

Probabilistic Analysis of Network Flow Algorithms

Richard M. Karp †

Rajeev Motwani †

Noam Nisan †

Computer Science Division,
573 Evans Hall, University of California,
Berkeley, CA 94720.

ABSTRACT

This paper is concerned with the design and probabilistic analysis of certain algorithms for the maximum-flow problem and related capacitated transportation problems. These algorithms run in linear time and, under certain assumptions about the probability distribution of edge capacities, obtain an optimal solution with high probability. The design of our algorithms is based on the following general method, which we call the mimicking method, for solving problems in which some of the input data is deterministic and some is random with a known distribution:

- (i) Replace each random variable in the problem by its expectation; this gives a deterministic problem instance that has a special form, making it particularly easy to solve;
- (ii) Solve the resulting deterministic problem instance;
- (iii) Taking into account the actual values of the random variables, mimic the solution of the deterministic instance to obtain a near-optimal solution to the original problem;
- (iv) Fine-tune this suboptimal solution to obtain an optimal solution.

We present linear time algorithms to compute a feasible flow in directed and undirected capacitated transportation problem instances. The algorithms are shown to be successful with high probability when the probability distribution of the input data satisfies certain assumptions. We also consider the maximum flow problem with multiple sources and sinks. Under probabilistic assumptions about the input, we show that with high probability the minimum cut isolates either the sources or the sinks, and we give a linear-time algorithm that produces a maximum flow with high probability.

† Supported by NSF Grant DCR-8411054



1. Introduction

Probabilistic analysis of combinatorial problems has been the subject of many recent investigations [19]. The starting point for such analysis is the assumption that the problem instances are drawn from a certain probability distribution. Under this assumption one studies the behaviour of the solution to the combinatorial problem or the performance of certain algorithms. The analysis often establishes that certain quick-and-dirty algorithms produce optimal or near-optimal solutions with high probability. In this paper we will be concerned with devising quick-and-dirty algorithms of this type for certain combinatorial problems.

The problems considered here are all closely related to the maximum flow problem. In particular, we present fast (linear time) algorithms for the maximum flow problem and certain versions of the transportation problem; these algorithms are guaranteed to succeed with high probability if the probability distribution of inputs satisfies certain assumptions. The main contribution of this work is the formulation and application of a new technique for solving problems with probabilistic inputs. We call this technique the probabilistic mimicking of deterministic solutions. The mimicking paradigm works as follows:

STAGE 1 {DETERMINISTIC RELAXATION} Suppose we are given a problem instance $P(\vec{X}, \vec{Y})$ with parameters (or input) \vec{X} and \vec{Y} . The vector \vec{X} represents the random values and the vector \vec{Y} represents the deterministic values. Construct a deterministic relaxation of P by replacing each random variable by its *expected* value. We now have a deterministic problem instance $P(E[\vec{X}], \vec{Y})$.

STAGE 2 {SOLUTION OF DETERMINISTIC RELAXATION} Construct a solution to the new problem instance $P(E[\vec{X}], \vec{Y})$. This instance is typically of a special form, for which a highly efficient algorithm is available.

STAGE 3 {MIMICKING PROCESS} Construct a solution to the original problem instance $P(\vec{X}, \vec{Y})$ by *mimicking* the solution constructed in the previous step. The exact form of the mimicking process will depend on the problem under consideration. The idea is to use the solution of the deterministic relaxation as a guide in solving the original problem instance.

STAGE 4 {FINE TUNING} In this stage we fine-tune the solution obtained in the previous stage to come up with an optimal solution to the original problem instance. The fine-tuning process is highly efficient because, with high probability, the solution resulting from the mimicking process is already close to the optimal.

The paper is organized as follows. In section 2 we define the problems under consideration and briefly discuss the previous work in this area. We also present an algorithm which will be used for the fine-tuning stage of our algorithms. In section 3 we present algorithms for the undirected transportation problem and the max-flow problem. In section 4 we present an algorithm for the directed transportation problem. Finally, in section 5, we discuss further work along these lines.

2. Preliminaries

2.1. Problem Definition

In this section we define the three problems under consideration. These are the *Maximum Flow Problem*, the *Supply Demand Problem* and the *Capacitated Transportation Problem*. We also present the classic *Min-Max* theorems associated with each of these problems (see [20] for more details). The problems are defined for the case of directed graphs only. These definitions apply to undirected cases also if we look upon each undirected edge (u,v) as representing the two directed edges (u,v) and (v,u) .

Maximum Flow Problem: Let D be a digraph with vertices $V(D)=S \cup I \cup T$, where S represents the *source* vertices, T represents the *sink* vertices and I represents the *intermediate* vertices. Let $E(D)$ represent the set of directed arcs in the digraph D . The *capacity* function, $c:E(D) \rightarrow \mathbb{R}^+$, assigns a non-negative real number $c(u,v)$ to every arc (u,v) in $E(D)$. Let $C=\{c(u,v)\}$ denote the set of all capacities. An instance of the max-flow problem consists of the digraph D and the set of capacities C .

Given an instance of a max-flow problem a function, $f:E(D) \rightarrow \mathbb{R}^+$, is called a *flow* function if it satisfies the following constraints,

$$\sum_{(v,u) \in E} f(v,u) = \sum_{(w,v) \in E} f(w,v) \quad \forall v \in I$$

$$0 \leq f(u,v) \leq c(u,v) \quad \forall (u,v) \in E$$

Let $f(A,B) = \sum_{a \in A, b \in B} f(a,b)$ where $A, B \subseteq V$. The value of a flow function, $val(f)$, is defined as follows,

$$val(f) = f(S, V-S) - f(V-S, S) = f(V-T, T) - f(T, V-T)$$

The max-flow problem is to find a maximum-value flow function on a given instance of a flow problem.

This problem has been studied quite extensively [20] and can be solved in polynomial time. The fastest known algorithm is due to Goldberg & Tarjan [14]. It runs in $O(nm \log(n^2/m))$ time, where $n = |V|$ and $m = |E|$. In the case of dense graphs the best max-flow algorithm known requires $O(n^3)$ time. We assume that the reader is familiar with the basic theory of network flows. This includes the notion of a *cut* (and the *capacity* of a cut) as well as the following theorem.

MAX-FLOW MIN-CUT THEOREM : The maximum value of an (S,T) -flow is equal to the minimum capacity of an (S,T) -cut.

Supply-Demand Problem: Let D be a digraph with vertex set $V=S \cup I \cup T$, where S represents the *source* vertices, T represents the *sink* vertices and I represents the *intermediate* vertices. Let E represent the set of edges in D . Associated with every source vertex $s \in S$ is a non-negative number a_s , called the *supply* at s . Similarly, associated with

every *sink* vertex $t \in T$ is a non-negative number b_t called the *demand* at t . Each edge $(u,v) \in E$ has a capacity $c(u,v)$ associated with it. The supply-demand problem is to find a *flow* which meets the demands at T from the supplies available at S . More formally, we are looking for a flow $f: E \rightarrow \mathbb{R}^+$ which satisfies the following conditions,

$$\begin{aligned} 0 &\leq f(u,v) \leq c(u,v) && \forall u,v \in V \\ f(s, V - \{s\}) - f(V - \{s\}, s) &\leq a_s && \forall s \in S \\ f(v, V - \{v\}) - f(V - \{v\}, v) &= 0 && \forall v \in I \\ f(V - \{t\}, t) - f(t, V - \{t\}) &\geq b_t && \forall t \in T \end{aligned}$$

A flow which satisfies all of the above conditions is called a *feasible* flow. Let $a(X) = \sum_{v \in X} a_v$ when $X \subseteq S$, similarly define $b(X)$ for $X \subseteq T$. Let $c(X, \bar{X})$ denote the capacity of the (X, \bar{X}) cut in G . The following theorem [20] gives necessary and sufficient conditions for a feasible flow to exist in an instance of the supply-demand problem.

SUPPLY-DEMAND THEOREM: A feasible flow exists in the supply-demand problem if and only if the following inequality holds for every subset $X \subseteq V$.

$$b(T \cap \bar{X}) - a(S \cap \bar{X}) \leq c(X, \bar{X})$$

This theorem requires that for every $X \subseteq V$ the excess of demands over supplies in \bar{X} must be less than the total capacity of the edges leading out of X . There is an easy transformation from an instance of a supply-demand problem to an instance of a max-flow problem and vice versa.

Capacitated Transportation Problem: The *capacitated transportation problem* (hereafter called the transportation problem) is a special case of the supply-demand problem. In this problem we have no intermediate vertices, i.e. $I = \emptyset$, and the graph is bipartite between S and T , i.e. all edges join a vertex in S to a vertex in T . We will only consider the cases where the sum of all the supplies is equal to the sum of all the demands. The Supply-Demand Theorem when applied to the transportation problem yields the following,

TRANSPORTATION THEOREM : A feasible flow exists in an instance of the transportation problem if and only if the following condition holds for every subset $X \subseteq S$,

$$\sum_{z \in X} a_z \leq \sum_{t \in T} \min(b_t, c(X, \{t\}))$$

Probabilistic Formulation: For each of these problems we will make two probabilistic assumptions:

- (a) each edge is present in the graph with probability $p(n)$, independent of the other edges.
- (b) the capacities of the edges are i.i.d. random variables which have bounded support.

In certain cases the probabilistic assumptions may only apply to some subset of the edges in the complete graph. For example, we may assume that a certain subset of the edges is always present (or absent) while the rest obey the first probabilistic assumption. We may also consider the case where certain edges have fixed (or deterministic) capacities while the rest obey the second probabilistic assumption.

2.2. Previous Work and Main Results

The study of random graphs was initiated by Erdős and Rényi in 1959 [5]. Subsequently, it has attracted a lot of attention and a rather elegant body of results has been built up [2]. The theory of random graphs is concerned with graphs drawn from certain probability spaces. A commonly used probability space is called $G_{n,p}$, where n is a positive integer and $0 \leq p \leq 1$. This probability space consists of all graphs on the vertex set $V = \{1, 2, \dots, n\}$. In a graph drawn from $G_{n,p}$, the probability that an edge is present is p independently for each edge. In other words, the probability of any graph with e edges is $p^e(1-p)^{N-e}$, where $N = \binom{n}{2}$. Similarly, we define the probability space $B_{n,p}$ which consists of all bipartite graphs on the vertex set $V = S \cup T$, where $|S| = |T| = n$. Independently for each edge, the probability that a graph drawn from $B_{n,p}$ contains that edge is p .

