



Working Paper

Classifying customer-provider relationships in the internet

Author(s):

Erlebach, Thomas; Hall, A.; Schank, T.

Publication Date:

2002

Permanent Link:

<https://doi.org/10.3929/ethz-a-004403617> →

Rights / License:

[In Copyright - Non-Commercial Use Permitted](#) →

This page was generated automatically upon download from the [ETH Zurich Research Collection](#). For more information please consult the [Terms of use](#).

T. Erlebach, A. Hall, T. Schank

*Classifying Customer-Provider Relationships
in the Internet*

*TIK-Report
Nr. 145, July 2002*

Thomas Erlebach, Alexander Hall, Thomas Schank
Classifying Customer-Provider Relationships in the Internet
July 2002
Version 1
TIK-Report Nr. 145

Computer Engineering and Networks Laboratory,
Swiss Federal Institute of Technology (ETH) Zurich

Institut für Technische Informatik und Kommunikationsnetze,
Eidgenössische Technische Hochschule Zürich

Gloriastrasse 35, ETH Zentrum, CH-8092 Zürich, Switzerland

Classifying Customer-Provider Relationships in the Internet*

Thomas Erlebach, Alexander Hall[†]
Computer Engineering and Networks Lab
ETH Zürich
CH-8092 Zürich, Switzerland
{erlebach|hall}@tik.ee.ethz.ch

Thomas Schank
Department of Computer & Information Science
Universität Konstanz
D-78457 Konstanz, Germany
schank@fmi.uni-konstanz.de

July 2002

Abstract

The problem of inferring customer-provider relationships in the autonomous system topology of the Internet leads to the following optimization problem: given an undirected graph G and a set P of paths in G , orient the edges of G such that as many paths as possible are valid, meaning that they do not contain an internal node with both incident edges on the path directed away from that node. The complexity of this problem was left open by Subramanian et al. (“Characterizing the Internet hierarchy from multiple vantage points,” INFOCOM 2002). We show that finding an orientation that makes all paths valid (if such an orientation exists) can be done in linear time and that the maximization version of the problem is \mathcal{NP} -hard and cannot be approximated within $1/n^{1-\epsilon}$ for n paths unless $\mathcal{NP}=\text{co-}\mathcal{RP}$. We present constant-factor approximation algorithms for the case where the paths have bounded length and prove that the problem remains \mathcal{APX} -hard in this case. Finally, we report experimental results demonstrating that the approximation algorithm yields very good solutions on real data sets.

1 Introduction

The Internet is a huge, complex network whose present state is the outcome of a distributed growth process without centralized control. Because of the importance of the Internet as our

*Research supported by EU Thematic Network APPOL II, IST-2001-32007, and FET Open project COSIN, IST-2001-33555, with funding provided by the Swiss Federal Office for Education and Science (BBW).

[†]Supported by the joint Berlin/Zurich graduate program Combinatorics, Geometry, and Computation (CGC), financed by ETH Zurich and the German Science Foundation (DFG).

basic communication infrastructure, significant research efforts have recently been devoted to the discovery and analysis of its topology. The Internet topology can be considered either on the level of individual routers and hosts or on the level of *autonomous systems* (subnetworks under separate administrative control). In this paper, we focus on the autonomous system (AS) topology of the Internet. The AS topology can be represented as an undirected graph: each vertex corresponds to an AS, and two vertices are joined by an edge if there is at least one physical link between the corresponding ASs. The AS topology has been investigated by a number of authors, see e.g. [6, 13, 4, 15].

Different ASs exchange routing information using the Border Gateway Protocol (BGP). If an AS X is connected to another AS Y , it *announces* the routes to a certain set $A(X, Y)$ of destination addresses to Y . This implies that when Y has a packet with a destination in the set $A(X, Y)$, it can forward the packet to the AS X , because X knows a route to that destination. An AS may choose not to announce all routes that it knows to all of its neighbors; this decision is determined by the routing policy employed by the AS.

