I N F O R M A T I K

Polynomial Time Algorithms for Network Information Flow

Peter Sanders, Sebastian Egner, and Ludo Tolhuizen

MPI–I–2003–1–008 March 2003

FORSCHUNGSBERICHT RESEARCH REPORT

M A X - P L A N C K - I N S T I T U T
F Ü R
I N F O R M A T I K

Stuhlsatzenhausweg 85 66123 Saarbrücken Germany

Authors' Addresses

Peter Sanders
Stuhlsatzenhausweg 85
Max-Planck-Institut für Informatik,
66123 Saarbrücken, Germany
email: `sanders@mpi-sb.mpg.de`

Sebastian Egner, Ludo Tolhuizen
Philips Research Laboratories
Prof. Holstlaan 4
5656 AA Eindhoven, The Netherlands
email:`{sebastian.egner,ludo.tolhuizen}@philips.com`

Abstract

The famous max-flow min-cut theorem states that a source node s can send information through a network (V, E) to a sink node t at a rate determined by the min-cut separating s and t. Recently it has been shown that this rate can also be achieved for multicasting to several sinks provided that the intermediate nodes are allowed to reencode the information they receive. We give polynomial time algorithms for solving this problem. We additionally underline the potential benefit of coding by showing that multicasting without coding sometimes only allows a rate that is a factor $\Omega(\log |V|)$ smaller.

This work is partially supported by DFG grant SA 933/1-1.

Keywords

network flow, multicasting, coding, minimum cut, maximum flow, finite field, randomized algorithm, derandomization

1 Introduction

In this paper, we consider the problem of multicasting from a source s to multiple sinks in a network $G = (V, E)$ with maximum data rate, assuming that channels in the network are free of transmission errors. Ahlswede et al. [1] have shown that coding within the network can nevertheless increase the obtainable data rate. It is even possible to match the upper bound implied by the minimum cut between the source and a sink. In Section 4 we show that sometimes the (data) rate (or bandwidth) achievable without coding must be a factor $\Omega(\log |V|)$ smaller. This area of *network information flow* is conceptually interesting because it brings together the seemingly unrelated concepts of coding and network flows.

So far, a limitation of multicasting with coding was that only exponential time algorithms for constructing appropriate coding schemes were known [1, 12, 11]. But we give polynomial time algorithms here (independently, Jaggi, Chou, and Jain [8] have obtained a similar result). In contrast, maximizing multicast data rate *without* coding is at least as hard as the minimum Steiner tree problem [3, 9], i.e., even for undirected graphs $1+\epsilon$ approximation is hard for some $\epsilon > 0$ and after intensive research, no algorithm achieving an approximation factor better than 1.55 is known [14].

Overview: We continue the introduction with basic notation in Section 1 and a short review of related work in Section 1. As our main result, Section 2 develops polynomial time algorithms for unit capacity directed acyclic graphs. Section 3 gives generalizations for capacitated and cyclic graphs that achieve almost optimal rate. The logarithmic gap to multicasting without coding is demonstrated in Section 4. Section 5 gives an alternative algorithm for multicasting with coding that trades faster coding scheme construction for higher requirements for the size of the finite field used. This algorithm also illustrates relations between different approaches to reasoning about multicasting with coding. Section 6 discusses some open questions.

Basic Notation: Consider an acyclic, unit capacity network $G = (V, E)$ where parallel edges are allowed. Node $s \in V$ is the *source* node, $T \subseteq V$ the set of *sink* nodes, h is the size of the min-cut separating s from some $t \in T$. We use the notation $\Gamma^-(v)$ and $\Gamma^+(v)$ for the set of edges reaching and leaving node v respectively; $\mathrm{head}(e)$ denotes the node at which edge e starts.

We use linear coding over a finite field $\mathbb{F}$. It will turn out that $|\mathbb{F}| = \mathcal{O}(|T|)$ is sufficient. The source node s gets h input symbols from $\mathbb{F}$. The symbol carried by an edge e is an element of $\mathbb{F}$. Let $m_e : \Gamma^-(\mathrm{head}(e)) \to \mathbb{F}$ denote the *local coding vector*, i.e., if $y(e)$ denotes the symbol carried by edge e, we have $y(e) = \sum_{e' \in \Gamma^-(\mathrm{head}(e))} m_e(e')y(e')$. Our task is to determine the coefficients $m_e(e')$ such that all sinks can reconstruct the original information from the symbols reaching them.

The symbol $\delta_{u,v}$ stands for 1 if $u = v$ and 0 otherwise.

Related Work: Ahlswede et al. [1] have shown that the source can multicast information at a rate approaching h to all the sinks as the symbol size approaches infinity. They give a simple example that shows that without coding this rate is not always achievable. Li et al. [12] show that linear coding can be used for multicasting with rate h and finite symbol size. Our algorithms can be viewed as fast implementations of the approach by

Li et al. The main difference is that Li et al. have to check an exponential number of edge sets to verify that the coding coefficients choosen for a particular edge are correct. We reduce this to a single edge set per sink node by making explicit use of precomputed flows to each sink. Koetter and Medard [11] give an elegant algebraic characterization of linear coding schemes that achieve the maximal data rate. They show that finite fields of size $\mathcal{O}(|T| \cdot h)$ are sufficient and give a polynomial time algorithm to verify a given linear network coding scheme. They also give an algorithm for constructing coding schemes. However, this algorithm involves checking a polynomial identity $F(\mathbf{x}) = 0$ for a multivariate polynomial F with an exponential number of coefficients. Jaggi, Chou, and Jain [8] have independently and concurrently developed and algorithm similar to our deterministic algorithm from Section 2 that runs in time $\mathcal{O}(|E|(|T|^4 + h^3))$ plus the time for $|T|$ maximum flow computations.

There is a number of results for multicasting without coding in undirected networks. The single source problem then amounts to (fractionally) packing Steiner trees connecting $\{s\} \cup T$. The fractional Steiner tree packing problem is equivalent to the minimum weight Steiner tree problem (MWST) via duality of the corresponding linear programs [3, 9]. In particular, approximation algorithms for MWST yield approximation algorithms for packing Steiner trees with the same approximation ratio using the ellipsoid method. Without coding, multiple multicast requests can be decomposed into a corresponding number of Steiner tree packing problems using linear programming. Using randomized rounding one obtains approximate integral solutions to multiple multicast requests [15, 3]. The approximation quality is a constant factor plus a logarithmic additive term. Jansen and Zhang [10] avoid using the ellipsoid method at the cost of a slightly worse approximation ratio. Using this approach, running times as low as $|E|^2 \cdot k \cdot \mathrm{polylog}(|E|)$ can be achieved where k is the number of unit data rate multicast requests.

2 Polynomial Time Coding

This section is concerned with establishing the following result.