We shall require methods for finding perfect matchings in random graphs. A perfect matching in an undirected graph $G = (V, E)$ is a subgraph in which each vertex in V has degree one. Consider an instance of the transportation problem in which the underlying bipartite graph has all its edges directed from the sources to the sinks. The problem of finding a perfect matching in a bipartite graph is the special case of this directed transportation problem where all the supplies, demands and capacities are set to one. The problem of finding a perfect matching in a random bipartite graph drawn from $B_{n,p}$ is a special case of the probabilistic transportation problem.

Erdős and Rényi [6,7] proved the following theorem about the existence of perfect matchings in random bipartite graphs.

THEOREM (Erdős & Rényi): Let $p(n) = (\ln n + c)/n$, where c is a positive constant, then for $B \in B_{n,p}$

$$\lim_{n \rightarrow \infty} \text{Prob}[B \text{ has a perfect matching}] = e^{-2e^{-c}}$$

Angluin and Valiant [1] gave a fast algorithm, called the Proposal Algorithm, to construct a bipartite matching. This algorithm works in $O(n \log n)$ time and, with high probability, constructs a perfect matching in a random bipartite graph, provided that $p(n)$ is large enough. It builds up the matching progressively by having the source vertices propose to (or pair up with) the sink vertices in stages. At each stage, the lowest-numbered free (or unpaired) source vertex pairs up with a sink vertex randomly chosen from among its neighbours. This sink vertex is then removed from the list of neighbours of the source

vertex. The sink vertex breaks off any previous pairing with a source vertex.

THEOREM (Angluin & Valiant) : For all $\alpha > 0$ there exists $\beta > 0$ such that if $p(n) > (\beta \log n)/n$ then for $B \in B_{n,p}$,

$$\text{Prob}[\text{Proposal Algorithm fails on } B] = O(n^{-\alpha})$$

Consider now the Max-Flow problem where $|S| = |T| = 1$ and the edge capacities are i.i.d. random variables. This problem was considered earlier by Frank & Hakimi [10] and Frank & Frisch [11]. They studied the random variable $val(f)$ under the above assumptions and obtained results about its probability distribution. Karp [18], Grimmett & Welsh [15] and Grimmett & Suen [16] obtained strong asymptotic results for complete graphs with i.i.d. edge capacities. In particular, they showed that the minimum cut is almost surely the set of edges incident on the source or those incident on the sink. These results are all existential and do not yield any fast algorithms to construct the maximum flow.

Doulliez & Jamouille [4] proposed a decomposition method to determine the existence of a feasible flow in an instance of the transportation problem when the edge capacities are i.i.d. random variables. There are many results concerning bounds on the probability of existence of a feasible flow in probabilistic transportation problems [22]. Hassin & Zemel [17] studied the probabilistic version of the transportation problem where the underlying graph is complete and the edge capacities are random variables. A collection of random variables $\{c_k : k \in K\}$ is said to be *proper* with constant α iff for each $y \geq 0$ and each $k \in K$, $\text{Prob}[c_k \leq y \mid S] \leq \alpha y$. Here S represents any conditioning event concerning the variables $\{c_l : l \in K, l \neq k\}$. Hassin & Zemel considered the case where the edge capacities form a proper collection and the supplies/demands are bounded from above. Under these assumptions, they presented an algorithm which, with high probability, computes a feasible flow in linear time.

Our main results are as follows. We first consider the undirected transportation problem where the edge capacities are random variables and there is a bound on the size of the supplies/demands. We also consider the directed transportation problem where the edge capacities are random variables and the supplies/demands satisfy a *realizability* condition (see section 2.3). For both kinds of transportation problems we present linear time algorithms which compute a feasible flow with high probability. We look at the max-flow problem where the edge capacities are i.i.d. random variables. Here we show that the minimum cut is almost surely the cut isolating the sources or the sinks. Again, we present a linear time algorithm that solves such max-flow problem instances, with high probability.

2.3. Realization of 0-1 Matrices

The following notion of realizability of 0-1 matrices will prove useful in the description of our algorithms. Let $R=(r_1, r_2, \dots, r_m)$ and $C=(c_1, c_2, \dots, c_n)$ denote two vectors with non-negative integral entries. The pair (R, C) is said to be *realizable* if and only if there exists a 0-1 $m \times n$ matrix $M=M(R, C)$ with R and C as its row- and column-sum vectors, respectively. The matrix $M(R, C)$ is called a *realization* of the pair (R, C) . The realizability problem is to compute the realization of a pair of row- and column-sum vectors.

The realizability problem is closely related to the capacitated transportation problem. Consider instances of the transportation problem where the underlying bipartite graph is complete with all edges directed from S to T and of capacity one. It is easily seen that finding a feasible flow in such transportation problem instances with the integral supply-demand vectors (A, B) is equivalent to finding the realization of the pair (A, B) . Ryser [23] and Gale [13] gave necessary and sufficient conditions for a pair (R, C) to be realizable. There is a simple greedy algorithm [9, 12, 13] to construct the realization matrix provided it is feasible. This algorithm works in linear time.

Another notion of realizability will also be used in our algorithms. Let R and C be as before. Also, let $\hat{R}=R-D \cdot \bar{1}$ and $\hat{C}=C-D \cdot \bar{1}$, where D is a positive integer and $\bar{1}$ represents the "all ones" vector of the appropriate dimension. The pair (R, C) is said to be D -*realizable* if the pair (\hat{R}, \hat{C}) is *realizable*. The D -*realization* of the pair (R, C) will be denoted by $M_D(R, C)=M(\hat{R}, \hat{C})$.

2.4. A Fine-Tuning Algorithm

The final stage of our mimicking paradigm involves the fine tuning of the solution obtained by the mimicking process. The algorithm described below will be used for this purpose. Consider an instance of the transportation problem with the underlying undirected bipartite graph $G(S \cup T, E)$ satisfying the following conditions.

- (a) $|S| = |T| = n$
- (b) $0 < L(n) \leq a_i, b_j \leq U(n) < (nL(n)/\log n)^{0.5}$
- (c) $\sum_{i \in S} a_i = \sum_{j \in T} b_j$
- (d) for all $i \in S, j \in T$ the supplies and demands a_i, b_j are positive integers
- (e) each edge independently is present with probability $p(n)$
- (f) the capacity of every edge which is present is 1

The algorithm described below uses the Proposal Algorithm [1] (see section 2.2) as a subroutine. It first decomposes the transportation problem instance into $U(n)$ instances of the matching problem for random bipartite graphs. The algorithm constructs $U(n)$ bipartite graphs ensuring that each vertex occurs in exactly as many of these graphs as its supply/demand. The edges from the graph G are equiprobably distributed among the $U(n)$ subgraphs. It is shown that, with high probability, the Proposal Algorithm will find

a perfect matching in each of these subgraphs. The algorithm sends one unit of flow along each edge which is in the union of the perfect matchings found by the Proposal Algorithm. This is the required flow.

$$\text{Let } A_i = \sum_{r=1}^i a_r \text{ and } B_j = \sum_{r=1}^j b_r, \text{ where } A_0 = B_0 = 0.$$

The Fine Tuning Algorithm

STEP (1) Construct U bipartite graphs $G(0), \dots, G(U-1)$ with vertex sets $S(0), \dots, S(U-1)$ and $T(0), \dots, T(U-1)$ as follows. Place each vertex $s_i \in S$ in the vertex sets $S((A_{i-1}+1) \bmod U), \dots, S(A_i \bmod U)$. Similarly, place each vertex $t_j \in T$ in the vertex sets $T((B_{j-1}+1) \bmod U), \dots, T(B_j \bmod U)$.

STEP (2) Colour each edge $(s_i, t_j) \in E$ independently and equiprobably with one of U colours.

STEP (3) For each edge $(s_i, t_j) \in E$, place it in $E(G(k))$ if and only if it has colour k and $s_i \in S(k), t_j \in T(k)$.

STEP (4) Using the Proposal Algorithm [1] find a perfect matching in each of the U subgraphs generated in the previous steps. Let $F \subseteq E$ denote the union of the U perfect matchings.

STEP (5) Saturate (in the forward direction) all the edges in F . This flow will be a feasible solution to the transportation problem instance under consideration.

The following theorem results,

FINE TUNING THEOREM : For all $\alpha > 0$, there exists $\beta > 0$ such that $p(n) > \frac{\beta U^2 \log n}{Ln}$

implies that the Fine-Tuning Algorithm finds a feasible flow in the transportation problem

instance in $O(n^2)$ time with probability $1 - O\left[U \left[\frac{U}{Ln}\right]^\alpha\right]$.

PROOF : The following obvious facts will constitute the proof of this theorem.

Fact 1 : $\lfloor Ln/U \rfloor \leq n_k = |S_k| = |T_k| \leq n, \forall k \in [0, U-1]$.

Fact 2 : $\lim_{n \rightarrow \infty} n_k \rightarrow \infty$.

Fact 3 : Each $s_i(t_j)$ occurs once each in exactly $a_i(b_j)$ of the U subgraphs constructed in step (1).

It is clear that in each of the U subgraphs the edges are present independently of other edges in the same subgraph (although there is dependence between two such subgraphs). Let p_k be the probability that an edge is present in the graph $G(k)$.

Fact 4 : $p_k = p(n)/U(n) \geq (\beta \log n_k)/n_k$.

The Proposal Algorithm will succeed in finding a perfect matching in $G(k)$ in $O(n \log n)$ time with probability $1 - O((U/(Ln))^\alpha)$, when $p_k > \frac{\beta \log n_k}{n_k}$. The relation between α and β is the same as for the Proposal Algorithm. The probability that the

Proposal Algorithm does not succeed on all $G(k)$ is bounded from above by the sum of the probabilities of failure on each of the U matching problems. This completes the proof.