While the AS topology provides information about the connections between ASs, it has been pointed out in [8] that it is also important to know the economic relationships between ASs, because these relationships determine the routing policies and thus the paths that packets can potentially take in the network. If an AS X is connected to an AS Y , then Y can be a customer, provider, or peer of X . If Y is a customer of X , then X announces all its routes to Y . If Y is a provider or a peer of X , then X announces only the routes to destinations in its own AS and to destinations announced by its own customers; it does (usually) not announce routes to destinations that it can reach only through another provider or through peers. By knowing customer-provider relationships between ASs, one could have a sound basis for investigations of routing issues in the Internet.

Since data about these relationships is not easy to obtain directly, however, a natural idea is to infer AS relationships from the routing paths observed in the network (these paths can be determined from BGP information). Subramanian et al. [14] pursue this approach and propose a formulation of this inference problem as a combinatorial optimization problem, called the Type-of-Relationship problem: given a graph and a set of paths in the graph, classify the edges of the graph into customer-provider relationships and peer-peer relationships such that as many of the paths as possible are valid (consistent with this classification). Here, a path is valid if it consists of zero or more customer-provider edges, followed by at most one peer-peer edge, followed by zero or more provider-customer edges. Subramanian et al. write that they suspect the problem to be \mathcal{NP} -hard but have been unable to prove this, and they propose a heuristic algorithm.

1.1 Our results

In this paper, we consider two versions of the Type-of-Relationship problem: the problem of computing a classification such that *all* paths are valid (if such a classification exists), denoted ALLTOR, and the problem of computing a classification that maximizes the number of valid paths, denoted MAXTOR. We prove that ALLTOR can be solved in linear time by reducing it to 2SAT and we show that MAXTOR is \mathcal{NP} -hard, thus settling the complexity of the Type-of-Relationship problem left open in [14]. Our hardness proof implies that MAXTOR cannot be approximated within $1/n^{1-\varepsilon}$ (for any $\varepsilon > 0$) for general instances with n paths unless $\mathcal{NP}=\text{co-}\mathcal{RP}$. Motivated by the characteristics of instances arising in practice, we then consider the case where the given paths are short. First, we show that MAXTOR can be approximated within a factor of $(k+1)/2^k$ if all paths

have length at most k . Then we give approximation algorithms achieving a better approximation ratio for $k = 2, 3, 4$. We also prove that MAXTOR is \mathcal{APX} -complete for paths of bounded length. \mathcal{APX} -completeness (see e.g. [2]) implies that MAXTOR cannot be approximated within a certain constant factor unless $\mathcal{P} = \mathcal{NP}$, even for instances with short paths. Finally, we have implemented our approximation algorithm and obtained very encouraging results on real data sets.

Independently of our work, Di Battista, Patrignani and Pizzonia [5] have recently also obtained the result that ALLTOR can be solved in linear time and that MAXTOR is \mathcal{NP} -hard. Interestingly, they also use 2SAT to solve the ALLTOR problem and present an \mathcal{NP} -hardness proof by reducing MAX2SAT. Their reduction is based on the same idea with which we are able to prove \mathcal{APX} -hardness. Di Battista et al. do not consider the question of approximability or inapproximability of MAXTOR. Instead, they give an \mathcal{NP} -hardness result for the problem of maximizing peer-peer relationships and present a new heuristic algorithm for MAXTOR. They do not give approximation bounds for the algorithm, but they show that it performs well on real data sets.

The remainder of the paper is structured as follows. Definitions and preliminaries are given in Section 2. Section 3 deals with the ALLTOR problem. The hardness results for general MAXTOR are given in Section 4. In Section 5, the approximability of MAXTOR instances with paths of bounded length is studied. Section 6 describes our experimental results.

2 Preliminaries

We are given a simple, undirected graph $G = (V, E)$ whose nodes correspond to the autonomous systems of the Internet and whose edges correspond to physical connections between those autonomous systems. Furthermore, we are given a set P of simple, undirected paths in G . (We allow that P contains a path several times, i.e., P can actually be a multi-set.) These paths are possible data routes in the Internet; they are usually obtained from BGP routing tables. Informally, the goal is to classify the edges of G as customer-provider or peer-peer relationships in a “correct” way by using only the information obtained from the set of paths P . Here, the basic assumption is that the routing policies of the ASs depend on these relationships in the way described in Section 1, thus motivating the following definition (given in [14]).

