. To implement lines $1-2$ of Relabel we maintain a list of vertices of $W$ at each distance value between $d_{\min }(W)$ and $d_{\max }(W)$. We also mark each awake vertex, to implement line 6 of Min_VC_Split and lines 3 and 5 of Relabel. These data structures are initialized each time a dormant set is awakened, in total time $O\left(n_{T}(m+n)\right)=O\left(n_{T} m\right)$.

We record the value $d_{\text {min }}(D \cap T)$ for each dormant set $D$. Thus, each join operation uses time $O(1)$, giving time $O\left(n_{T} n\right)$ total for joins.

Each push takes $O(1)$ time. The Saturation Bound implies there are $O\left(n_{T} m\right)$ saturating pushes. It follows from this discussion that excluding nonsaturating pushes, the total time for Min_VC_Split is $O\left(n_{T} m\right)$.

Now suppose Min_VC_Split is executed on a unit capacity digraph $G$. (In such a graph every edge has capacity one but parallel edges are allowed.) This is the case needed for all results of Section 2 except weighted connectivity. For such graphs we can consider the residual graph also to be unit capacity. That is, instead of defining the residual capacity of $e$ as $r(e)$, the residual graph contains $r(e)$ copies of $e$, each having capacity 1. (Thus, each edge of $G$ corresponds to a unique edge in any residual graph.) The Saturation Bound still holds. Furthermore in such residual graphs every push is saturating, i.e., there are no nonsaturating pushes. We have thus completely analyzed the time for such graphs and can conclude the following:

Theorem 3.2. A minimum s-split of a vertex cover $T$ can be found in time $O\left(n_{T} m\right)$ in a digraph with unit edge capacities, for $n_{T}=|T| \leq n$. The space is $O(m)$.

We turn to the implementation of Min_VC_Split on capacitated digraphs. We first achieve time $O\left(n_{T} m \log \left(n^{2} / m\right)\right)$ by incorporating dynamic trees, and then we improve the logarithmic factor. The modifications and analysis are similar to the maximum flow algorithm of [10]. We begin by summarizing that implementation in the context of Min_VC_Split. Then we sketch the changes to our algorithm from [10] and the changes in the analysis of our algorithm. We conclude by sketching how to improve the time bound to $O\left(n_{T} m \log \left(n_{T}^{2} / m\right)\right.$ ). We assume the reader is familiar with [10].
The dynamic tree algorithm of [10] maintains a forest of dynamic trees. Consider a dynamic tree edge $v w . v w$ is the current edge of $v$, with $w$ the parent of $v$ in the dynamic tree. Furthermore edge $v w$ satisfies the condition of line 6 of Min_VC_Split, so line 7 can push flow along it. The algorithm pushes flow along paths in dynamic trees, specifically along paths from a vertex $v$ to the root of its dynamic tree. The algorithm also pushes flow along single edges not in a dynamic tree. It maintains the following overflowing vertex invariant: A ny overflowing vertex is the root of the dynamic tree containing it. This is done by using a routine send to push flow in a dynamic tree. send $(v)$ repeatedly pushes flow from $v$ to the root of its dynamic tree until either $v$ has no excess or $v$ becomes a dynamic tree root. send also cuts any newly saturated edge, to preserve the definition of dynamic tree edges. Line 7 of Min_VC_Split either links the dynamic trees of $v$ and $w$ and then does $\operatorname{send}(v)$, or pushes flow along edge $v w$ and then does send $(w)$.

W hen a vertex $v$ is relabelled, i.e., the distance value $d(v)$ is increased as in line 5 of Relabel, the algorithm performs an operation cutchildren(v), where cut-children(v) cuts every dynamic tree edge corre-
sponding to a child of $v$. This preserves the definition of dynamic tree edges.

The algorithm chooses overflowing vertices to process (line 5 of Min_VC_Split) using a queue of overflowing vertices (the FIFO pre-flow-push algorithm). Each time an overflowing vertex $v$ is chosen from the head of the queue, pushes from $v$ are done (as described above for line 7 of Min_VC_Split) until either the excess of $v$ becomes 0 or a relabel operation is done to increase $d(v)$. A vertex is added to the end of the queue when it becomes overflowing as a result of a push; also if $v$ is relabelled it is added to the end of the queue.