●

We make two remarks about the generality of this algorithm. First, note that the algorithm uses each edge in the forward direction only. The algorithm would work equally well in the case where the edges are directed, though it would not send any flow along edges which are directed from T to S . Also, the algorithm does not actually require that $|S| = |T| = n$. A careful examination of the proof shows that it would be sufficient to have $(U^2 \log n)/(Lp(n)) \leq |S|, |T| \leq n$, where n now is a measure of the size of the problem. We need to impose these bounds on the sizes of S and T to ensure that n_k (in Fact 1) cannot be too small or too large.

3. The Undirected Transportation Problem

We now present an algorithm to solve certain instances of the transportation problem where the underlying graph is undirected. Let $D(m)$ be such that $D(m) < m/2(\log m)^{0.5}$. Consider an instance of the transportation problem which satisfies the following conditions (Y).

(Y.1) the underlying graph $G(S \cup T, E)$ is undirected and $|S| = |T| = m$

(Y.2) $\forall u \in S, a_u < p \cdot D(m)$ and $\forall v \in T, b_v < p \cdot D(m)$

(Y.3) $\sum_{u \in S} a_u = \sum_{v \in T} b_v$, and all supplies and demands are non-negative integers

(Y.4) the probability that an edge from $S \times T$ is present is p (a constant)

(Y.5) if an edge (u, v) is present then its capacity $c(u, v) = 1$

We will show that the transportation problem instances satisfying conditions (Y) are feasible with high probability. This will be done constructively by specifying a linear time algorithm which succeeds in finding the feasible flow with high probability. The following theorem results.

UNDIRECTED TRANSPORTATION THEOREM : The Undirected Transportation Algorithm finds a feasible flow for transportation problem instances satisfying (Y) in linear time with probability $1 - O(n^{-\gamma})$, where $\gamma > 0$ is a constant which depends on p .

The expected value of the capacity of any edge is p . Consider the deterministic relaxation of the above problem. It would correspond to finding a feasible flow for the supplies a_i and b_j in the case where the underlying undirected graph is complete with all edges having capacity p . It is not very hard to show that the deterministic problem has a feasible solution. However, we will need to find a feasible flow using the edges in the forward direction only. In this directed case, the problem need not have a feasible solution at all, e.g. consider the case where $a_1 = b_1 = 2$ and all other supplies and demands are 0. To get around this problem we will add a large number $E(m)$ to each supply and demand. Now we will be able to find a feasible flow for the deterministic relaxation using edges in the forward direction only. It has been observed [17] that a transportation problem

instance is feasible provided the values of the supplies and demands are sufficiently uniform. Uniformity of the supplies and demands requires that the supply (demand) at a source (sink) is in proportion to the net capacity of the edges incident at that source (sink). In a sense, adding $E(m)$ to each supply and demand corresponds to making their values more uniform, since the expected total capacity of the edges incident at each vertex is equal.

The Undirected Transportation Algorithm is based on the mimicking paradigm. The first stage of this algorithm constructs a solution to the deterministic relaxation of the original problem. To simplify the algorithm, we will scale up the deterministic relaxation by a factor of $c=1/p$. This corresponds to multiplying all capacities, supplies and demands by a factor of c . Let us choose $E(m)=m/2\log m$. Define $\hat{a}_u=ca_u+E(m)$ for all $u \in S$, and $\hat{b}_v=cb_v+E(m)$ for all $v \in T$. Also, let \hat{a} be the s -dimensional vector (\hat{a}_u) and \hat{b} the t -dimensional vector (\hat{b}_v) . The deterministic problem is an undirected transportation problem with (\hat{a},\hat{b}) as the supply-demand vectors and all edges of capacity $c \cdot p=1$. We will use all edges in the forward direction only and so the deterministic problem is exactly that of finding a realization of (\hat{a},\hat{b}) .

The first stage of the algorithm constructs a realization of the pair (\hat{a},\hat{b}) , viz. the $m \times m$ 0-1 matrix $M(\hat{a},\hat{b})$. This corresponds to the deterministic solution of the transportation problem instance with all random variables (in this case, the edge capacities) replaced by their expected values. In the second step we mimic the solution of the deterministic relaxation by saturating the edges (if present) which correspond to the non-zero entries in $M(\hat{a},\hat{b})$. The flow is sent in the forward direction, i.e. from S to T . Finally, using the edges in the backward direction we fine-tune the solution to obtain a feasible flow for the transportation problem instance under consideration.

The Undirected Transportation Algorithm

STEP (1) (*Deterministic Relaxation and Solution*) Construct the $s \times t$ 0-1 matrix $M(\hat{a},\hat{b})$. The pair (\hat{a},\hat{b}) must have only integral entries for the realization to be feasible. This can be ensured by rounding (up or down) each \hat{a}_u, \hat{b}_v while still satisfying the condition $\sum_{u \in S} \hat{a}_u = \sum_{v \in T} \hat{b}_v$.

STEP (2) (*Mimicking Process*) Saturate (in the forward direction) all existing edges which correspond to the 1's in $M(\hat{a},\hat{b})$. This yields a $s \times t$ 0-1 flow matrix N . Let the row- and column-sum vectors of this matrix be \bar{a} and \bar{b} , respectively.

STEP (3) (*Fine Tuning*) At this point each $u \in S$ has sent an excess of $\bar{a}_u - a_u$ units of flow, while each $v \in T$ has received an excess of $\bar{b}_v - b_v$ units of flow. Using the Fine-Tuning algorithm and all edges in the backward direction, send out $\bar{b}_v - b_v$ units of flow from each $v \in T$. Ensure that each $u \in S$ receives exactly $\bar{a}_u - a_u$ units of flow.

STEP (4) Combine the flows constructed in the two previous steps to obtain the desired feasible flow.

3.1. Analysis of the Undirected Transportation Algorithm

We will prove the Undirected Transportation Theorem with the help of the lemmata presented below. Lemma 1 establishes that the first step of the algorithm will succeed by showing that the pair (\hat{a}, \hat{b}) is realizable. The realization algorithm will construct the matrix $M(\hat{a}, \hat{b})$ in linear time given its feasibility.

LEMMA 1 : The pair (\hat{a}, \hat{b}) is realizable. -

PROOF : Consider a transportation problem instance I with $|S| = |T| = m$. Assume that all edges from $S \times T$ are present. Let each edge have capacity 1 and be directed from S to T . Let \hat{a}_u be the supply at $u \in S$, and \hat{b}_v the demand at $v \in T$. The Integrality Theorem for flow problems can also be applied to this transportation problem instance. It implies that, in a feasible transportation problem instance, if all supplies, demands and capacities are integral then there is an integral feasible flow. The rounding process in STEP (1) ensures that the supplies and demands are integral, while all capacities are 1 in I . Thus, we have that I is feasible if and only if the pair (\hat{a}, \hat{b}) is realizable. We invoke the Transportation Theorem to establish the feasibility of I .

By the Transportation Theorem, I is feasible if and only if $\forall X \subseteq S$

$$\sum_{u \in X} \hat{a}_u \leq \sum_{v \in T} \min(\hat{b}_v, c(X, \{v\})) \quad (3.1)$$

In this case, $c(X, \{v\}) = x$ for all $v \in T$, where $x = |X|$. Note that the following bounds hold for \hat{a} and \hat{b} ,

$$\begin{aligned} E(m) \leq \hat{a}_u \leq E(m) + D(m), & \quad \forall u \in S \\ E(m) \leq \hat{b}_v \leq E(m) + D(m), & \quad \forall v \in T \end{aligned}$$

We now perform a case analysis on the value of x to establish the validity of (3.1).

Case 1 [$x \leq E(m)$] :

In this case (1) is equivalent to the following,

$$\sum_{u \in X} \hat{a}_u \leq xm$$

Since $m > E(m) + D(m)$ (for large m) the above inequality, and hence (3.1), is valid.

Case 2 [$x > E(m)$] :

In this case the inequality (3.1) is implied by the following,

$$\sum_{u \in X} \hat{a}_u \leq \sum_{v \in Y} \hat{b}_v + x(m - y), \quad \forall Y \subseteq T, y = |Y| \quad (3.2)$$

Using the bounds on \hat{a} , \hat{b} and the fact that $\sum_{u \in S} \hat{a}_u = \sum_{v \in T} \hat{b}_v$, it is clear that (3.2) is valid if the following inequality holds.

$$(m - x)E(m) + (m - y)(x - E(m) - D(m)) \geq 0, \quad \forall y \leq m \quad (3.3)$$

Again, using the bounds on \hat{a} and $\hat{\delta}$, we see that (3.2) holds if $y \leq m - E(m) - D(m)$. Moreover, (3.3) is valid if $x > E(m) + D(m)$. The only case left is where $y > m - E(m) - D(m)$ and $E(m) \leq x \leq E(m) + D(m)$. But for this range of values of x and y we see that (3.3) is valid provided $m \geq D(m) + E(m)$ and $D(m)^2 \leq E(m)(m - E(m))$. Given our choice of $E(m)$ and the bound on $D(m)$ this is clearly true. ●

We now turn to STEP (2) of the algorithm. The next lemma bounds the values of the excess flow sent by the sources and the excess flow received by the sinks. Let $F(m) = D(m) + E(m)$.