Definition 1 *For a classification of the edges of G into customer-provider and peer-peer relationships, a path $p \in P$ is valid if*

- *it starts with zero or more customer-provider edges*
- *followed by zero or one peer-peer edge*
- *followed by zero or more provider-customer edges.*

The problem of classifying the edges into peer-peer or customer-provider relationship such that a maximum number of paths are valid was posed as the Type-of-Relationship (ToR) problem in [14]. We consider two versions of this problem: In the first variant, denoted ALLTOR, the goal is to decide whether there is an edge classification such that all paths in P are valid, and to compute such a classification if it exists. In the second variant, denoted MAXTOR, the goal is to compute an edge classification such that a maximum number of paths in P are valid.



Figure 1: Representation of customer-provider and peer-peer relationship.

For convenience we represent a customer-provider edge as a directed edge from customer to provider and a peer-peer edge as a bidirected edge between the two peers as shown in Figure 1.

Figure 2 gives some examples of valid and invalid paths. Using the directed and bidirected edge formulation for customer-provider and peer-peer relationships, we can rephrase the validity of a path as follows, where we take a bidirected edge to be pointing away from both of its endpoints.

Lemma 2 *For a given edge classification (represented as directed and bidirected edges), a path p is valid if and only if it does not contain a node from which two edges of p are pointing away.*

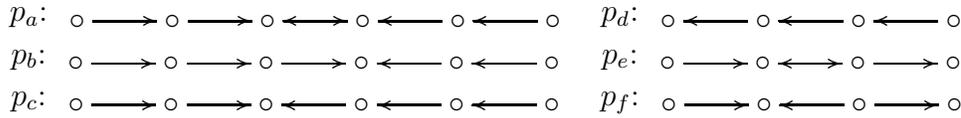


Figure 2: Paths p_a, \dots, p_d are valid, p_e and p_f are not valid.

Finally, we argue that peer-peer edges can be completely disregarded in this formulation of the problem. To see this, assume that for a graph G and a set of paths P , the relationships between the ASs have been classified as directed or bidirected edges. Now replace every bidirected edge by a directed edge in arbitrary direction. Using Lemma 2, we see immediately that all the valid paths remain valid after this operation. Figure 2 illustrates this transition for both directions from path p_a to path p_b and from path p_a to p_c , respectively. Therefore, we can assume without loss of generality that the solutions for MAXTOR and ALLTOR computed by an algorithm do not contain any bidirected edges.¹

An assignment of directions to the edges of an undirected graph is called an *orientation* of the graph. We can now state the ALLTOR and MAXTOR problems simply as follows:

ALLTOR: Given a graph G and a set of paths P , decide whether there is an orientation of G such that all paths in P are valid, and compute such an orientation if it exists.

MAXTOR: Given a graph G and a set of paths P , compute an orientation of G such that a maximum number of paths in P are valid.

In the case of MAXTOR, we will be interested in approximation algorithms. Here, an algorithm for MAXTOR is a ρ -approximation algorithm if it runs in polynomial time and always computes an orientation of G such that the number of valid paths is at least ρ times the number of valid paths in the optimal orientation. Note that $\rho \leq 1$.

The length of a path is defined as the number of edges on the path. All nodes on a path except its two endpoints are called *internal* nodes of the path.

¹From a practical point of view, it might be interesting to find an alternative formulation of the problem where peer-peer edges are actually necessary.

3 The ALLTOR problem

In this section we show that the ALLTOR problem is solvable in linear time. To this end we show that an instance of ALLTOR can be reduced to a 2SAT problem.

Lemma 3 *A path $p \in P$ of length k can be split up into $k - 1$ paths p_i , $1 \leq i \leq k - 1$, each of length 2, such that for any orientation of G , p is valid if and only if all p_i are valid.*