E ach dynamic tree is maintained to contain at most $k$ vertices for some parameter value $k$. Thus, each dynamic tree operation uses time $O(\log k)$. Choosing $k=\log \left(n^{2} / m\right)$ gives the time bound $O\left(n m \log \left(n^{2} / m\right)\right.$ ) to find a maximum value flow [10].

The dynamic tree implementation of Min_VC_Split is the same as [10]. It does some additional operations that keep each dynamic tree contained in one set of the partition $\mathscr{P}$, and other additional operations corresponding to the extra operations of Min_VC_Split compared to [10]. The additional operations are as follows.
Consider Relabel(v). Recall that $v$ is the root of its dynamic tree since it is overflowing. If the condition of line 3 of Relabel holds then cut-children $(v)$ is done before Create_Dormant $(\{v\})$ or Join $(v)$ is called in line 4. The call to Create_Dormant in line 2 transfers every dynamic tree of $W$ whose root is at distance $\geq d(v)$ to the new dormant set $D_{\gamma}$. ( $D_{\gamma}$ equals the set of all vertices in these trees. In proof, $v$ 's dynamic tree is the only dynamic tree of $W$ that contains a vertex at distance $d(v)$, by the test of line 1 . This implies any dynamic tree containing a vertex of $D_{\gamma}$ has its root at distance $\geq d(v)$, since each edge $u w$ of a dynamic tree has $d(u)=d(w)+1$ and the root has the smallest distance value.)

Next consider New_Sink. Line 1 does cut-children $(t)$ when $t$ is deleted from $W$. Line 6 does cut-children $(v)$ when $d(v)$ is increased by 2.

Lastly note that at the end of New_Sink the overflowing vertex invariant can be violated: A vertex $w \in W$ can have positive excess without being a dynamic tree root. (This is caused by the edge saturations in line 2 of New_Sink. These saturations could occur in previous executions of New_Sink if $W$ was just awakened in line 5.) To restore the invariant New_Sink performs send $(w)$ for each such $w$.
We turn to the analysis of the dynamic tree algorithm. Cut-children operations do $O\left(n_{T} m\right)$ extra dynamic tree cut operations. This follows since after an operation cut-children $(v), v$ becomes dormant or $d(v)$ increases. New_Sink does $\leq m$ new send operations. Since this accounts for all new operations, the analog of Lemma 5.1 of [10] holds (with the same proof): The dynamic tree algorithm runs in time $O\left(\left(n_{T} m+\right.\right.$
$\left.n_{Q A}\right) \log k$ ), where $n_{Q A}$ is the total number of times a vertex is added to the queue of overflowing vertices.
There are $O\left(n_{T} n\right)$ passes over the queue. The argument is essentially the same as [10]: When an overflowing vertex $v$ is chosen from the head of the queue (line 5 of Min_VC_Split) it is processed until either its excess becomes 0 or Relabel $(v)$ is done, which either makes $v$ dormant or increases $d(v)$. D efine the potential function of [10], $\Phi=\max \{d(v): v \in W$ is overflowing\}. If a pass over the queue increases $\Phi$ some distance $d(v)$ increases by at least the same amount. If a pass does not change $\Phi$ some distance $d(v)$ increases (since any vertex $x$ with $d(x)=\Phi$ at the start of the pass either gets 0 excess or becomes dormant). Thus, $O\left(n_{T} n\right)$ passes increase $\Phi$ or keep it the same.

It remains to show that $O\left(n_{T} n\right)$ passes decrease $\Phi$. The number of passes that decrease $\Phi$ is at most the total increase in $\Phi$. We have shown that all passes increase $\Phi$ by a total amount $O\left(n_{T} n\right)$. In addition $\Phi$ increases by a total amount $O\left(n_{T} n\right)$ when dormant sets are awakened, by the Transfer Bound. Hence, the total increase in $\Phi$ is $O\left(n_{T} n\right)$, as desired.

The analog of Lemma 5.2 of [10] is that $n_{Q A}=O\left(n_{T}\left(m+n^{2} / k\right)\right)$. The proof of [10] applies without change. Now taking $k=n^{2} / m$ achieves time $O\left(n_{T} m \log \left(n^{2} / m\right)\right)$ for Min_VC_Split.