LEMMA 2 : In the Undirected Transportation Algorithm

$$\forall \gamma > 0 \exists l > 0 \text{ s.t. } \text{Prob}[\exists u \in S : |\bar{a}_u - p \cdot \hat{a}_u| \geq l(F(m)\log F(m))^{0.5}] = O(m^{-\gamma})$$

$$\forall \gamma > 0 \exists l > 0 \text{ s.t. } \text{Prob}[\exists v \in T : |\bar{b}_v - p \cdot \hat{\delta}_v| \geq l(F(m)\log F(m))^{0.5}] = O(m^{-\gamma})$$

PROOF : Let X be the sum of r Bernoulli trials, where each trial assumes the value 1 with probability p and value 0 with probability $1-p$. The Chernoff bound [3] as applied to the tail of binomial distribution states that for any β , $0 \leq \beta \leq 1$,

$$\text{Prob}[X_v \geq (1+\beta)rp] \leq \exp(-\beta^2 rp / 2)$$

The number of non-zero entries in each row/column of the matrix $M(\hat{a}, \hat{\delta})$ is large, at least $E(m)$. An application of the Chernoff bound for the tail of the binomial distribution completes the proof of the two lemma. ●

At this stage the following bounds hold (with high probability).

$$E(m) - l(F(m)\log F(m))^{0.5} \leq \bar{a}_u - a_u \leq E(m) + l(F(m)\log F(m))^{0.5}, \forall u \in S \quad (3.4)$$

$$E(m) - l(F(m)\log F(m))^{0.5} \leq \bar{b}_v - b_v \leq E(m) + l(F(m)\log F(m))^{0.5}, \forall v \in T \quad (3.5)$$

STEP (3) of the algorithm uses the edges (in the backward direction) to route the excess flow back from sinks to sources. It is clear that a_u , \hat{a}_u , b_v and $\hat{\delta}_v$ are all integral as required for the application of the Fine-Tuning Algorithm. The bounds on the excess, (3.4) and (3.5), together with the Fine-Tuning Theorem yields the following lemma.

LEMMA 3 : The Fine-Tuning Algorithm succeeds in Step (3) with probability $1 - O(m^{-\alpha})$, where $\alpha > 0$ is a constant which depends on p .

The probability of failure of each stage of the algorithm is now suitably bounded from above. We now observe that the probability of failure of the entire algorithm is bounded from above by the sum of the probabilities of failure of the various steps in the algorithm (even though they may not be independent). This completes the proof of the Undirected Transportation Theorem. Note that we cannot choose $E(m)$ to be larger than $m/\log m$ if we wish to employ the Fine-Tuning Theorem as above. The proof of Lemma 1 requires that $E(m)$ should not be much smaller than $D(m)$. This imposes an upper bound of $m/(\log m)^{0.5}$ on $D(m)$.

We observe that the Undirected Transportation Algorithm can also be applied to certain cases where $|S| \neq |T|$. The proof of the Undirected Transportation Theorem is easily seen to extend to the case where $|S| \neq |T|$ provided $|S|, |T|$ are large enough. More precisely, we require that $|S|, |T|$ should be substantially larger than $D(m)$ and $E(m)$. Note that to be able to apply the Fine-Tuning Algorithm in the case where $|S| \neq |T|$, we again need that $|S|$ and $|T|$ should be substantially larger than $E(m)$.

3.2. The Undirected Max-Flow Problem

We now make use of the Undirected Transportation Theorem to devise an algorithm to solve a probabilistic version of the undirected max-flow problem. Consider an instance of the max-flow problem satisfying the conditions (Z).

- (Z.1) the underlying graph is undirected
- (Z.2) $|S| = |T| = r$ and $|I| = n$, where $r \leq n$
- (Z.3) each edge is present with a probability p (a constant)
- (Z.4) if an edge (u, v) is present then its capacity $c(u, v) = 1$

We assume, without loss of generality, that the $(S, V-S)$ cut has a smaller capacity than the $(V-T, T)$ cut. If this is not the case then we can interchange the roles of the source and sink vertices and reverse the direction of each edge. We present the following theorem,

MIN-CUT THEOREM : In instances of the probabilistic max-flow problem satisfying conditions (Z),

$$\text{Prob}[\text{The } (S, V-S) \text{ cut is the minimum cut}] = 1 - O(n^{-\alpha})$$

where $\alpha > 0$ is a constant which depends on p .

This theorem is proved by presenting a linear time algorithm which constructs a flow saturating the $(S, V-S)$ cut. Clearly, the value of the max-flow is less than or equal to the capacity of this cut. This establishes the validity of the above theorem.

The Max-Flow Algorithm works in two stages. In the first stage an instance of the probabilistic transportation problem is created, such that a feasible solution to that instance will yield a flow saturating the $(S, V-S)$ cut in the original problem. In the second stage it uses the Undirected Transportation Algorithm to solve the transportation problem instance generated in the first stage.

3.3. The Max-Flow Algorithm

Our aim is to find a flow which will saturate the $(S, V-S)$ cut. Clearly, we can ignore edges drawn from $S \times S$ and $T \times T$. Moreover, all edges from $S \times T$ can be saturated (in the forward direction) without affecting the flow through the remaining edges. We now consider only the edges from $S \times I, I \times I$ and $I \times T$. The algorithm will find a flow which will saturate all edges from $S \times I$ in the forward direction.

For each $v \in I$ define $\Delta_v = c(S, v) - c(v, T)$, this is the excess of the capacity of incoming edges over the capacity of the outgoing edges.. Let $I^+ = \{v \in I : \Delta_v > 0\}$ and $I^- = \{v \in I : \Delta_v < 0\}$. Further define $\Delta^+ = \sum_{v \in I^+} \Delta_v$ and similarly Δ^- . We now describe the first stage of the Max-Flow Algorithm.

STEP (1) Saturate all edges from $S \times I$ by sending flow from sources to intermediate nodes.

STEP (2) Saturate all edges from $I \times T$ by sending flow from intermediate nodes to the sinks.

STEP (3) Since $c(S, I) \leq c(I, T)$ we have $\Delta^+ + \Delta^- \leq 0$. At this point, for each $v \in I$, the excess of in-flow over out-flow is exactly Δ_v . Arbitrarily reduce flow along edges drawn from $I^- \times T$ until the net flow across the (S, I) cut equals the net flow across the (I, T) cut. Ensure that, for each $v \in I^-$, the new excess flow, say δ_v , remains non-positive. Let $\delta^- = \sum_{v \in I^-} \delta_v$, then we have $\Delta^+ + \delta^- = 0$ and $\Delta_v \leq \delta_v \leq 0$, for each $v \in I^-$.

STEP (4) Let G^1 be the bipartite subgraph of the original underlying graph G which is induced by the vertex set $S^1 = I^+$ and $T^1 = I^-$. Construct an instance of the CTP with G^1 as the underlying graph, supply Δ_v for $v \in S^1$ and demand $-\delta_v$ for $v \in T^1$.

It is easy to see that the first stage works in linear time. In the second stage of the Max-Flow Algorithm we find a feasible flow for this transportation problem instance using the linear-time Undirected Transportation Algorithm described earlier. The final flow is the sum of the two flows constructed in the two stages of the max-flow algorithm. The following theorem results.

MAX-FLOW THEOREM : The Max-Flow Algorithm finds a maximum flow (for instances satisfying conditions (Z)) in linear time with probability $1 - O(n^{-\alpha})$, where $\alpha > 0$ is a constant which depends on p .

Observe that the above algorithm will find a flow saturating the $(S, V-S)$ cut if the second stage succeeds. This flow will be a maximum flow. To complete the proof of this theorem we need to show that the transportation problem instance generated by this algorithm satisfies the conditions of the Undirected Transportation Theorem. We first prove the following bound on the size of the excess at each intermediate node. This bound also applies to the size of the supplies and demands of the transportation problem instance generated in the first stage. We also show that $|I^+|$ and $|I^-|$ are both fairly close to $n/2$.

LEMMA 4 : In the Max-Flow Algorithm

$$\forall \gamma > 0 \exists k > 0 \text{ s.t. } \text{Prob}[\exists v \in I : |\Delta_v| > k(n \log n)^{0.5}] = O(n^{-\gamma}).$$

PROOF : The result holds for the case where $r \leq n^{0.5}$ since the maximum value of Δ_v is then bounded by $n^{0.5}$. We now consider the case where $r > n^{0.5}$.

Consider some $v \in I$. Let X_v denote the number of edges from $S \times \{v\}$ which are actually present in the underlying graph G . Similarly, let Y_v denote the number of edges from $\{v\} \times T$. Clearly, both X_v and Y_v are the sum of r independent Bernoulli trials where each trial assumes value 1 with probability p and value 0 with probability $1-p$. We make use of the Chernoff bound [3] as applied to the tails of binomial distributions. For any β , $0 \leq \beta \leq 1$, we have,

$$\text{Prob}[X_v \geq (1+\beta)rp] \leq \exp(-\beta^2 rp/2)$$

Choosing $\beta = ((b \log r)/r)^{0.5}$, for some positive constant b , and applying the bound to both X_v and Y_v we have,

$$\text{Prob}[|\Delta_v| = |X_v - Y_v| \geq 2\beta rp] \leq 4r^{-b/2}$$

Summing over all $v \in I$ and choosing $b = 2(1+2\gamma)/p$ and $k = (8p(1+2\gamma))^{0.5}$ we have,

$$\text{Prob}[\exists v \in I: |\Delta_v| \geq k(r \log r)^{0.5}] = O(r^{-2\gamma})$$

Since $n^{0.5} \leq r \leq n$ we have the desired result. ●

LEMMA 5 : In the Max-Flow Algorithm

$$\forall \gamma > 0 \exists f > 0 \text{ s.t. } \text{Prob}[|I^+| - n/2| > f(n \log n)^{0.5}] = O(n^{-\gamma}).$$

PROOF : In our definition of I^+ and I^- we ignored the vertices $v \in I$ for which $\Delta_v = 0$. Since these vertices can be assigned to either set, they can be used to balance the sizes of I^+ and I^- . To simplify the following description we will assume that such vertices will be assigned to either I^+ or I^- equiprobably.

By symmetry, $\text{Prob}[v \in I \text{ is assigned to } I^+] = 1/2$. Independence follows from the observation that all edge capacities are independently distributed. An application of the Chernoff bound yields the desired result. ●

It is clear that the transportation problem instances generated in the first stage satisfies all requirements of the Undirected Transportation Theorem, with one exception. The number of sources (I^+) and sinks (I^-) in the transportation problem instance will not be equal. However, as we remarked earlier, the Undirected Transportation Algorithm can be still be used provided the number of source and sinks is large enough. The bound from Lemma 5 shows that this is indeed the case.

4. A Directed Transportation Algorithm

We now present an algorithm to solve certain instances of the transportation problem where the underlying graph is directed. In particular, we consider instances of the transportation problem satisfying the following conditions (X).