Proof: The construction of the $k - 1$ paths works as follows: path p_1 consists of the first two edges of p . Path p_i consist of edge i and edge $i + 1$ of path p . In this way, path p_i overlaps with path p_{i+1} by one edge. Figure 3 shows how a path of length 3 is split up into two paths of length 2. The correctness of the lemma follows directly from Lemma 2. \square

	directed path p_1	e_1	e_2	p_1 valid	2SAT clause
p :	$\circ \xrightarrow{e_1} \circ \xrightarrow{e_2} \circ \xrightarrow{e_3} \circ$	in	in	yes	$e_1 \vee e_2$
p_1 :	$\circ \xrightarrow{e_1} \circ \xrightarrow{e_2} \circ$	in	out	yes	$e_1 \vee \overline{e_2}$
p_2 :	$\circ \xleftarrow{e_2} \circ \xleftarrow{e_3} \circ$	out	in	yes	$\overline{e_1} \vee e_2$
	$\circ \xleftarrow{e_2} \circ \xrightarrow{e_3} \circ$	out	out	no	$\overline{e_1} \vee \overline{e_2}$

Figure 3: Decomposition of paths (left), and truth table for a path of length 2 (right).

By Lemma 3, it suffices to consider the ALLTOR problem for instances with paths of length 2. (Paths of length 1 are always valid.) The truth table on the right-hand side of Figure 3 shows a one-to-one correspondence between the validity of a path of length 2 and the logical or of two literals. This suggests the use of 2SAT to solve the ALLTOR problem, resulting in the following algorithm:

1. Orient the edges of $G = (V, E)$ arbitrarily.
2. Split all paths $p \in P$ into $\sum_{p \in P} (\ell(p) - 1)$ paths of length 2 as shown above, where $\ell(p)$ is the length of path p .
3. Construct a 2SAT instance where each edge $e_i \in E$ corresponds to a variable and each path p (of length 2) corresponds to a clause in the following way: like in the truth table of Figure 3, e_i appears negated if the corresponding edge is pointing away from the internal node v of p and not negated if it is pointing towards v .
4. Solve the resulting 2SAT instance.
5. If the 2SAT instance is not satisfiable, the ALLTOR instance is not solvable. Otherwise, flip the directions of all edges whose corresponding variable has been assigned *false*. This gives an orientation where all paths are valid.

Correctness of the algorithm can be verified easily by looking at the four possible configurations for input paths of length two and the assignment of the variables in the corresponding clause. Steps 1, 2, 3 and 5 can obviously be performed in time linear in the size of the input. 2SAT is well known to be solvable in linear time, so the running-time for Step 4 is linear as well.

Theorem 1 *ALLTOR can be solved in linear time.*

4 The general MAXTOR problem

In this section we show that MAXTOR is \mathcal{NP} -hard and even cannot be approximated well. To achieve this we will give an approximation-preserving polynomial reduction from the well known \mathcal{NP} -hard maximum independent set problem (denoted MAXIS) to MAXTOR. First, consider a graph G with two paths as shown in Figure 4. It can be verified that in this graph there is no orientation of the edges such that both paths are valid: Each of the three shared edges of both paths would require to flip the direction in one of the paths to make them valid. But since a change of direction of the edges in a valid path can happen only once, one of the two paths must be invalid.

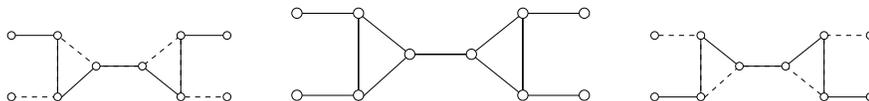


Figure 4: A graph (middle) with two paths (left and right) that cannot both be valid in any orientation.

Lemma 4 *There exists a graph G with two paths such that only one of the two paths can be valid.*

We will use this construction as a gadget to obtain a reduction from MAXIS to MAXTOR. Let an instance of MAXIS be given by an undirected graph $H = (V_H, E_H)$. We create an instance (G, P) of MAXTOR by mapping every node in H to a path in P such that any two paths in P

- are edge-disjoint if there is no edge between them in H and
- cannot be valid simultaneously in any orientation if there is an edge between them in H .

Obviously, such an instance (G, P) can be constructed in polynomial time using the gadget of Figure 4.

Now observe that there is a one-to-one correspondence between orientations of G with t valid paths and independent sets in H with cardinality t . In particular, if we could solve MAXTOR in polynomial time or approximate it in polynomial time with ratio ρ , we could also solve MAXIS in polynomial time or approximate it with ratio ρ , respectively. Since MAXIS is known to be \mathcal{NP} -hard and not even approximable with ratio $1/n^{1-\epsilon}$ on graphs with n nodes for any $\epsilon > 0$ unless $\mathcal{NP} = \text{co-}\mathcal{RP}$ [10], we obtain the following hardness result for MAXTOR.