Theorem 1. *Consider an acyclic unit capacity multigraph (V, E) and let h denote the minimum cut between the source s and any sink $t \in T$. The Linear Information Flow (LIF) algorithm constructs linear codes for all nodes in expected time $\mathcal{O}(|E| \cdot |T| \cdot h^2)$ and worst case time $\mathcal{O}(|E| \cdot |T| \cdot (h^2 + |T|^2))$ such that the following properties hold:*

- *Any finite field of size $|\mathbb{F}| \geq 2 \cdot |T|$ can be used to represent symbols sent along edges. For the deterministic case, $|\mathbb{F}| \geq |T|$ is sufficient.*
- *The source gets h information symbols as input.*
- *A node needs time $\mathcal{O}(\min(|T|, |\Gamma^-(\mathrm{head}(e))|))$ to compute the symbol to be sent along edge e. The source needs time $\mathcal{O}(h)$ for each edge.*
- *Each sink can reconstruct all h information symbols in time $\mathcal{O}(h^2)$.*

Building on the notation given in Section 1 we explain the effects of linear coding in terms of linear algebra. We introduce h parallel edges $e_1, \ldots, e_h$ from some new node s' to s that

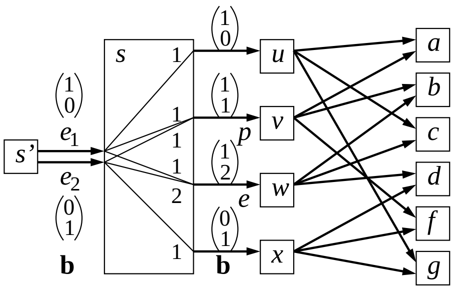

Figure 1: An example for multicasting with linear coding from s to $T = \{a, b, c, d, f, g\}$. We have $h = 2$. Assume that all the flows are decomposed into a topmost path and a bottommost path. The thin lines within s give nonzero coefficients for local coding vectors. The **b** vectors give the resulting global coding vectors. Let us assume that $\mathbb{F} = \mathrm{GF}(3)$ and that the edges leaving s are considered from top to bottom. Then the only feasible linear combinations for m_e are $[m_e(e_1) = 1, m_e(e_2) = 2]$ or $[m_e(e_1) = 2, m_e(e_2) = 1]$. As further examples for our notation we have $\Gamma^-(b) = \{(v, b), (w, b)\}$, $\mathrm{head}(e) = s$, $T(e) = \{b, c, d\}$, $P(e) = \{e_1, e_2\}$, $f^d_{\leftarrow}(e) = e_1$, and $f^c_{\leftarrow}(e) = e_2$. Before m_e is fixed, $C_b = \{p, e_2\}$ and correspondingly $B_b = \left\{\binom{1}{1}, \binom{0}{1}\right\}$.

carry the input symbols for the source s. Since linear coding is used, the information carried by an edge e is a linear combination of the input symbols of s. We can characterize the effect of all the local coding vectors on edge e independently of a concrete input using *global coding vectors* $\mathbf{b}(e) \in \mathbb{F}^h$. where $\mathbf{b}(e_i) = [0^{i-1}, 1, 0^{h-i}]$ and $\mathbf{b}(e) = \sum_{p \in \Gamma^-(\mathrm{head}(e))} m_e(p)\mathbf{b}(p)$ for $e \in E$. These vectors are well defined because the network is acyclic. Using elementary linear algebra it can be seen that a linear coding scheme can be used for multicasting from s to T if and only if for all $t \in T$, the vectors $\{\mathbf{b}(p) : p \in \Gamma^-(t)\}$ span the vector space $\mathbb{F}^h$. Reconstructing the original information can then be achieved by solving a linear system of equations of polynomial size. The intuition is that a linear code mixes the information received from different edges but it does not loose essential information as long as there is a bijective mapping between the input and the data reaching the sink.

The challenge is now to find the local coding vectors efficiently, possibly using a small finite field that allows fast arithmetics. Our algorithm achieves this goal by making explicit use of a maximum flow algorithm. Initially, it computes s-t flows f^t of magnitude h for each $t \in T$ and decomposes these flows into h edge disjoint paths from s to t. If there were only a single sink node, our task would be simple now. We could route the i-th input symbol along the i-th edge disjoint path. If an edge e is on some flow path W from s to t, let $f^t_{\leftarrow}(e)$ denote the predecessor edge of edge e on path W. In our single sink example, we could choose a nonzero coefficient for $m_e(f^t_{\leftarrow}(e))$ and zero for all other coefficients.

With multiple sinks, our approach is to superimpose multiple s-t flows. The algorithm steps through the nodes $u \in V$ in topological order. This ensures that the global coding vectors of all edges reaching u are known when the local coding vectors of the edges leaving u are determined. The algorithm defines the coefficients of m_e for one edge $e \in \Gamma^+(u)$ after the other. There might be multiple flow paths to different sinks through edge e. Let $T(e)$ denote the set of sinks using e in some flow f^t and let $P(e) = \{f^t_{\leftarrow}(e) : t \in T(e)\}$ denote the set of predecessors edges of e in some flow path. Nonzero coefficients for m_e are only chosen for edges in $P(e)$. To ensure that all sinks can reconstruct the input, the algorithm of Li et al. [12] verifies that the global coding vector $\mathbf{b}(e)$ is linearly independent of an exponential number of sets of other global coding vectors. Our algorithm can dramatically simplify this task by exploiting the flows. It turns out that only $\mathcal{O}(|T|)$ edge sets need to be checked for each $e \in E$.

We maintain the invariant that for each sink $t \in T$ there is a set of h edges C_t such that the global coding vectors $B_t := \{\mathbf{b}(c) : c \in C_t\}$ are a basis of $\mathbb{F}^h$, i.e., the original input can be reconstructed from the information carried by the edges in C_t. C_t contains one edge from each path in f^t, namely the edge whose global coding vector was defined most recently. Thus, when the computation completes, $C_t \subseteq \Gamma^-(t)$ and the invariant ensures that sink t gets all the information.

We initially establish the invariant by assigining the artificial input edges $\{e_1, \ldots, e_h\}$ with $\mathbf{b}(e_i) = [0^{i-1}, 1, 0^{h-i}]$ to C_t. When the linear combination m_e for a new edge e has been defined, we have to exchange $f^t_\leftarrow(e)$ by e in all the C_t with $t \in T(e)$. Hence, to maintain the invariant, it suffices to check for all $t \in T(e)$ whether B_t still spans $\mathbb{F}^h$. Figure 1 gives an example for the algorithm and its notation.