(X.1) the underlying graph $G(S \cup T, E)$ is complete and every edge is directed from S to T

(X.2) $|S| = |T| = n$

(X.3) the edge capacities are i.i.d. random variables drawn from the set $\{0,1, \dots, K\}$, where $K > 1$ is some constant

(X.4) the expected value of the edge capacities is at least $1+\epsilon$, where ϵ is a positive constant

(X.5) the pair (a,b) is $(D+1)$ -realizable, for some constant $D > 0$

We present an algorithm, the Directed Transportation Algorithm, which will solve such instances of the transportation problem with high probability. This leads to the following theorem,

DIRECTED TRANSPORTATION THEOREM : The Directed Transportation Algorithm finds a feasible flow for transportation problem instances satisfying (X) in linear time with probability $1-O(n^{-\gamma})$, $\gamma > 0$.

Before we describe the algorithm and prove the Directed Transportation Theorem we present two combinatorial theorems which are useful in the analysis of the Directed Transportation Algorithm.

4.1. A Combinatorial Process

Consider the following combinatorial process. The *state* of the process is an arbitrary placement of n particles, call them $P = \{1,2,3,\dots,n\}$, on integer points of the real line. There may be more than one particle at a given position. The *initial state* has all n particles at the origin. A *state transition* is divided into two distinct steps. Let S be a subset of P such that $|S| = 2k$. The first step in a transition moves every particle in S , for some arbitrary S , one position in the negative direction (say to the left). In the second step, the k leftmost particles are each moved *two* positions to the right (or in the positive direction). It can be shown that no particle will ever move out of the interval $[-A \log n, 2]$, where A is some positive constant.

Let K be a positive integer and θ a positive real number. Consider now the following generalization of the combinatorial process. The definition of the state (as well as the initial state) of the process is as before. The first step of a transition, as before, moves all particles in S , for some arbitrary S , one step each to the left. The second step of the transition involves the choice of n arbitrary integers, $\{d_1, d_2, \dots, d_n\}$, such that $\sum_{j=1}^n d_j = |S|$ and $0 \leq d_i \leq K$ for each i . We will refer to the requirement that $\sum_{j=1}^n d_j = |S|$ as the *balance* constraint. It is also required that $\sum_{j=1}^t d_j \geq (1+\theta)t$ for all $t < t_0$, where t_0 is the index of the rightmost non-zero d_i . This last condition will be referred to as the *prefix* constraint. The second step of a transition now moves the t^{th} -leftmost particle d_t positions to the right, for $t = 1, 2, \dots, t_0$.

The conditions imposed on the second step of a transition ensure that the net rightward movement of a group of t leftmost particles is larger than t . This constraint prevents any particle from straying too far away to the left. The following theorem can

be proved about the generalized process.

INTERVAL THEOREM : All particles remain in the interval $[-C \log n, D]$ of the real line, where $C=C(K, \theta)$ and $D=D(K, \theta)$ are positive constants.

The following notation and lemma will be required for the proof of this theorem. Let $p_i(\tau)$ denote the location of the i^{th} particle after τ transitions have taken place. The state of the process at time step τ will be given by the set of locations, $p_i(\tau)$, occupied by the particles i , $1 \leq i \leq n$. The main tool for the analysis of this combinatorial process will be the following notion of the *moment* at an integer point on the real line.

DEFINITION : Let p be any integer point on the real line. The *left moment* at p at any time is the sum of the distances from p of all particles to the left of p on the real line.

$$LM(p, \tau) = \sum_i \max(p - p_i(\tau), 0)$$

In the following lemma, we will show that the moment satisfies an invariant inequality at each time step. Using this invariant, we will be able to establish that no particle can move too far away from the origin.

LEMMA 6 : There exists A and ρ depending only on θ and K such that,

- (a) $A > 0$ and $0 < \rho < 1$, and
- (b) at each time step τ and for each integer l , $LM(l, \tau) \leq nA\rho^{-l}$.

PROOF : The proof will be by induction on the time step τ . We will assume that the lemma holds for all l at the time step $\tau-1$, and prove it for all l at time step τ . We will start by proving the induction step, and later show the base case (i.e. at time step 0). The value of the constants ρ and A will also be specified later.

Assume that the left moment at every integer position satisfies the required inequality at time step $\tau-1$. We now show that it must satisfy the required inequality after the completion of the τ^{th} transition (henceforth referred to as the current transition). Observe that it suffices to prove the invariant for a particular position l on the real line, without any loss of generality. Therefore, we are only required to show that $LM(l, \tau) \leq nA\rho^{-l}$. The main idea of the proof is to bound the new moment at l by a positive linear combination of the moments at the previous time step. To simplify our notation we will consider everything relative to this location l . We will use the following notation.

NOTATION :

- Let *cell* i denote the location $l-i$ on the real line, where $i \leq K-1$.
- Let *cell* K denote all locations to the left of $l-(K-1)$.
- Let A_i denote the number of particles in *cell* i , before the current transition.
- Let a_i denote the number of particles which were moved one unit to the left from *cell* i in the first stage of the current transition.
- Let β_i denote the sum of the distances that particles from *cell* i were moved to the right in the second stage of the current transition. (Thus, $\beta_i = \sum_{j \in S_i} d_j$, where S_i is the

set of particles that resided in cell i after the first stage of the current transition.)

• Let $M(h)$ denote the moment at cell h before the current transition and $M'(h)$ denote the moment at cell h after the current transition. (By definition, $M(h) = LM(l-h, r-1)$ and we have to prove that $M'(0) \leq nA\rho^{-l}$.)

In general, we will only be interested in particles which lie at cell 0 or to its left. Whenever we refer to a particle at cell i it will be assumed that $0 \leq i \leq K$, unless otherwise stated.

The moment at cell h , $-1 \leq h \leq K$, before the current transition is given by the following equation.

$$M(h) = M(K) + \sum_{i=h+1}^K (i-h)A_i \leq nA\rho^{-l+h} \quad (4.1)$$

Let L denote the moment gained at cell 0 during the first stage of the current transition. Also, let R denote the moment lost at cell 0 during the second stage of the current transition. Thus, we have that,

$$M'(0) = M(0) + L - R \quad (4.2)$$

Observe that each particle to the left of the cell 0 which was moved a unit to the left (in the first stage of the current transition) will contribute to the increase in the moment at cell 0. It is now easy to see that L is given by the following equation.

$$L = \sum_{i=0}^K a_i \quad (4.3)$$

Consider now a particle which was moved c positions to the right in the second stage of the current transition. It is possible that this particle ended up at a position to the right of cell 0. In that case, its contribution to the decrease in the moment at cell 0 would be less than c . This complicates the computation of the value of R . However, we do know that a particle which was moved from cell i to a position to the right of cell 0 will cause a decrease of i in the moment at cell 0. Further, it is known that a particle can move at most K positions to the right in a single transition. This implies that the particles in cell i (before the second stage of the current transition) must cause a decrease of at least $\frac{i}{K}\beta_i$ in the moment at cell 0. Thus, we have the following lower bound on the value of R .

$$R \geq \frac{1}{K} \sum_{i=1}^K i\beta_i \quad (4.4)$$

It will be convenient to express the above lower bound on R in terms of A_i and a_i . This can be done as follows. Consider the rightmost particle which was moved to the right in the second stage of the current transition. Let t denote the cell to which this particle belonged at the end of the first stage of the current transition. If cell t was to the right of cell 1 then set $t=1$. It is clear that if $t > 1$ then $\beta_1 = \beta_2 = \dots = \beta_{t-1} = 0$. The

balance constraint on this combinatorial process requires that the net leftward movement in the first stage of any transition be exactly equal to the net rightward movement in the second stage of that transition. Therefore, we have the following inequality.

$$\sum_{i=t}^K \beta_i \geq \sum_{i=0}^K a_i \quad (4.5)$$

Note that this inequality need not be tight since there may be particles to the right of cell 0 which were moved to the left in the first stage of the current transition. It is not very hard to see that the inequality also holds in the case where particles to the right of cell 0 were moved rightwards in the second stage.

The number of particles at cell i and to its left after the first stage is $a_{i-1} + \sum_{j=i}^K A_j$.

We invoke the *prefix* constraint on the transitions to derive the following inequality for $t+1 \leq i \leq K$.

$$\sum_{j=i}^K \beta_j \geq (1+\theta)(a_{i-1} + \sum_{j=i}^K A_j) \quad (4.6)$$

We now use equations (4.4), (4.5) and (4.6) to derive the following lower bound on R , after some algebraic manipulation.

$$R \geq \frac{t}{K} \sum_{i=0}^K a_i + \frac{1+\theta}{K} \sum_{i=t}^{K-1} a_i + \frac{1+\theta}{K} \sum_{i=t+1}^K (i-t)A_i \quad (4.7)$$

We are now in a position to give an upper bound on the value of $M'(0)$. Substituting equations (4.1), (4.3) and (4.7) into (4.2) we obtain the following inequality.

$$M'(0) \leq M + \sum_{i=1}^K iA_i + \left[1 - \frac{t}{K}\right] \sum_{i=0}^K a_i - \frac{1+\theta}{K} \sum_{i=t}^{K-1} a_i - \frac{1+\theta}{K} \sum_{i=t+1}^K (i-t)A_i \quad (4.8)$$

In order to show that this is bounded by $nA\rho^{-t}$ we have to consider three different cases.

Case I [$t=K$]: In this case the RHS of inequality (4.8) turns out to be exactly $M(0)$ and, thus, by the induction hypothesis, $M'(0) \leq M(0) \leq nA\rho^{-t}$.