Theorem 2 *MAXTOR is \mathcal{NP} -hard and cannot be approximated within a factor of $1/n^{1-\epsilon}$ on instances with n paths for any $\epsilon > 0$ unless $\mathcal{NP} = \text{co-}\mathcal{RP}$.*

Thus it is not feasible to get a good approximation ratio for the general MAXTOR problem. We might hope, however, to be able to exploit the structure of real AS graphs or the observed routing paths in order to give an algorithm with good approximation ratio for a restricted case of MAXTOR.

5 Approximating MAXTOR instances with bounded path length

Motivated by the negative result concerning the approximability of the general MAXTOR problem, we now consider a restricted version of the problem. Looking at real world data [14, 1], we noticed that most of the paths are actually quite short (see also Section 6). This leads us to consider instances of MAXTOR where the maximum length of the paths is bounded by a constant. In the following, we first present a very simple randomized approach achieving constant approximation ratio for this case. Then we show how to use an approximation algorithm for MAX2SAT to achieve significantly better approximation ratios for instances containing only paths of length at most k for $k = 2, 3, 4$. Finally, we prove \mathcal{APX} -completeness for instances where the path length is bounded by an arbitrary constant, by an approximation preserving reduction from the MAX2SAT problem, which is well known to be \mathcal{APX} -complete [2].

5.1 A simple constant factor approximation algorithm

For paths of constant length there is a very easy randomized approximation algorithm: just select the directions of the edges independently at random. If each edge is oriented in one of the two possible ways with probability $1/2$, a path of length k is valid with probability

$$p_k = \frac{k+1}{2^k}.$$

To see this, note that in a valid path the direction of the edges can change only once at one of the $k-1$ internal nodes or not at all. There are $k-1$ possibilities for the direction of all edges in the path in the former case and 2 possibilities in the latter case. Altogether there are 2^k possibilities to orient the edges. If all paths have length at most k , we get by linearity of expectation

$$\mathbb{E}(A_{rand}) \geq \frac{k+1}{2^k}n \geq \frac{k+1}{2^k}Opt,$$

where A_{rand} is the value of the approximate solution and Opt the value of the optimum. The algorithm can easily be derandomized in the standard way (method of conditional probabilities), giving the following theorem.

Theorem 3 *The MAXTOR problem with paths no longer than k edges can be approximated within a factor of $\frac{k+1}{2^k}$ of the optimum in polynomial time.*

For example, this gives ratio 0.75 for paths of length at most 2, 0.5 for paths of length at most 3, and $5/16 = 0.3125$ for paths of length at most 4. We can also state the following corollary.

Corollary 5 *The MAXTOR problem with average path length bounded by k can be approximated within a factor of $\frac{2k+1}{2 \cdot 4^k}$ of the optimum in polynomial time.*

Proof: Consider an instance of MAXTOR with n paths. If the average path length is k , there are at least $n/2$ paths with length at most $2k$. Applying the simple randomized algorithm described above to these $n/2$ paths, we obtain a solution with at least

$$\frac{n}{2} \cdot \frac{2k+1}{4^k}$$

valid paths. □

5.2 Better ratios for paths of short length

To obtain better ratios for paths of length at most 4 or less we use the approximation algorithm for MAX2SAT due to Lewin, Livnat and Zwick [12]. They achieve a ratio of $r = 0.940$ by enhancing the semidefinite programming (SDP) based approach first presented in the seminal paper by Goemans and Williamson [9]. Given an instance (G, P) of MAXTOR, we construct an instance of MAX2SAT from the paths as described in Section 3 (by splitting every path into paths of length two and adding a clause for each length 2 path) and apply the MAX2SAT approximation algorithm to the resulting instance. Then we orient the edges of G according to the assignment returned by the MAX2SAT algorithm. It may seem surprising that this approach gives a good ratio because there is not necessarily a one-to-one correspondence between paths and clauses.