This time bound suffices for the weighted connectivity algorithms of Section 2. It is desirable to replace $n$ by $n_{T}$ in the logarithmic factor for applications where $n_{T}$ is appreciably smaller. (This includes the graphs generated by our reduction from the general weighted vertex cover split problem to the bipartite case, since the reduction can add $\Theta(m)$ new vertices.) To achieve this improved time bound we incorporate the bipush approach of [2] into our dynamic tree algorithm. We change our additional send operations into bi-sends. We do not change our additional cut-children operations (since performing cut-children for the root of a dynamic tree does not create any new leaves of nontrivial dynamic trees). All other changes are as specified in [2]. The analysis is the same as [2].

Theorem 3.3. A minimum $s$-split of a vertex cover $T$ can be found in time $O\left(n_{T} m \log \left(2+n_{T}^{2} / m\right)\right)$ in a capacitated graph, for $n_{T}=|T| \leq n$. The space is $O(m)$.

## REFERENCES

1. R. K. A huja, T. L. M agnanti, and J. B. Orlin, "Network flows: Theory, A lgorithms, and A pplications," Prentice H all, E nglewood Cliffs, NJ, 1993.
2. R. K. A huja, J. B. Orlin, C. Stein, and R. E. Tarjan, Improved algorithms for bipartite network flow, SIAM J. Comput. 23 (1994), 906-933.
3. M. Becker, W. D egenhardt, J. D oenhardt, S. H ertel, G. K aninke, W. K eber, K. M elhorn, S. Näher, H. R ohnert, and T. Winter, A probabilistic algorithm for vertex connectivity of graphs, Inform. Process. Lett. 15 (1982), 135-136.
4. J. Cheriyan and J. H. Reif, Directed $s$ - $t$ numberings, rubber bands, and testing digraph $k$-vertex connectivity. Combinatorica 14 (1994), 435-451.
5. D. Coppersmith and S. Winograd, On the asymptotic complexity of matrix multiplication, SIAM J. Comput. 11 (1982), 472-492.
6. S. Even, A n algorithm for determining whether the connectivity of a graph is at least $k$, SIAM J. Comput. 4 (1975), 393-396.
7. S. E ven and R. E. Tarjan, Network flow and testing graph connectivity, SIAM J. Comput. 4 (1975), 507-518.
8. Z. Galil, Finding the vertex connectivity of graphs, SIAM J. Comput. 9 (1980), 197-199.
9. H. N. Gabow, Efficient splitting off algorithms for graphs, "Proc. 26th A nnual ACM Symp. on Theory of Comp., 1994," 696-705.
10. A. V. Goldberg and R. E. Tarjan, A new approach to the maximum-flow problem, J. Assoc. Comput. Mach. 35 (1988), 921-940.
11. G. Gallo, M.D.Grigoriadis, and R.E.Tarjan, A fast parametric maximum flow algorithm and applications, SIAM J. Comput. 18 (1989), 30-55.
12. M. R. Henzinger, A static 2-approximation algorithm for vertex connectivity and incremental approximation algorithms for edge and vertex connectivity, J. Algorithms 24 (1997), 194-220.
13. J. Hao and J. B. Orlin, A faster algorithm for finding the minimum cut in a directed graph, J. Algorithms 17 (1994), 424-446.
14. V. King, S. Rao, and R. Tarjan, A faster deterministic maximum flow algorithm, $J$. Algorithms 17 (1994), 447-474.
15. L. Lovász, "Combinatorial Problems and Exercises," 2nd ed., North-H olland, New Y ork, 1993.
16. N. Linial, L. Lovász, and A. Wigderson, Rubber bands, convex embeddings and graph connectivity, Combinatorica 8 (1988), 91-102.
17. H. Nagamochi and T. Ibaraki, A linear-time algorithm for finding a sparse $k$-connected spanning subgraph of a $k$-connected graph, Algorithmica 7 (1992), 583-596.

[^0]:    ${ }^{1}$ This research was supported in part by an NSF CAREER A ward, Grant CCR -9501712.