Case II [$t=K-1$]: In this case inequality (4.8) can be seen to imply the following, using the fact that $0 \leq a_i \leq A_i$.

$$M'(0) \leq M(K) + \sum_{i=1}^K iA_i + \frac{1}{K} \sum_{i=0}^{K-2} A_i - \frac{\theta}{K} A_K$$

At this point we invoke equation (4.1) and make use of the induction hypothesis to obtain the following bound.

$$M'(0) \leq nA\rho^{-t-1} \left[\left[1 - \frac{1}{K}\right] + \left[\frac{1-\theta}{K}\right] \rho^K + \frac{\theta}{K} \rho^{K+1} \right] = nA\rho^{-t-1} f(\rho)$$

The last inequality completes the analysis of Case II provided $f(\rho) \leq \rho$. It can be shown that this is indeed the case provided ρ is sufficiently close to 1. To see this, it is enough to verify that the function $g(\rho)$,

$$g(\rho) = f(\rho) - \rho = \left[1 - \frac{1}{K}\right] + \left[\frac{1-\theta}{K}\right]\rho^K + \frac{\theta}{K}\rho^{K+1} - \rho$$

is negative when $\rho < 1$, for ρ close enough to 1. This can be verified by observing that g is continuous at 1, $g(1) = 0$, and $g'(1)$ is positive.

Case III [$1 \leq t \leq K-2$]: In this case inequality (4.8) can be seen to imply the following, using the fact that $a_i \leq A_i$.

$$M'(0) \leq \left[1 - \frac{1}{K}\right] \left[M(K) + \sum_{i=0}^K [i+1] A_i \right] + \left[\frac{1}{K} - \frac{\theta}{K} [K-t] \right] \left[M(K) + A_K \right] + \frac{\theta}{K} [K-t] M(K)$$

At this point we invoke equation (4.1) and make use of the induction hypothesis to obtain the following bound.

$$M'(0) \leq nA\rho^{-l-1} \left[\left[1 - \frac{1}{K}\right] + \left[\frac{1}{K} - \frac{\theta}{K} [K-t] \right] \rho^K + \frac{\theta}{K} [K-t] \rho^{K+1} \right] \leq nA\rho^{-l-1} \rho = nA\rho^{-l}$$

The last inequality also holds when $\rho < 1$, for ρ close enough to 1. This can be verified in the same manner as in case II.

This concludes the proof of the induction step. Note that ρ is chosen to lie close enough to 1, so as to satisfy all the inequalities derived above. Similarly, A will be chosen such that the base case is satisfied.

Base Case: In the initial state all particles were at the origin. This implies that, at time step 0, the left moment at all integer points $p \leq 0$ were 0, thus trivially satisfying the invariant. The left moment at location p , for $p > 0$, is simply pn in the initial state. We choose A such that, for all $p > 0$, it is the case that $np \leq nA\rho^{-p}$. Clearly, it must be the case that $A \geq p\rho^p$. Such an A exists since the function $h(p) = p\rho^p$ is bounded for $p > 0$. ●

We are now ready to prove the Interval Theorem.

PROOF (Interval Theorem): We first show that no particle can move too far to the left. This may be verified by considering the leftmost point on the real line which has a non-zero moment. Let $p_l < 0$ be the leftmost point on the real line at which a particle may be placed in this combinatorial process. A particle in location p_l contributes 1 to the left moment at $p_l + 1$. Thus, we have the following inequalities.

$$1 \leq LM(p_l + 1) \leq nA\rho^{-p_l - 1}$$

Therefore, it must be the case that $p_l \geq -C \log n$ where $C(K, \theta)$ is a positive constant.

Now we show that a particle cannot move too far to the right. Let $p_r > 0$ be the rightmost position occupied by a particle during this combinatorial process. Consider the

first time a particle is moved onto the position p_r . Clearly, this particle must have previously occupied a position at or to the right of $p_r - K$. The *prefix* condition requires that there be at most $\frac{n}{1+\theta}$ particles to the left of a particle which is moved to the right. This implies that there must be at least $\theta n / (1+\theta)$ particles at or to the right of the position $p_r - K$. Therefore, there must be a right moment of at least $(p_r - K)\theta n / (1+\theta)$ about the origin. It is easy to see that, due to the *balance* constraint, the right moment must be equal to the left moment at the origin at all times. The left moment at the origin is always less than nA . Thus, we get the following inequality.

$$(p_r - K) \frac{\theta n}{1+\theta} \leq nA$$

This means that $p_r \leq D$, where $D(K, \theta)$ is an appropriately chosen constant. ●

4.2. A Generalization of the Mendelsohn-Dulmage Theorem

Mendelsohn and Dulmage [21] proved the following theorem about bipartite graphs.

MENDELSON–DULMAGE THEOREM : Let $G(S \cup T, E)$ be a bipartite graph and let M_1 and M_2 be two matchings in G . Then there exists a matching $M_3 \subseteq M_1 \cup M_2$, such that M_3 covers all the nodes of S covered by M_1 and all the nodes of T covered by M_2 .

The proof of this theorem is constructive and leads to an algorithm which runs in time $O(|S| + |T|)$. We generalize this theorem as follows. Let K be a fixed positive integer. By a K -matrix we mean a $m \times n$ matrix whose entries are drawn from the set $\{0, 1, \dots, K\}$.

GENERALIZED MENDELSON–DULMAGE THEOREM : Let M^1 be a K -matrix whose row sums, a_i^1 , and column sums, b_j^1 , satisfy the following condition,

$$l_i^1 \leq a_i^1 \text{ and } b_j^1 \leq u_j^1$$

Similarly, let M^2 be a K -matrix whose row sums, a_i^2 , and column sums, b_j^2 , satisfy the following condition,

$$a_i^2 \leq u_i^2 \text{ and } l_j^2 \leq b_j^2$$

Then there exists a K -matrix M^3 such that

(a) the row and column sums of M^3 satisfy the following conditions

$$l_i^3 \leq a_i^3 \leq u_i^3$$

$$l_j^3 \leq b_j^3 \leq u_j^3$$

(b) for each i, j we have $M_{ij}^3 \leq \max(M_{ij}^1, M_{ij}^2)$

Moreover, a matrix M^3 satisfying the above conditions can be constructed in $O(Kmn)$ time.

PROOF : First, observe that the only case we need to consider is where for each i, j we have $\min(M_{ij}^1, M_{ij}^2) = 0$. If this is not the case then let D be the K -matrix where $\forall i, j D_{ij} = \min(M_{ij}^1, M_{ij}^2)$. Subtract D from both M^1 and M^2 . Also, modify the upper and lower bounds on the rows (and columns) of M^1, M^2 by subtracting the row-sums (and column-sums) of the matrix D . Clearly, a solution to the new problem, when added to D , gives a solution to the original problem involving M^1 and M^2 .

We now restrict our attention to the case where $\forall i, j \min(M_{ij}^1, M_{ij}^2) = 0$. Let $a_i = \max(a_i^1, a_i^2)$ and $b_j = \max(b_j^1, b_j^2)$, for all i, j . Define a vertex set S such that for each row i there are a_i vertices $s(i, 1), s(i, 2) \dots s(i, a_i)$. Similarly, define the vertex set T such that for each column j there are b_j vertices $t(j, 1), t(j, 2) \dots t(j, b_j)$. We will construct matchings X^1 and X^2 on the bipartite vertex set $S \cup T$ corresponding to the two matrices M^1 and M^2 , respectively. We describe the construction of the matching X^1 only. The other matching, X^2 , can be constructed analogously.

For each non-zero entry M_{ij}^1 , introduce M_{ij}^1 edges into X^1 connecting vertices corresponding to row i with vertices corresponding to column j . It is easy to ensure that each vertex in $S \cup T$ has at most one edge incident on it. This process yields a matching of cardinality $\sum_{i=1}^m a_i^1 = \sum_{j=1}^n b_j^1$.

We now invoke the Mendelsohn-Dulmage Theorem to construct a third matching, $X^3 \subseteq X^1 \cup X^2$, which covers all the vertices in S covered by X^1 and those in T covered by X^2 . To derive the matrix M^3 from the matching X^3 , set the value of M_{ij}^3 to be the number of edges in M^3 which connect vertices corresponding to row i with those corresponding to column j .

Clearly, the row-sums of M^3 respect the upper bounds since the number of vertices corresponding to each row do so. The row-sums of M^3 are seen to respect the lower bounds since the matching X^3 covers at least as many vertices of each row as the row-sums of M^1 . Similar reasoning shows that the column-sums of M^3 respect both the upper and lower bounds. This establishes condition (a) of the theorem.

To verify the validity of condition (b), recall that $\min(M_{ij}^1, M_{ij}^2) = 0$. This implies that the number of edges in $X^1 \cup X^2$ connecting vertices corresponding to a row i with the vertices corresponding to a column j is less than $\max(M_{ij}^1, M_{ij}^2)$. Thus, the value of M_{ij}^3 cannot exceed $\max(M_{ij}^1, M_{ij}^2)$.

Finally, note that the entire proof is constructive. Moreover, the construction described above requires only $O(Kmn)$ time. This concludes the proof. ●

4.3. The Directed Transportation Algorithm

The key idea behind the Directed Transportation Algorithm can also be formulated in terms of the mimicking paradigm. We first set aside a small fraction of the edge capacities for the purposes of the Fine-Tuning Algorithm. Next, we construct the $(D+1)$ -

realization of the supply/demand vectors. This corresponds to the solution of the deterministic relaxation of the probabilistic transportation problem. The solution to the relaxed problem is then mimicked to obtain a partial solution to the original problem. The mimicking process is considerably more sophisticated than that used for the transportation problem. Finally, we use the reserved capacity to fine-tune the solution to obtain a feasible flow.

We now present a brief outline of the mimicking process used by this algorithm. The mimicking process works in a row-by-row fashion, i.e., the algorithm computes the flow matrix for the probabilistic instance by mimicking in order the rows of the flow matrix for the deterministic instance. It is ensured that the row-sums of the solution created by this process are exactly equal to the desired values, viz. \hat{a}_i for row i . Consider the stage where the first $i-1$ rows have already been mimicked. This means that we have created a partial flow matrix for which the entries of the first $i-1$ rows have already been determined. We now describe how the entries of the i^{th} row will be computed. At this point there may be a discrepancy (P_j) in the column-sums (for the first $i-1$ rows) between the deterministic and the mimicking solutions. Let P_j denote the excess of the j^{th} partial column-sum in the mimicking solution over that in the deterministic solution. While determining the values for the i^{th} row we will consider the columns in increasing order of discrepancy. The edges corresponding to the entries in row i are saturated until the desired row-sum is achieved. The behaviour of the column discrepancies is quite analogous to the combinatorial process outlined in section (4.1). To make the analogy complete it will be necessary to introduce a certain amount of fictitious capacity, as will be explained later. Let $\theta=1/N$, where N is a fixed positive integer such that $0<\theta<\epsilon$. Also, let D be the constant $D(K,\theta)$ determined by the Interval Theorem (section 4.1).