First, a bit of notation: Let Opt denote the optimum value of the considered MAXTOR instance and A_k the value of our approximate solution for an instance with path length bounded by k . Let g_k be the number of paths in P that have length exactly k . Note that paths of length 1 can be ignored, since they are always valid. Hence, we can assume $g_1 = 0$. Furthermore, note that for deriving lower bounds on the approximation ratio A_k/Opt , it suffices to consider only instances in which all paths have length exactly k : If an instance contains a shorter path, this path can easily be lengthened by adding an appropriate number of extra nodes and edges to the path at one of its ends. At most $k \cdot n$ edges and nodes are added, and a solution to this modified instance clearly gives a solution to the original instance with at least the same number of valid paths.

In the simplest case, when all paths have length 2, we can directly transfer the ratio $r = 0.940$ from MAX2SAT to MAXTOR. This is because each path is represented by exactly one clause which is satisfied if and only if the path is valid.

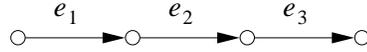


Figure 5: A path of length 3 which leads to the two clauses $e_1 \vee \bar{e}_2$ and $e_2 \vee \bar{e}_3$. See Section 3 for details about the definition of the clauses given a graph and paths.

Now consider a path of length 3. It is represented by two 2SAT clauses. The variable corresponding to the edge in the middle appears in both clauses, once negated and once not negated (see Figure 5 for an example). Therefore, one of the two clauses is always satisfied. Clearly, both clauses are satisfied if and only if the path is valid. This gap of either one or two clauses being satisfied can be used to derive a bound on the approximation ratio. Consider a MAXTOR instance with $n = g_3$ paths of length 3. Note that an optimal solution of MAX2SAT has the value $Opt + g_3$ and directly gives an optimal solution to the MAXTOR problem (with value Opt). The MAX2SAT approximation algorithm satisfies $A_3 + g_3$ clauses with

$$\frac{A_3 + g_3}{Opt + g_3} \geq r. \quad (1)$$

Because there is always a solution to MAX2SAT such that at least $3/4$ of the clauses are satisfied and there are $2 \cdot g_3$ clauses, we have $A_3 + g_3 \geq 3/2 \cdot g_3$ or $g_3 \leq 2 \cdot A_3$. Applying this to (1) leads to

$$\begin{aligned} A_3 &\geq r \cdot (Opt + g_3) - g_3 \geq r \cdot Opt + 2(r - 1) \cdot A_3 \\ A_3 &\geq \frac{r}{3 - 2r} Opt \end{aligned}$$

giving approximation ratio $r/(3 - 2r) \geq 0.839$. Note that this is a considerable improvement over the ratio $1/2$ obtained for paths of length three by Theorem 3.

In a MAX2SAT instance derived from paths of length 4 there will be three clauses for each of the paths. We refer to the three clauses of a path as a *triple*. With the same argumentation as for length 3 paths, we have that for each triple at least one clause is always satisfied and all three are satisfied if and only if the corresponding path is valid. This shows that any solution of this instance satisfies

$$x + y + g_4$$

clauses, where g_4 , y and x are the number of triples with at least one, at least two and exactly three satisfied clauses, respectively. Note that x yields the number of valid paths and $x \leq y \leq g_4$.

We now compare a solution $S = A_4 + y + g_4$ computed by the MAX2SAT approximation algorithm with a solution $S' = Opt + y' + g_4$ derived from an optimal solution of the corresponding MAXTOR instance. By [7] we know that $S/S' \geq r$. From this we can bound the approximation ratio A_4/Opt . In the worst case $y = g_4$ and $y' = Opt$. This gives

$$\frac{A_4 + 2 \cdot g_4}{2 \cdot Opt + g_4} \geq r$$

$$A_4 \geq 2r \cdot Opt + (r - 2) \cdot g_4 \geq 2r \cdot Opt + 4(r - 2)A_4,$$

because there are $3 \cdot g_4$ clauses of which at least $3/4$ are satisfied in the approximate solution, i.e. $A_4 + 2 \cdot g_4 \geq 9/4 \cdot g_4$ or $g_4 \leq 4 \cdot A_4$. Solving for A_4 we get

$$A_4 \geq \frac{2r}{9 - 4r} \cdot Opt.$$

This gives an approximation ratio of $2r/(9 - 4r) \geq 0.358$, which is a slight improvement compared to $5/16 = 0.3125$.