It remains to explain how to find coefficients for m_e that maintain the invariant. We argue that random coefficients for edges in $P(e)$ do the job if $\mathbb{F} \geq 2 \cdot |T|$: Lemma 2 below shows that for a fixed sink, the failure probability is only $1/|\mathbb{F}|$. Summing over all sinks, we see that the failure probability is at most $|T|/|F| \leq 1/2$. This straight-forward construction yields a Monte-Carlo algorithm for finding a *single* local coding vector, i.e., the construction can fail. To find *all* coding vectors fast using a small field, this is not sufficient. Here the knowledge of the flows encoded in the C_ts and the invariant come in and allow us to convert the Monte Carlo Algorithm into a Las Vegas algorithm, i.e., a constant expected number of trials followed by $|T(e)|$ independence tests suffices to find a feasible local coding vector.

Lemma 2. *For any $e \in E$ and $t \in T(e)$ assume that B_t is a basis of $\mathbb{F}^h$. Then with probability $1/|\mathbb{F}|$, a random local coding vector $m_e : P(e) \to \mathbb{F}$ fulfills the property that $(B_t \setminus \{\mathbf{b}(f^t_\leftarrow(e))\}) \cup \{\mathbf{b}(e)\}$ is a basis of $\mathbb{F}^h$ where $\mathbf{b}(e)$ is the corresponding global coding vector $\sum_{p \in P(e)} m_e(p)\mathbf{b}(p)$.*

Proof. If we fix the coefficients $m_e(p)$ for $p \in P(e) \setminus \{f^t_\leftarrow(e)\}$ then there is exactly one choice of $m_e(f^t_\leftarrow(e))$ for which $\mathbf{b}(e)$ is linearly dependent of $B_t \setminus \{\mathbf{b}(f^t_\leftarrow(e))\}$, i.e., there are exactly $|\mathbb{F}|^{|P(e)|-1}$ local coding vectors that violate the property for sink t. Since there are $|\mathbb{F}|^{|P(e)|}$ choices for local coding vectors, the probability that a random choice violates the property is $|\mathbb{F}|^{|P(e)|-1}/|\mathbb{F}|^{|P(e)|} = 1/|\mathbb{F}|$. □

What we have said so far already yields an algorithm running in polynomial expected time. In what follows, we further refine the algorithm to obtain a fast and more concrete implementation (Figure 2) and a deterministic way of choosing the linear combinations m_e.

2.1 Testing Linear Independence Fast

The mathematical basis for our refinement of the LIF algorithm is the following lemma which states that the problem of testing linear independence from an $h-1$ dimensional vector space can be reduced to a single scalar product.

Lemma 3. *Consider a basis B of $\mathbb{F}^h$ and vectors $\mathbf{b} \in B$, $\mathbf{a} \in \mathbb{F}^h$ such that $\forall \mathbf{b}' \in B : \mathbf{b}' \cdot \mathbf{a} = \delta_{\mathbf{b},\mathbf{b}'}$. Then, any vector $\mathbf{x} \in \mathbb{F}^h$ is linearly independent of $B \setminus \{\mathbf{b}\}$ if and only if $\mathbf{x} \cdot \mathbf{a} \neq 0$.*

Function LIF(V, E, s, T)

$h := \min_{t \in T} \min \{|C| : C \text{ is } s\text{-}t \text{ cut}\}$ —— $= \min_{t \in T} |\text{max flow from } s \text{ to } t|$
insert a new source s' into V —— help to establish the invariant
insert h parallel edges $\{e_1, \ldots, e_h\}$ from s' to s into E

let f^t denote a set of h edge disjoint paths from s to t —— the chosen flow from s to t
(* We use the notation $f^t_{\leftarrow}(e)$, $T(e)$, and $P(e)$ to access the flows ***)**
let $\mathbb{F}$ be the field of size 2^m, $m = \lfloor \log_2 \max \{|T(e)| : e \in E\} \rfloor + 1$

forall i: $\mathbf{b}(e_i) := [0^{i-1}, 1, 0^{h-i}]$ —— the i-th unit vector of $\mathbb{F}^h$
forall $t \in T$ **do**
  $C_t := \{e_1, \ldots, e_h\}$ —— t is supplied through C_t
  $B_t := \{\mathbf{b}(e_1), \ldots, \mathbf{b}(e_h)\}$ —— the coding vectors span $\mathbb{F}^h$
  forall $c \in C_t$: $\mathbf{a}_t(c) := \mathbf{b}(c)$ —— inverse vectors

foreach vertex $v \in V \setminus \{s'\}$ in topological order **do**
  forall outgoing edges e of v **do**
    (* Invariant: $\forall t \in T : |C_t| = h$ and $\forall c, c' \in C_t : \mathbf{b}(c) \cdot \mathbf{a}_t(c') = \delta_{c,c'}$ ***)**
    choose a linear combination $\mathbf{b}(e) = \sum_{p \in P(e)} m_e(p)\mathbf{b}(p)$ such that —— **(*)**
      $\forall t \in T(e) : \mathbf{b}(e)$ is linearly independent of $B_t \setminus \{\mathbf{b}(f^t_{\leftarrow}(e))\}$
    forall $t \in T(e)$ **do**
      $C'_t := (C_t \setminus \{f^t_{\leftarrow}(e)\}) \cup \{e\}$ —— advance the set of edges C_t,
      $B'_t := (B_t \setminus \{\mathbf{b}(f^t_{\leftarrow}(e))\}) \cup \{\mathbf{b}(e)\}$ —— update B_t correspondingly, and
      $\mathbf{a}'_t(e) := (\mathbf{b}(e) \cdot \mathbf{a}_t(f^t_{\leftarrow}(e)))^{-1} \mathbf{a}_t(f^t_{\leftarrow}(e))$ —— update $\mathbf{a}_t$ correspondingly
      forall $c \in C_t \setminus \{f^t_{\leftarrow}(e)\}$: $\mathbf{a}'_t(c) := \mathbf{a}_t(c) - (\mathbf{b}(e) \cdot \mathbf{a}_t(c))\mathbf{a}'_t(e)$
      $(C_t, B_t, \mathbf{a}_t) := (C'_t, B'_t, \mathbf{a}'_t)$

return $(h, \{m_e : e \in E\}, \{(C_t, \mathbf{a}_t) : t \in T\}, \mathbb{F})$.

Figure 2: Linear Information Flow (LIF) coding with fast testing of linear independence. Given a network (V, E), a source s and a set of sinks T, the algorithm constructs linear codes for intermediate nodes such that the rate from s to T is maximal.