The Directed Transportation Algorithm

STEP (1) Set aside a small fraction of the edge capacities for use by the Fine-Tuning Algorithm. For each edge with a non-zero capacity, $\tilde{\alpha}(i,j)$, set aside one unit of capacity with probability δ independent of the other edges, where $0<\delta<\epsilon-\theta$. Let the expected value of the new capacities, $\{c(i,j)\}$, be $1+\epsilon'$; then $\theta<\epsilon'$.

STEP (2) (*Deterministic Relaxation*) Let $a=(a_i)$ and $b=(b_j)$ denote the supply and demand vectors, respectively. In linear time, construct the $(D+1)$ -realization of the pair (a,b) . Let $M=M(\hat{a},\hat{b})$ be the resulting matrix.

STEP (3) (*Mimicking Process*) Construct a flow $X=(x_{ij})$, row-by-row, such that $\sum_{j=1}^n x_{ij}=\hat{a}_i$ (for each i) and $\sum_{r=1}^i x_{rj}\approx\sum_{r=1}^i M_{rj}$ (for each j). The flow X can be constructed as follows. Suppose we are currently processing row i . Let $P_j=\sum_{r=1}^{i-1} x_{rj}-\sum_{r=1}^{i-1} M_{rj}$, for each column j . ($P_j=0$ if $i=1$).

STEP (3.1) $P_j \leftarrow P_j - M_{ij}$, for each j .

STEP (3.2) Let c_1, c_2, \dots, c_n denote the capacities of edges going from source i to the sinks in increasing order of P_j . Increase these capacities from c_l to c'_l so as to ensure that $\sum_{l=1}^t c'_l \geq (1+\theta)t$ and that $c_k \in [0, K]$, for each t, k . This may be done by choosing c'_l to be $\max\{c_l, [(1+\theta)l - \sum_{k=1}^{l-1} c'_k]\}$. The extra capacity introduced in this fashion will be called the *fictitious* capacity. Let U_i denote the amount of fictitious capacity required for row i ; i.e., $U_i = \sum_l (c'_l - c_l)$.

STEP (3.3) Send out \hat{a}_i units of flow out of source i by considering the edges in increasing order of P_j . Send out c'_j units of flow along the j^{th} such edge, until a total of \hat{a}_i units of flow have been shipped out. It may be observed that, due to the introduction of the fictitious capacities, the flow along an edge may exceed the actual capacity. Let x_{ij} denote the flow sent along the edge (i, j) .

STEP (3.4) $P_j \leftarrow P_j + x_{ij}$, for each j .

STEP 4 Repeat the mimicking process of STEP (3) with the roles of the rows and columns interchanged. In other words, construct a flow by mimicking the deterministic solution in a column-by-column fashion, using row discrepancies and introducing fictitious capacities as in Step (3). Let Y be the flow obtained in this manner. Also, let V_j denote the total amount of fictitious capacity introduced in column j .

STEP 5 Construct a flow X' which satisfies all capacity constraints by appropriately reducing the flow X along edges with fictitious capacities. Similarly, construct a flow Y' from the flow Y .

STEP 6 Consider the two K -matrices X' and Y' . Using the Generalized Mendelsohn-Dulmage algorithm (section 4.2), compute a third K -matrix Z such that the row- and column-sums of Z satisfy the bounds satisfied by the row- and column-sums of X' and Y' . The Generalized Mendelsohn-Dulmage algorithm ensures that each entry of Z matrix is no more than the larger of the corresponding entries in the X' and Y' matrices and hence no more than the corresponding edge's capacity.

STEP 7 (Fine Tuning) Let \bar{a} and \bar{b} denote the row and column sum vectors for the flow matrix Z . In STEP (1) the capacity set aside for each edge (i, j) was $\tilde{\alpha}(i, j) - c(i, j)$. Using these capacities and the Fine-Tuning Algorithm, construct a flow Z' with supply and demand vectors $(a - \bar{a})$ and $(b - \bar{b})$, respectively. The sum of the two flows, Z and Z' , is a feasible flow for the transportation problem instance under consideration.

4.4. Analysis of the Directed Transportation Algorithm

The proof of the Directed Transportation Theorem will be presented via the following lemmata. But, first, observe that the algorithm runs in linear time since each step requires at most $O(n^2)$ operations, and the size of input is $\Omega(n^2)$. We now proceed to

show that each step of the algorithm succeeds with high probability.

LEMMA 7 : At the end of STEP (1), the remaining capacities have expected value at least $1+\epsilon'=1+\epsilon-\delta$.

PROOF : Consider an edge (i, j) . The original capacity of this edge is $c(i, j)$ and the capacity at the end of STEP (1) is $\tilde{\alpha}(i, j)$. Let $p_k = \text{Prob}[\text{edge } (i, j) \text{ has } c(i, j) = k]$, for each $k \in [0, K]$.

$$E[\tilde{\alpha}(i, j)] = \sum_{k=1}^K p_k [\delta(k-1) + (1-\delta)k] \geq E[c(i, j)] - \delta \geq 1 + \epsilon - \delta$$

●

A generalized random walk analysis yields the following bounds on the amount of fictitious capacity introduced by the algorithm.

LEMMA 8 : For all $\gamma > 0$ there exists $s > 0$ such that $\max_i U_i \leq s \log n$ with probability $1 - O(n^{-\gamma})$.

PROOF : We make use of the analysis of a generalized one-dimensional random walk [8] to prove this result. Consider the following random walk process. The particle is initially at some integral position $z > 0$. The r^{th} step is given by the random variable X_r which takes only integral values. Let S_r denote the position of the particle after r steps. Then, $S_0 = z$ and $S_t = \sum_{r=1}^t X_r + z$ for $t > 0$. Let $u_z(a)$ denote the probability of the particle going to a position ≤ 0 before it goes to a position $\geq a$, for some fixed integer $a > z$.

Suppose the following conditions hold,

- (a) (X_r) are i.i.d. random variables
- (b) $m = E[X_1] > 0$
- (c) $X_r \in [-\nu, \mu]$, where ν, μ are positive integers

Feller [8] proves that,

$$u_z(a) \leq \frac{\sigma^{a+\mu-1} - \sigma^z}{\sigma^{a+\mu-1} - 1}$$

where σ is the unique positive root (other than 1) of the characteristic function of the probability distribution of X_r . It can be shown that $0 < \sigma < 1$ when $m > 0$. Note that σ is a constant which only depends on the distribution of (X_r) .

Consider now the fictitious capacity introduced for row i in STEP (3.2). Let $U_i(t) = (1+\theta)t - \sum_{r=1}^t c_r$, then $U_i = \max_t U_i(t)$. Recall that $\theta = 1/N$ where N is a positive integer such that $0 < 1/N < \epsilon'$. We want to show that U_i cannot be too large.

Let us relate the fictitious capacity to the random walk as follows. Let $X_r = Nc_r - (N+1)$. We now have that $S_t = -NU_i(t) + z$, here z is the initial position for the random walk which will be specified later. Since the random variables (c_r) are drawn from

the set $\{0,1,2, \dots, K\}$ and $E[c_r]=1+\epsilon'$, we have,

- (a) $X_r \in \{Nd-(N+1) : d \in \{0,1,2, \dots, K\}\}$
- (b) $E[X_r]=\epsilon'N-1 > 0$
- (c) $\nu=-(N+1), \mu=N(K-1)-1$

Suppose we select $a=n+z$ and $z=k \log n$, for some constant $k > 0$. Now, $u_z(a)$ was defined as being the probability that the value of S_t falls below 0 before it has ever risen above a . Equivalently, it is the probability that the value of U_i rises above z/N before it has ever fallen below $-(a-z)/N$. Clearly, the value of U_i can never fall below $-(a-z)/N$ unless $\sum_{r=1}^i c_r \geq n \geq \hat{a}_i$. The processing of row is over if n units of flow have been shipped out. This establishes that $u_z(a)$ is the probability that U_i is greater than $k \log n / N$. We now have,

$$\text{Prob}[U_i > \frac{k}{N} \log n] = u_z(a) \leq \frac{\sigma^{a+\mu-1}-\sigma^z}{\sigma^{a+\mu-1}-1} \leq \sigma^z$$

Summing the probability of failure over all n rows and given that $0 < \sigma < 1$, we have for some $\gamma > 0$,

$$\text{Prob}[\max_i U_i \leq \frac{k}{N} \log n] = 1 - O(n^{-\gamma})$$

•

The proof of the next lemma is identical.

LEMMA 9 : For all $\gamma > 0$ there exists $s > 0$ such that $\max_j V_j = O(\log n)$ with probability $1 - O(n^{-\gamma})$.

Consider now the flow X constructed in STEP (4). Since we used a certain amount of fictitious capacity, the *actual* flow X' is slightly less than X . The next lemma gives bounds on the amount of actual flow leaving a source as well as the actual flow entering a sink.

LEMMA 10 : Consider the flow matrix X' . The i^{th} row-sum of X' (the net flow out of source i) is $a_i - U_i - (D+1)$, while the j^{th} column sum of X' (the net flow into sink j) is less than $b_j - 1$.

PROOF : We draw an analogy between the construction of the flow X and the combinatorial process described in section (4.2). The value of P_j , after the processing of row $i-1$, corresponds to the position of particle j after $i-1$ state transitions. The set $\{j \mid M_{ij}=1\}$ corresponds to the set S arbitrarily chosen for the i^{th} state transition. The flow x_{ij} routed to the column j with the t^{th} smallest P_j corresponds to the distance d_t moved by the t^{th} leftmost particle. The use of fictitious capacities ensures that, for each t , $\sum_{k=1}^t c'_k \geq (1+\theta)t$ as required by the combinatorial process.