Note that this approach cannot be carried over to paths of length greater than 4. It uses the fact that at least $3/4$ of the clauses can be satisfied, which does not help if each path is represented by 4 or more clauses and the path is valid if and only if all of them are satisfied.

Summarizing the results discussed above, we get the following theorem.

Theorem 4 *The MAXTOR problem with paths no longer than k edges can be approximated within a factor c_k of the optimum in polynomial time, where c_k has the following form:*

- $c_2 := 0.940$
- $c_3 := 0.839$
- $c_4 := 0.358$
- $c_k := \frac{k+1}{2^k}$ for $k > 4$

5.3 \mathcal{APX} -hardness

So far we have given constant-factor approximation algorithms for instances of MAXTOR with paths of bounded length. Now we show that even instances containing only paths of length 2 cannot be approximated better than some constant unless $\mathcal{P} = \mathcal{NP}$. Both results together give that MAXTOR with constant maximal path length is \mathcal{APX} -complete.

To this aim, we reduce MAX2SAT and apply an \mathcal{APX} -hardness result by Håstad [11]. We start by reducing a MAX2SAT instance to a MAXTOR instance (G, P) with paths of length 3. Then we explain how this instance can be modified such that it contains only paths of length 2.

Assume that we are given a MAX2SAT instance with variables $x_i, i \in \{1 \dots n'\}$ and clauses $c_j, j \in \{1 \dots m'\}$. For each variable x_i we add two nodes $\overline{x_i}, x_i$ to G and an edge $e_i = \{\overline{x_i}, x_i\}$ between them. If in a solution to MAXTOR this edge is directed towards x_i this corresponds to $x_i = \text{true}$. Otherwise, if it is directed towards $\overline{x_i}$, this means $x_i = \text{false}$. For each clause $c_k = l_i \vee l_j$, with literals $l_i \in \{x_i, \overline{x_i}\}$ and $l_j \in \{x_j, \overline{x_j}\}$, one edge $\{l_i, l_j\}$ is added. Additionally a path of length 3 is added to P along the nodes $\overline{l_i}, l_i, l_j, \overline{l_j}$. Figure 6 shows a simple example with one clause, and the graph on the left-hand side of Figure 7 shows an example with two clauses.

MAX2SAT clause	path	graph	directed path	truth value
$(x_1 \vee \overline{x_2})$	$\overline{x_1}, x_1, \overline{x_2}, x_2$			$x_1 = 1, x_2 = 0$

Figure 6: Example of how an instance of MAXTOR is constructed from a MAX2SAT instance. On the right-hand side, a possible assignment of directions to the edges and the corresponding truth values are shown.

The paths and edges are defined such that a path is valid in a solution to the MAXTOR instance if and only if the corresponding clause is satisfied. This is clear because a path is valid if and only if either e_i is directed towards l_i or e_j towards l_j , which corresponds to a truth assignment where $l_i \vee l_j$ is satisfied. In particular, note that given a satisfied clause the edge $\{l_i, l_j\}$ can always be directed such that the whole path is valid. Conversely, for an unsatisfied clause no such direction of $\{l_i, l_j\}$ exists.

Thus, maximizing the number of valid paths also maximizes the number of satisfied clauses. With [11] we get that MAXTOR is not approximable within ratio $q = 0.955$ for instances with paths of length at most k for any constant $k \geq 3$.

To obtain a similar result for paths of length 2, we modify the instance as follows: each path $\overline{l_i}, l_i, l_j, \overline{l_j}$ is replaced by two overlapping paths $\overline{l_i}, l_i, l_j$ and $l_i, l_j, \overline{l_j}$. See the graph in the middle of Figure 7 for an example. Clearly, in a MAXTOR solution one of the two paths is always valid. The corresponding clause is satisfied if and only if both paths are valid. So an optimal solution to the MAXTOR instance with Opt valid paths gives an optimal assignment to the variables such that $Opt - m'$ clauses are satisfied. An approximate solution to MAXTOR giving A_2 valid paths leads to $A_2 - m'$ satisfied clauses. With [11] we know that