Proof. Let $\mathbf{x} = \sum_{\mathbf{b}' \in B} x(\mathbf{b}')\mathbf{b}'$ be the unique representation of $\mathbf{x}$ in the basis B. We get

$$\mathbf{x} \cdot \mathbf{a} = \sum_{\mathbf{b}' \in B} x(\mathbf{b}')\mathbf{b}' \cdot \mathbf{a} = \sum_{\mathbf{b}' \in B} x(\mathbf{b}')\delta_{\mathbf{b},\mathbf{b}'} = x(\mathbf{b}).$$

Now, $\mathbf{x}$ is linearly independent of $B \setminus \{\mathbf{b}\}$ if and only if $x(\mathbf{b}) \neq 0$, i.e, if and only if $\mathbf{x} \cdot \mathbf{a} = 0$. □

The LIF-algorithm given in Figure 2 maintains vectors $\mathbf{a}_t(c)$ for each sink t and edge $c \in C_t$ that can be used to test linear independence of $B_t \setminus \{\mathbf{b}(c)\}$. The invariant now becomes

$$\forall t \in T : |C_t| = h \text{ and } \forall c, c' \in C_t : \mathbf{b}(c) \cdot \mathbf{a}_t(c') = \delta_{c,c'} \ . \tag{1}$$

This invariant implies both the linear independence of B_t and the desired property of $\mathbf{a}_t$.

The algorithm in Figure 2 implements the outline of the LIF algorithm given above. To prove correctness we have to verify the loop invariant.

Lemma 4. *The Loop-Invariant (1) holds for* $(C_t, B_t, \mathbf{a}_t)$.

Proof. Proof by induction. Before the loop over the vertices, the Loop-Invariant (1) is trivially satisfied. Now assume as induction hypothesis that Loop-Invariant (1) holds for $(C_t, B_t, \mathbf{a}_t)$. We show that it holds for $(C'_t, B'_t, \mathbf{a}'_t)$.

In C_t we replace edge $f^t_\leftarrow(e)$ by edge e, hence the size of C'_t is the same as the size of C_t. According to the algorithm, $\mathbf{b}(e)$ is chosen linearly independent to $B_t \setminus \{\mathbf{b}(f^t_\leftarrow(e))\}$. Hence, $\mathbf{b}(e)\cdot\mathbf{a}_t(f^t_\leftarrow(e)) \neq 0$ by Lemma 3, and $\mathbf{a}'_t(e)$ is well defined. Finally, we verify $\mathbf{b}(c)\cdot\mathbf{a}'_t(c') = \delta_{c,c'}$ for all $c, c' \in C'_t$ by a short calculation:

$$\begin{aligned}
\mathbf{b}(e) \cdot \mathbf{a}'_t(e) &= \mathbf{b}(e) \cdot \left(\mathbf{b}(e) \cdot \mathbf{a}_t(f^t_\leftarrow(e))\right)^{-1} \mathbf{a}_t(f^t_\leftarrow(e)) = 1, \\
\mathbf{b}(e) \cdot \mathbf{a}'_t(c) &= \mathbf{b}(e) \cdot \mathbf{a}_t(c) - (\mathbf{b}(e) \cdot \mathbf{a}_t(c))(\mathbf{b}(e) \cdot \mathbf{a}'_t(e)) = 0 \quad \text{for } c \neq e \\
\mathbf{b}(c) \cdot \mathbf{a}'_t(e) &= \left(\mathbf{b}(e) \cdot \mathbf{a}_t(f^t_\leftarrow(e))\right)^{-1} (\mathbf{b}(c) \cdot \mathbf{a}_t(f^t_\leftarrow(e))) = 0 \quad \text{for } c \neq e \\
\mathbf{b}(c) \cdot \mathbf{a}'_t(c') &= \mathbf{b}(c) \cdot (\mathbf{a}_t(c') - (\mathbf{b}(e) \cdot \mathbf{a}_t(c'))\mathbf{a}'_t(e)) \\
&= \mathbf{b}(c) \cdot \mathbf{a}_t(c') - (\mathbf{b}(e) \cdot \mathbf{a}_t(c'))(\mathbf{b}(c) \cdot \mathbf{a}'_t(e)) \\
&= \mathbf{b}(c) \cdot \mathbf{a}_t(c') = \delta_{c,c'} \quad \text{for } c, c' \neq e.
\end{aligned}$$

□

Remark. If the vectors in B_t are arranged as the rows of a matrix $\mathbf{B}_t$ and the columns $\mathbf{a}_t$ are correspondingly arranged as a matrix $\mathbf{A}_t$, then the invariant is equivalent to $\mathbf{A}_t = \mathbf{B}_t^{-1}$. In this notation, the method of updating the inverse vectors $\mathbf{a}_t$ in the LIF algorithm is a special case of the Sherman-Morrison formula [13, Section 2.7].

What we have said so far, suffices to establish the complexity of the randomized variant of LIF:

Lemma 5. *If Line* **(*)** *in Figure 2 is implemented by choosing random* $m_e : P(e) \to \mathbb{F}$ *until the condition "*$\forall t \in T(e) : \mathbf{b}(e)$ *is linearly independent of* $B_t \setminus \{\mathbf{b}(f^t_\leftarrow(e))\}$*" is satisfied, then the algorithm can be implemented to run in time* $\mathcal{O}(|E| \cdot |T| \cdot h^2)$ *and the returned information allows decoding in time* $\mathcal{O}(h^2)$ *at each sink.*

Proof. Using a single graph traversal, we can find a flow augmenting path from s to $t \in T$ in time $\mathcal{O}(|E|)$ [2]. We apply this routine round robin to each sink until for some sink no augmenting paths are left. Hence, we can find h augmenting paths for each sink in time $\mathcal{O}(|T| \cdot |E| \cdot h)$.[1]

The algorithm then constructs the finite field $\mathbb{F}$ of 2^m elements for m chosen such that $|\mathbb{F}| < 4|T|$. This involves the creation of a lookup table of $|\mathbb{F}|$ entries for successors of

[1] We can also use Dinitz' algorithm [5] to find many paths in time $\mathcal{O}(|E|)$. For large h this also yields improved asymptotic time bounds for the flow computation part [6].

elements (Conway's "Zech-logarithm" [7]). Using this table, any arithmetic operation in $\mathbb{F}$ can be computed in constant time.[2]

Initializing C_t, B_t, and $\mathbf{a}_t(c)$ takes time $\mathcal{O}(|T| \cdot h^2)$. The two main loops collectively iterate over all edges so that there is a total number of $|E|$ iterations. Computing $P(e)$ takes time $\mathcal{O}(|T|)$ if the flows f^t maintain pointers to the predecessors of edges in the path decomposition of f^t.

Finding a random local coding vector m_e takes time $\mathcal{O}(|P(e)|) = \mathcal{O}(|T|)$. Computing $\mathbf{b}(e)$ and testing linear independence using the vectors $\mathbf{a}_t(c)$ takes time $\mathcal{O}(h \cdot |T|)$. Since the success probability is constant, the expected cost for finding a linearly independent m_e is $\mathcal{O}(1) \cdot \mathcal{O}(h \cdot |T|) = \mathcal{O}(h \cdot |T|)$. Computing $\mathbf{a}'_t(c)$ for all t and all $c \in C_t$ takes time $\mathcal{O}(|T| \cdot h^2)$.