It follows that the net actual flow out of source i is exactly $\hat{a}_i - U_i = a_i - U_i - (D+1)$. We know that the net flow sent into sink j is $\hat{b}_j + P_j$. Invoking the Interval Theorem (section 4.1) we have that $P_j \leq D$. It follows that the net actual flow into sink j is no more than $\hat{b}_j + D = b_j - 1$. ●

LEMMA 10 : Consider the flow matrix Y' . The i^{th} row-sum of Y' (the net flow out of source i) is at most $a_i - 1$, while the j^{th} column sum of Y' (the net flow into sink j) is exactly $\hat{b}_j - V_j = b_j - V_j - (D+1)$.

PROOF : Similar to the proof for Lemma 10. ●

LEMMA 12 : Consider the flow matrix $Z(\bar{a}, \bar{b})$ constructed in STEP (6). It satisfies the following conditions,

$$\begin{aligned} a_i - U_i - (D+1) &\leq \bar{a}_i \leq a_i - 1, \quad \forall i \\ b_j - V_j - (D+1) &\leq \bar{b}_j \leq b_j - 1, \quad \forall j \end{aligned}$$

PROOF : The proof follows from Lemma 10, Lemma 11 and the application of the *GMD* theorem. ●

To complete the proof of the Directed Transportation Theorem we need to show that STEP (7) succeeds with high probability.

LEMMA 13 : The Fine-Tuning Algorithm succeeds in STEP (7) with probability $1 - O(n^{-\gamma})$, $\gamma > 0$.

PROOF : From Lemma 8, Lemma 9 and Lemma 12 we have,

$$\begin{aligned} 1 &\leq a_i - \hat{a}_i = O(\log n), \quad \forall i \\ 1 &\leq b_j - \hat{b}_j = O(\log n), \quad \forall j \end{aligned}$$

In STEP (1) we set aside a certain fraction of the edge capacities. The reserved capacity for an edge (i, j) is 1 with probability at least δ/K , it is 0 with probability at most $1 - (\delta/K)$. Moreover, the algorithm sets aside the capacities independently for each edge. The Fine-Tuning Theorem is now applicable to the supplies $a - \bar{a}$ and the demands $b - \bar{b}$. This completes the proof. ●

4.5. Application to the Directed Max-Flow Problem

We now apply the Directed Transportation Theorem to the solution of a probabilistic version of the directed max-flow problem. Consider the instances of the max-flow problem where the underlying graph is directed and the following conditions (M) are satisfied.

(M.1) $S = \{s\}$, $T = \{t\}$ and $I = \{1, 2, \dots, n\}$

(M.2) $\forall i \in I$, $c(s, i) = a_i$ and $c(i, t) = b_i$

(M.3) $\forall i, j \in I$ where $i \neq j$, the capacities $c(i, j)$ are i.i.d. random variables drawn from the set $\{0, 1, 2, \dots, K\}$ with expectation at least $1 + \epsilon$, where $\epsilon > 0$

(M.4) The pair (a, b) is $(D+1)$ -realizable for some constant $D > 0$

We will exhibit a linear time reduction from such instances of the max-flow problem to instances of the directed transportation problem satisfying the conditions for the Directed Transportation Theorem. This leads to a linear time algorithm for the max-flow problem and proves the following,

COROLLARY 1: Let I be an instance of the directed max-flow problem satisfying conditions (M). With high probability, the cut consisting of all edges leading out of the source (or into the sink) is a minimum cut. Moreover, there exists a linear time algorithm to find a maximum flow in I , with high probability of success.

We now specify the linear time reduction which will prove the above corollary. Let I be an instance of the max-flow problem satisfying conditions (M). We will reduce I to an instance C of the directed transportation problem. Construct a directed bipartite graph B as follows. Let $S' = \{s_1, s_2, \dots, s_n\}$ and $T' = \{t_1, s_2, \dots, s_n\}$ denote the bipartite vertex set for B . Associate with each source vertex s_i the supply a_i and with each sink vertex t_j associate the demand b_j . Every edge (s_i, t_j) is present and all edges are directed from S' to T' . Let $c(s_i, t_j) = c(i, j)$ for $i \neq j$.

Suppose we choose $c(s_i, t_i) = \infty$, for each i . It is then easy to see that a feasible flow for C can always be transformed (in linear time) to a feasible flow for I which saturates all edges leading out of s . Clearly, this would be a maximum flow for I . However, we will choose $c(s_i, t_i)$ to be i.i.d. random variables like the remaining capacities. This is done to ensure that the Directed Transportation Algorithm can be used to solve the instance C .

The Directed Transportation Theorem, when applied to C , implies that C is almost surely feasible. Also, the Directed Transportation Algorithm will almost surely find a feasible flow for C . We can now derive, in linear time, a feasible flow for I which will also be a maximum flow. This completes the proof of the corollary. Note that our reduction may seem weak since we do not exploit the fact that $c(s_i, t_i) = \infty$, for all i . As it turns out, however, we can prove that C is almost surely feasible anyway and, thus, find a maximum flow for I .

5. Further Work

Generalizations of current results:

The results presented above could be extended in many directions. It would be interesting to consider different distributions for the edge capacities. We could also look at the case of sparse graphs, i.e. graphs where the probability of an edge, $p(n)$, is small. Another possibility is to consider instances of the directed transportation problem where $K = K(n)$ and $\epsilon = \epsilon(n)$, where the former goes to infinity with n while the latter goes to zero as n approaches infinity.

Large Diameter Graphs:

It has been empirically observed that most network flow algorithms tend to have their worst performance on graphs of large diameter. It would be interesting to consider flows on random graphs which have large diameters. This could be done by considering sparse graphs or random *layered* graphs where the diameter can be increased by increasing the number of layers of vertices.

Mimicking Deterministic Solutions:

In our algorithms we "mimic" deterministic solutions of certain easy problems. We construct a deterministic problem where all random variables are replaced by their expected values. This gives rise to very special cases of the deterministic instances. These can easily be solved in linear time. We then try to mimic these deterministic solutions in the probabilistic instance, while trying to keep the "error" within bounds. The error in this mimicking process can be easily handled since it is shown to be small with high probability.

It would be extremely interesting to apply this technique to other problems. In particular, we might consider probabilistic instances of *Multi-commodity Flow* problems as a natural extension of the problems considered above. Another possibility is the *Minimum-Cost Flow* problem. Of course, there is no reason why this technique should work only for flow problems.

References

1. D. Angluin and L. Valiant, "Fast probabilistic algorithms for Hamiltonian circuits and matchings," *J. Comput. System Sci.*, vol. 19, pp. 155-193, 1979.
2. B. Bollobás, *Random Graphs*, Academic Press, 1985.
3. H. Chernoff, "A Measure of Asymptotic Efficiency for Tests Based on the Sum of Observations," *Annals of Math. Stat.*, vol. 23, pp. 493-509, 1952.
4. P. Doulliez and E. Jamoulle, "Transportation Networks with Random Arc Capacities," *Revue Francaise Automat. Informa. Recherche Operationelle*, vol. 3, pp. 45-60, 1972.
5. P. Erdős and A. Rényi, "On random graphs," *Publ. Math. Debrecen*, vol. 6, pp. 290-297, 1959.
6. P. Erdős and A. Rényi, "On random matrices," *Publ. Math. Inst. Hungar. Acad. Sci.*, vol. 8, pp. 455-461, 1964.
7. P. Erdős and A. Rényi, "On random matrices II," *Studia Sci. Math. Hungar.*, vol. 3, pp. 459-464, 1968.
8. W. Feller, *An Introduction to Probability Theory and Its Applications, Volume I*, pp. 363-366, John Wiley & Sons, 1950.
9. L. R. Ford and D. R. Fulkerson, *Flows in Networks*, Princeton University Press, 1962.

10. H. Frank and S. L. Hakimi, "Probabilistic Flows Through a Communication Network," *IEEE Trans. Circuit Theory*, vol. CT-12, pp. 413-414, 1965.
11. H. Frank and I. T. Frisch, *Communication, Transmission and Transportation Networks*, Addison-Wesley, 1971.
12. D. R. Fulkerson and H. J. Ryser, "Widths and Heights of (0,1)-Matrices," *Canad. J. Math.*, vol. 13, pp. 239-255, 1961.
13. D. Gale, "A Theorem on Flows in Networks," *Pacific J. Math.*, vol. 7, pp. 1073-1082, 1957.
14. A. V. Goldberg and R. E. Tarjan, "A New Approach to the Maximum Flow Problem," *Eighteenth STOC*, pp. 136-146, 1986.
15. G. R. Grimmett and D. R. A. Welsh, "Flow in Networks with Random Capacities," *Stochastics*, vol. 7, pp. 205-229, 1982.
16. G. R. Grimmett and H. S. Suen, "The Maximal Flow Through a Directed Graph with Random Capacities," *Stochastics*, vol. 8, pp. 153-159, 1982.
17. R. Hassin and E. Zemel, "Probabilistic analysis of the capacitated transportation problem," *Discussion Paper No. 660*, Center for Mathematical Studies in Economics and Management Science, Evanston, Illinois, 1985.
18. R. M. Karp, "The Probabilistic Analysis of Combinatorial Optimization Algorithms," *Tenth International Symp. on Math. Prog.*, 1979.
19. R. M. Karp, J. K. Lenstra, C. J. H. McDiarmid, and A. H. G. Rinnooy Kan, *Probabilistic Analysis of Combinatorial Algorithms: An Annotated Bibliography*, pp. 52-88, John Wiley & Sons, 1985.
20. E. L. Lawler, *Combinatorial Optimization: Networks and Matroids*, Holt, Rinehart and Winston, 1976.
21. N. S. Mendelsohn and A. L. Dulmage, "Some Generalizations of the Problem of Distinct Representatives," *Canad. J. Math.*, vol. 10, pp. 230-241, 1958.
22. A. Prékopa and E. Boros, "On the probability of the existence of a feasible flow in a transportation network," *RUTCOR Research Report No. 20-86*, The State University of New Jersey, Rutgers, 1986.
23. H. J. Ryser, "Combinatorial Properties of Matrices of Zeros and Ones," *Canad. J. Math.*, vol. 9, pp. 371-377, 1957.