Combining all the parts, we get the claimed expected time bound of $\mathcal{O}(|E| \cdot |T| \cdot h^2)$. Sink $t \in T$ can reconstruct the vector of input symbols $\mathbf{x}$ at s by computing $\mathbf{x} = \sum_{c \in C_t} \mathbf{a}_t(c) y(c)$, where $y(c)$ denotes the symbol received over edge $c \in C_t$. This takes time $\mathcal{O}(h^2)$. □

2.2 Deterministic Implementation

We now explain how the algorithm LIF in Figure 2 can be implemented deterministically. We develop a deterministic method for choosing the local coding vectors in step (*) using the following two lemmas that we formulate as pure linear algebra problems without using graph theoretic concepts.

Lemma 6. *Consider pairs $(\mathbf{x}_i, \mathbf{y}_i) \in \mathbb{F}^h \times \mathbb{F}^h$ with $\mathbf{x}_i \cdot \mathbf{y}_i \neq 0$ for $1 \leq i \leq n \leq |\mathbb{F}|$. Then there is a linear combination $\mathbf{u}$ of $\mathbf{x}_1, \ldots, \mathbf{x}_n$ such that $\forall i : \mathbf{u} \cdot \mathbf{y_i} \neq 0$.*

Proof. Let V denote the vector space spanned by $\mathbf{x}_1, \ldots, \mathbf{x}_n$ and let d denote its dimension. Then $K_i = \{\mathbf{x} \in V : \mathbf{x} \cdot \mathbf{y}_i = 0\}$ is a $d-1$ dimensional subspace of V for all i. Clearly, $\mathbf{u}$ satisfies $\mathbf{u} \cdot \mathbf{y_i} \neq 0$ iff $\mathbf{u} \notin K_i$. As the K_i have at least the null vector in common, the number of linear combinations $\mathbf{u}$ as required in the lemma is

$$U = \left| V \setminus \bigcup_{i=1}^{n} K_i \right| \geq |V| - \left(1 + \sum_{i=1}^{n} |K_i \setminus \{\mathbf{0}\}|\right) = |\mathbb{F}|^d - \left(1 + n\left(|\mathbb{F}|^{d-1} - 1\right)\right) = \left(1 - \frac{n}{|\mathbb{F}|}\right)|\mathbb{F}|^d + n - 1.$$

If $n < |\mathbb{F}|$ then $U \geq |\mathbb{F}|^{d-1} \geq 1$. If $n = |\mathbb{F}|$ then $U \geq |\mathbb{F}| - 1 \geq 1$ because a field has at least two elements. □

The linear combination $\mathbf{u}$ from Lemma 6 can be found by fixing one coefficient of the linear combination after the other in a greedy fashion.

Lemma 7. *A linear combination $\mathbf{u}$ proven to exist in Lemma 6 can be found in worst case time $\mathcal{O}(n^2(h + |\mathbb{F}|))$.*

Proof. The proof is by induction on n. For $n = 1$ we simply set $\mathbf{u} = \mathbf{x_1}$.

For the induction step from $n-1$ to n, we apply the induction hypothesis to obtain a vector $\mathbf{u}$ with $\mathbf{u} \cdot \mathbf{y}_i \neq 0$ for $i < n$. If $\mathbf{u} \cdot \mathbf{y}_n$ happens to be nonzero, we are done.

[2] If the table is considered too large, one can resort to the polynomial representation of field elements. In this case, no table is needed at the cost of additional factors in running time that are polylogarithmic in $|T|$.

Otherwise, we apply Lemma 6 to $(\mathbf{u}, \mathbf{y}_1), \ldots, (\mathbf{u}, \mathbf{y}_{n-1}), (\mathbf{x}_n, \mathbf{y}_n)$. As $n \leq |\mathbb{F}|$, there is a linear combination $\mathbf{u}'' = \beta\mathbf{u} + \gamma\mathbf{x}_n$ such that $\forall i \leq n : \mathbf{u}'' \cdot \mathbf{y}_i \neq 0$. In particular, $\mathbf{u}'' \cdot \mathbf{y}_n = \gamma\mathbf{x}_n \neq 0$ since $\mathbf{u} \cdot \mathbf{y}_n = 0$. Hence, $\gamma \neq 0$ and $\mathbf{u}' = \mathbf{u}''/\gamma$ is of the form $\alpha\mathbf{u} + \mathbf{x}_n$ and we have $\forall i \leq n : \mathbf{u}' \cdot \mathbf{y}_i \neq 0$. The coefficient α can be found by trying all field elements. All $|\mathbb{F}|$ trials can be made in time $\mathcal{O}(n(h + |\mathbb{F}|))$ by precomputing the scalar products $\mathbf{u} \cdot \mathbf{y}_i$ and $\mathbf{x}_n \cdot \mathbf{y}_i$ for $1 \leq i < n$. We get a total time of $\mathcal{O}(n^2(h + \mathbb{F}))$ for n iterations.[3] □

Lemma 7 can be used to find the linear combination m_e in the LIF algorithm: Apply Lemma 7 to $\{(\mathbf{b}(f^t_{\leftarrow}(e)), \mathbf{a}_t(f^t_{\leftarrow}(e))) : t \in T(e)\}$, i.e., let $\mathbf{b}(e) = \mathbf{u} = \sum_{t \in T(e)} u_t \mathbf{b}(f^t_{\leftarrow}(e))$ denote a vector with $\mathbf{b}(e) \cdot \mathbf{a}_t(f^t_{\leftarrow}(e)) \neq 0$. This value for $\mathbf{b}(e)$ can also be obtained by setting $m_e(p)$ to the sum of all u_t with $f^t_{\leftarrow}(e) = p$.

The deterministic part of Theorem 1 can now be proven analogously to the proof of Lemma 5. We just have to substitute the expected time $\mathcal{O}(|T| \cdot h)$ needed by the randomized routine for choosing m_e by the time $\mathcal{O}(|T(e)|^2 \cdot (|\mathbb{F}| + h)) = \mathcal{O}(|T|^2(|T| + h)$ needed to apply Lemma 7. We obtain a total execution time of $\mathcal{O}(|E|(|T| \cdot h^2 + |T|^2(|T| + h))) = \mathcal{O}(|E| \cdot |T| \cdot (h^2 + |T|^2))$.

3 Refinements

Pipelining: So far our analysis only covers sending h symbols from s to the sinks. In reality, we want to send h symbols per time step in a pipelined fashion. If nodes naively forward linear combinations of the symbols received in the previous time step, we can get a mix of the information sent during different time steps that is difficult to decode.

Still, this naive algorithms works fine for *layered graphs* where a node v can be mapped to a layer $\ell(v)$ such that $\ell(s) = 0$ and edges only exist between adjacent layers. In that case, the information received by a node at layer i at time step τ is a linear combination of the information sent from s at time step $\tau - i$. Acyclic graphs that are not layered, are easy to convert into layered graphs by replacing edges spanning $k + 1$ layers by a chain of k new nodes. Since these new nodes have indegree and outdegree 1, no additional coding is necessary. They can be emulated by the original network by simply introducing time delays, i.e., data sent along an edge spanning $k + 1$ layers is delayed by k steps.

Edge Capacities: We now consider acyclic graphs with integer edge capacities $c(e)$. By replacing each capacitated edge by $c(e)$ parallel unit capacity edges, Section 2 immediately yields pseudopolynomial time algorithms, i.e., algorithms with running time polynomial in $|V|$, $|E|$, and $\sum_{e \in E} c(e)$. However, $\sum_{e \in E} c(e)$ can be exponential in the number of bits needed to define the input graph. Hence, the question arises how to handle graphs with very large edge capacities.

Again, Section 2 (almost) suffices to solve this problem:

Theorem 8. *Let h denote the maximum flow in a network $G = (V, E)$ with edge capacities*

[3] A more detailed analysis shows that at most n trials are needed and that only a constant expected number of trials are needed if the coefficients are tried in random order and $|\mathbb{F}| \geq 2n$.

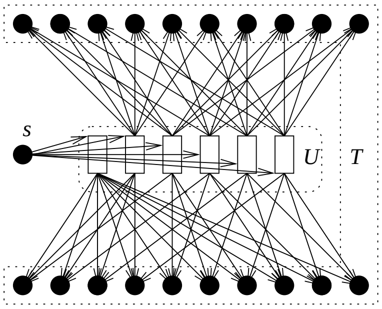

Figure 3: An example where three symbols per time step can be delivered. Without coding, the best we can do is to send three symbols every two time steps.

$c : E \to \mathbb{N}$. *For any $\epsilon > 0$ such that $\epsilon h \in \mathbb{N}$,*[4] *linear codes for network information flow can be found in time polynomial in $|V|$, $|E|$, and $1/\epsilon$ so that $(1-\epsilon)h$ symbols per time step can be communicated.*

Proof. In a preprocessing step, find the maximum flow f^t from s to each sink $t \in T$ and reduce $c(e)$ to $\max_{t\in T} f^t(e)$. In particular, no edge capacity exceeds h now. Let $w \leq |E|$ denote the maximum number of edge disjoint paths needed to realize any of flows f^t. Now build a unit capacity network $G' = (V, E')$ where each edge $e \in E$ corresponds to $c'(e) := \lfloor w \cdot c(e)/h\epsilon \rfloor$ parallel unit capacity edges in E'. Then find a multicast coding scheme for G'. This is possible in polynomial time since there are at most $w \cdot |E|/\epsilon$ edges in the unit capacity problem.

For coding and decoding in the capacitated instance, edge e is split into $c'(e)$ edges each with capacity $h\epsilon/w$. Each such edge transmits $h\epsilon$ symbols every w time steps using the encoding prescribed for the corresponding unit capacity edge. Thus, on each s-t path in any flow f^t, the unused capacity is at most $h\epsilon/w$, or $h\epsilon$ on all paths on one flow. Hence, the total used capacity is at least $h - h\epsilon = (1-\epsilon)h$. □

Cyclic Graphs: Ahlswede et al. [1] explain how to convert a cyclic graph $G = (V, E)$ with maximal obtainable rate h into an acyclic graph $G' = (V', E')$ with $|V'| = a \cdot |V|$, $|E'| = \mathcal{O}(a \cdot |E|)$ and rate $h' \geq (a - |V|)h$. They further explain how the communication on this graph can be emulated in a time steps on the cyclic graph G. This transformation can be used together with our polynomial time algorithm for finding coding schemes. By choosing $a = |V|/\epsilon$ we obtain a coding scheme that achieves a rate of $(1-\epsilon)h$.

4 The Gap to Multicasting Without Coding

Theorem 9. *There are unit capacity, acyclic networks where multicasting with coding allows a factor $\Omega(\log |V|)$ larger rate than multicasting without coding.*

Proof. Consider the network $G = (V, E)$ with vertices $V = \{s\} \cup U \cup T$ where $U = \{1, \ldots, 2h\}$, $T = \{t \subseteq U \mid |t| = h\}$ and edges $E = \{(s, u) \mid u \in U\} \cup \{(u, t) \mid t \in T, u \in t\}$. Note that $h = \Theta(\log |V|)$. Figures 1 and 3 give examples for $h = 2$ and $h = 3$ respectively.

[4]The requirement $\epsilon h \in \mathbb{N}$ avoids trivial rounding issues. By appropriately choosing our unit of time, we are quite flexible in choosing ϵ.

Since the min-cut has size h, the rate with coding is h.[5] Without coding, the rate is less than 2. To see this, suppose the source attempts to send $2n$ symbols $b_1, \ldots, b_{2n}$ to each of the sinks in T using n consecutive uncoded transmissions. Let U_i denote the subset of intermediate nodes receiving b_i. As the edges are unit capacity, the source can send at most $2hn$ symbols in total to the intermediate nodes, so $\sum_{i=1}^{2n} |U_i| \leq 2hn$. This implies that there is an i for which $|U_i| \leq h$.

For any h-element subset t of $U \setminus U_i$, the sink t does not receive the symbol b_i. $\square$

5 Faster Construction

We now outline an alternative algorithm for constructing linear network coding schemes. The algorithm is faster at the cost of using larger finite fields and hence possibly more expensive coding and decoding. Perhaps more importantly, this approach illustrates interesting connections between previous results and the present paper.

Theorem 10. *Linear network coding schemes using finite fields of size $|\mathbb{F}| \geq 2 \cdot |E| \cdot |T|$ and achieving rate h can be found in expected time $\mathcal{O}(|E| \cdot |T| \cdot h + |T| \cdot \mathcal{M}(h))$ where $\mathcal{M}(h) = \mathcal{O}(h^{2.376})$ denotes the time required for performing $h \times h$ matrix multiplications.*

Proof. (Outline) First find flows f^t decomposed into paths as before (time $\mathcal{O}(|E| \cdot |T| \cdot h)$). Then pick independent random local coding vectors m_e for all edges simultaneously. Compute all global coding vectors $\mathbf{b}(e)$ (time $\mathcal{O}\big(\sum_{e \in E} |P(e)| \cdot h\big) = \mathcal{O}(|E| \cdot |T| \cdot h)$). For each sink $t \in T$ let B_t denote the set of global coding vectors corresponding to edges ending a path in f^t. Check whether all the B_t span $\mathbb{F}^h$ (time $\mathcal{O}(|T| \cdot \mathcal{M}(h))$ using matrix inversion based on fast matrix multiplication [4]). If any of the tests fails, repeat.

Using Lemma 2 and the analysis of the "careful" algorithm from Section 2 it can be seen that the success probability is at least $1/2$: The careful algorithm would perform $\sum_{e \in E} |T(e)|$ independence tests, each would go wrong with probability $1/|\mathbb{F}|$. Hence, if we omit the tests the failure probability is at most $\sum_{e \in E} |T(e)|/|\mathbb{F}| \leq 1/2$. The expected number of repetitions will be constant. $\square$

This algorithm is quite similar to the one by Ahlswede et al. [1]. The main difference is that we choose random linear local coding functions rather than completely arbitrary random functions. Even the analysis of Ahlswede et al. could be adapted. Their analysis goes through if the arbitrary random functions are replaced by a random choice from a universal family of hash functions.[6] It is well known that random linear mappings form such a universal family. Exploiting the special structure of linear functions, this idea can be further developed into a polynomial time randomized algorithm achieving rate $(1-\epsilon)h$ for any constant $\epsilon > 0$. However, the analysis given here yields stronger results (rate h, exponentially smaller finite fields) because using flow arguments, we can reduce the exponential number

[5] In this network it suffices if the source s encodes the information using a code of length $2h$ that conveys h information symbols and allows reconstruction of the information from any h received symbols. This is called a $[2h, h]$ MDS code.

[6] A family $\mathcal{H} \subseteq A^B$ of functions from B to A is called *universal* if $\forall x, y \in B : \mathbf{Pr}[f(x) = f(y)] = 1/|A|$ for randomly chosen $f \in \mathcal{H}$.

of cuts considered in [1] to a polynomial number of edge sets that need to carry all the information.

Another interesting observation is that Koetter and Medard [11] arrive at a similar requirement for $|\mathbb{F}|$ as Theorem 10 using quite different algebraic arguments.

6 Discussion

Our polynomial time algorithms for multicasting in acyclic networks are quite simple and could perhaps be used in practice if the rate h is not too large (otherwise decoding could become too expensive). In addition, the algorithms are not limited to maximizing the rate. In a particular application, e.g. streaming multimedia content in a large computer network, the rate h' *required* could be fixed. In this situation, one could find a cheap or otherwise favorable subgraph with minimum cut at least h' and use our algorithms to construct codes to achieve multicast bandwidth h' from the server to each client.

There are a number of open problems for graphs with cycles. Is there a polynomial time algorithm that finds a coding scheme with rate *exactly* h? From a practical point of view even an approximate algorithm would be interesting if it allows coding schemes that are faster to decode. Currently, rate $(1-\epsilon)h$ requires decoding time $\Omega\left(h \cdot |V|/\epsilon\right)$ per decoded symbol which looks forbidding for large graphs.

Very little is known about the problem with multiple sources. It would be interesting to know whether we will again find that with coding optimal solutions are both easier to find and more efficient than solutions without coding.

Acknowledgements

We would like to thank Rudolf Ahlswede, Ning Cai, Philip Chou, Sidharth Jaggi, Irit Katriel, Piotr Krysta, Anand Srivastav, and Berthold Vöcking for fruitful discussions.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.

[2] R. K. Ahuja, R. L. Magnanti, and J. B. Orlin. *Network Flows*. Prentice Hall, 1993.

[3] B. Carr and S. Vempala. Randomized meta-rounding. In *32nd ACM Symposium on the Theory of Computing*, pages 58–62, 2000. also in Random Structures and Algorithms, 20(3):343–352, 2002.

[4] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9:251–280, 1990.

[5] E. A. Dinic. Algorithm for solution of a problem of maximum flow. *Soviet Math. Dokl.*, 11:1277–1280, 1970. now spelled 'Dinitz'.

[6] S. Even and E. Tarjan. Network flow and testing graph connectivity. *SIAM J. Comput.*, 4:507–518, 1975.

[7] K. Imamura. A method for computing addition tables in $GF(p^n)$. *IEEE Transactions on Information Theory*, 26:367–369, 1980.

[8] S. Jaggi, P. S. Chou, and K. Jain. Low complexity algebraic multicast network codes. In *International Symposium on Information Theory*. IEEE, July 2003. 1 page abstract, to appear.

[9] K. Jain, M. Mahdian, and M. R. Salavatipour. Packing Steiner trees. In *14th ACM-SIAM Symposium on Discrete Algorithms*, 2003. to appear.

[10] K. Jansen and H. Zhang. An approximation algorithm for the multicast congestion problem via minimum Steiner trees. In *3rd International Workshop on Approximation and Randomized Algorithms in Communication Networks (ARANCE)*, Rome, Italy, 2002.

[11] R. Koetter and M. Medard. An algebraic approach to network coding. In *Proceedings of INFOCOM*, 2002.

[12] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear network coding. *IEEE Transactions on Information Theory*, 2002. to appear.

[13] W. H. Press, S.A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 2nd edition, 1992.

[14] G. Robins and A. Zelikovsky. Improved steiner tree approximation in graphs. In *11th ACM Symposium on Parallel Architectures and Algorithms*, pages 770–779, 2000.

[15] S. Vempala and B. Vöcking. Approximating multicast congestion. In *10th Int. Symp. on Algorithms and Computation*, number 1741 in LNCS, pages 367–372, Chennai, 1999. Springer.

I N F O R M A T I K

Below you find a list of the most recent technical reports of the Max-Planck-Institut für Informatik. They are available by anonymous ftp from `ftp.mpi-sb.mpg.de` under the directory `pub/papers/reports`. Most of the reports are also accessible via WWW using the URL `http://www.mpi-sb.mpg.de`. If you have any questions concerning ftp or WWW access, please contact `reports@mpi-sb.mpg.de`. Paper copies (which are not necessarily free of charge) can be ordered either by regular mail or by e-mail at the address below.

Max-Planck-Institut für Informatik
Library
attn. Anja Becker
Stuhlsatzenhausweg 85
66123 Saarbrücken
GERMANY
e-mail: `library@mpi-sb.mpg.de`

MPI-I-2003-NWG2-002	F. Eisenbrand	Fast integer programming in fixed dimension
MPI-I-2003-NWG2-001	L.S. Chandran, C.R. Subramanian	Girth and Treewidth
MPI-I-2003-4-009	N. Zakaria	FaceSketch: An Interface for Sketching and Coloring Cartoon Faces
MPI-I-2003-4-008	C. Roessl, I. Ivrissimtzis, H. Seidel	Tree-based triangle mesh connectivity encoding
MPI-I-2003-4-007	I. Ivrissimtzis, W. Jeong, H. Seidel	Neural Meshes: Statistical Learning Methods in Surface Reconstruction
MPI-I-2003-4-006	C. Roessl, F. Zeilfelder, G. Nürnberger, H. Seidel	Visualization of Volume Data with Quadratic Super Splines
MPI-I-2003-4-005	T. Hangelbroek, G. Nürnberger, C. Roessl, H.S. Seidel, F. Zeilfelder	The Dimension of C^1 Splines of Arbitrary Degree on a Tetrahedral Partition
MPI-I-2003-4-004	P. Bekaert, P. Slusallek, R. Cools, V. Havran, H. Seidel	A custom designed density estimation method for light transport
MPI-I-2003-4-003	R. Zayer, C. Roessl, H. Seidel	Convex Boundary Angle Based Flattening
MPI-I-2003-4-002	C. Theobalt, M. Li, M. Magnor, H. Seidel	A Flexible and Versatile Studio for Synchronized Multi-view Video Recording
MPI-I-2003-4-001	M. Tarini, H.P.A. Lensch, M. Goesele, H. Seidel	3D Acquisition of Mirroring Objects
MPI-I-2003-2-004		?
MPI-I-2003-2-003	Y. Kazakov, H. Nivelle	Subsumption of concepts in $DL\ \mathcal{FL}_0$ for (cyclic) terminologies with respect to descriptive semantics is PSPACE-complete
MPI-I-2003-2-002	M. Jaeger	A Representation Theorem and Applications to Measure Selection and Noninformative Priors
MPI-I-2003-2-001	P. Maier	Compositional Circular Assume-Guarantee Rules Cannot Be Sound And Complete
MPI-I-2003-1-018		A Note on the Smoothed Complexity of the Single-Source Shortest Path Problem
MPI-I-2003-1-017	G. Schäfer, S. Leonardi	Cross-Monotonic Cost Sharing Methods for Connected Facility Location Games
MPI-I-2003-1-016	G. Schäfer, N. Sivadasan	Topology Matters: Smoothed Competitive Analysis of Metrical Task Systems”
MPI-I-2003-1-015	A. Kovács	Sum-Multicoloring on Paths
MPI-I-2003-1-014	G. Schäfer, L. Becchetti, S. Leonardi, A. Marchetti-Spaccamela, T. Vredeveld	Average Case and Smoothed Competitive Analysis of the Multi-Level Feedback Algorithm

MPI-I-2003-1-013	I. Katriel, S. Thiel	Fast Bound Consistency for the Global Cardinality Constraint
MPI-I-2003-1-012	D. Fotakis, R. Pagh, P. Sanders, P. Spirakis	- not published - Space Efficient Hash Tables with Worst Case Constant Access Time
MPI-I-2003-1-011	P. Krysta, A. Czumaj, B. Voecking	Selfish Traffic Allocation for Server Farms
MPI-I-2003-1-010	H. Tamaki	A linear time heuristic for the branch-decomposition of planar graphs
MPI-I-2003-1-009	B. Csaba	On the Bollobás – Eldridge conjecture for bipartite graphs
MPI-I-2003-1-008	P. Sanders	Polynomial Time Algorithms for Network Information Flow
MPI-I-2003-1-007	H. Tamaki	Alternating cycles contribution: a strategy of tour-merging for the traveling salesman problem
MPI-I-2003-1-006	M. Dietzfelbinger, H. Tamaki	On the probability of rendezvous in graphs
MPI-I-2003-1-005	M. Dietzfelbinger, P. Woelfel	Almost Random Graphs with Simple Hash Functions
MPI-I-2003-1-004	E. Althaus, T. Polzin, S.V. Daneshmand	Improving Linear Programming Approaches for the Steiner Tree Problem
MPI-I-2003-1-003	R. Beier, B. Vöcking	Random Knapsack in Expected Polynomial Time
MPI-I-2003-1-002	P. Krysta, P. Sanders, B. Vöcking	Scheduling and Traffic Allocation for Tasks with Bounded Splittability
MPI-I-2003-1-001	P. Sanders, R. Dementiev	Asynchronous Parallel Disk Sorting
MPI-I-2002-4-002	F. Drago, W. Martens, K. Myszkowski, H. Seidel	Perceptual Evaluation of Tone Mapping Operators with Regard to Similarity and Preference
MPI-I-2002-4-001	M. Goesele, J. Kautz, J. Lang, H.P.A. Lensch, H. Seidel	Tutorial Notes ACM SM 02 A Framework for the Acquisition, Processing and Interactive Display of High Quality 3D Models
MPI-I-2002-2-008	W. Charatonik, J. Talbot	Atomic Set Constraints with Projection
MPI-I-2002-2-007	W. Charatonik, H. Ganzinger	Symposium on the Effectiveness of Logic in Computer Science in Honour of Moshe Vardi
MPI-I-2002-1-008	P. Sanders, J.L. Träff	The Factor Algorithm for All-to-all Communication on Clusters of SMP Nodes
MPI-I-2002-1-005	M. Hoefer	Performance of heuristic and approximation algorithms for the uncapacitated facility location problem
MPI-I-2002-1-004	S. Hert, T. Polzin, L. Kettner, G. Schäfer	Exp Lab A Tool Set for Computational Experiments
MPI-I-2002-1-003	I. Katriel, P. Sanders, J.L. Träff	A Practical Minimum Scanning Tree Algorithm Using the Cycle Property
MPI-I-2002-1-002	F. Grandoni	Incrementally maintaining the number of l-cliques
MPI-I-2002-1-001	T. Polzin, S. Vahdati	Using (sub)graphs of small width for solving the Steiner problem
MPI-I-2001-4-005	H.P.A. Lensch, M. Goesele, H. Seidel	A Framework for the Acquisition, Processing and Interactive Display of High Quality 3D Models
MPI-I-2001-4-004	S.W. Choi, H. Seidel	Linear One-sided Stability of MAT for Weakly Injective Domain
MPI-I-2001-4-003	K. Daubert, W. Heidrich, J. Kautz, J. Dischler, H. Seidel	Efficient Light Transport Using Precomputed Visibility
MPI-I-2001-4-002	H.P.A. Lensch, J. Kautz, M. Goesele, H. Seidel	A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web
MPI-I-2001-4-001	H.P.A. Lensch, J. Kautz, M. Goesele, W. Heidrich, H. Seidel	Image-Based Reconstruction of Spatially Varying Materials
MPI-I-2001-2-006	H. Nivelle, S. Schulz	Proceeding of the Second International Workshop of the Implementation of Logics
MPI-I-2001-2-005	V. Sofronie-Stokkermans	Resolution-based decision procedures for the universal theory of some classes of distributive lattices with operators
MPI-I-2001-2-004	H. de Nivelle	Translation of Resolution Proofs into Higher Order Natural Deduction using Type Theory

MPI-I-2001-2-003	S. Vorobyov	Experiments with Iterative Improvement Algorithms on Completely Unimodel Hypercubes
MPI-I-2001-2-002	P. Maier	A Set-Theoretic Framework for Assume-Guarantee Reasoning
MPI-I-2001-2-001	U. Waldmann	Superposition and Chaining for Totally Ordered Divisible Abelian Groups
MPI-I-2001-1-007	T. Polzin, S. Vahdati	Extending Reduction Techniques for the Steiner Tree Problem: A Combination of Alternative-and Bound-Based Approaches