$$\frac{A_2 - m'}{Opt - m'} \leq q$$

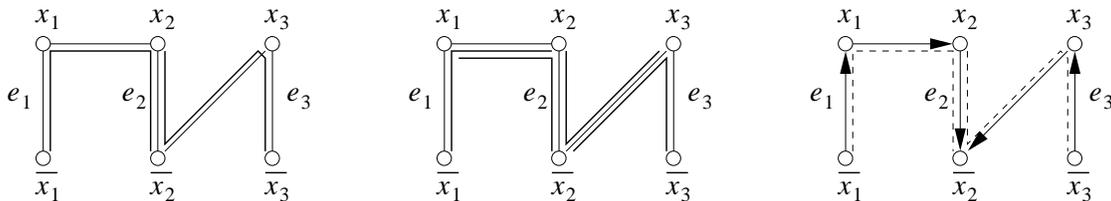


Figure 7: Network and paths resulting from the two clauses $x_1 \vee x_2, \overline{x_2} \vee x_3$: Constructed instance with length 3 paths (left), modified instance with length 2 paths (middle), and possible solution where both clauses are satisfied (right).

for at least one instance of MAX2SAT. With $Opt - m' \geq 3/4 \cdot m'$ (at least 3/4 of the clauses of any MAX2SAT instance can be satisfied) this yields

$$A_2 \leq \frac{4 + 3q}{7} \cdot Opt,$$

which gives the following theorem.

Theorem 5 *Unless $\mathcal{P} = \mathcal{NP}$, there is no approximation algorithm for MAXTOR with paths of length at most k that achieves ratio at least $\frac{4+3q}{7} \leq 0.981$ if $k = 2$ and ratio at least 0.955 if $k \geq 3$.*

In particular, with Theorem 4 we have that MAXTOR is \mathcal{APX} -complete for paths with constant maximum length.

6 Experimental results

In this section we discuss some experiments conducted with the approximation algorithm for MAXTOR presented in Section 5.2: the algorithm constructs a MAX2SAT instance from the given paths and uses an SDP-based algorithm to obtain an approximate solution. For our experiments we used the data that has been accumulated by Subramanian et al. [14] and is available on the WWW [1]. For three different dates (18 April 2001, 4 February 2002, and 6 April 2002) routing paths from 10, 9, respectively 14 autonomous systems were collected. From these 3.4 to 6.4 million paths in a network of about 10,000 nodes and 25,000 edges, the edge directions were deduced. It is notable that the path lengths are relatively short (see Table 1), i.e. the assumption in Section 5.2 that the path lengths are bounded by a small constant is quite realistic.

Table 1: Size of the network and number of paths at the three different dates. The maximal and average path lengths are given before (\max_1 and avg_1) and after (\max_2 and avg_2) the preprocessing described in the text.

Date	# of nodes	# of edges	# of paths	\max_1	avg_1	\max_2	avg_2
18 Apr 2001	10923	23935	3423422	12	3.52	10	1.70
4 Feb 2002	12788	27936	4988100	11	3.52	9	1.32
6 Apr 2002	13164	28510	6356435	11	3.55	8	1.39

In our implementation the paths are first preprocessed by directing edges that can be directed without conflicts. This shortens the paths considerably (cf. Table 1). Then it is checked with the ALLTOR algorithm described in Section 3 whether the graph can be oriented such that all paths are valid. This was not the case for any of the three dates. Finally, an approximate solution is calculated as described in Section 5.2: MAX2SAT is relaxed as in [9] and the rounding is done as in [7]. In [7] the improved approximation ratio could be achieved by adding new constraints and a modified rounding strategy. We could not add the constraints because the instance would have become too large. The new rounding strategy was adopted though. The freely available solver DSDP 4.5 [3] was used to solve the semidefinite programs. The overall results are given in Table 2. The computed orientations make roughly 98% of the given paths valid. This should be contrasted with the edge classifications computed by the heuristic algorithm of [14] that are also available at [1]. In these edge classifications, only about 80% of the paths are valid. This demonstrates that it pays off to use SDP-based approximation algorithms for MAXTOR in practice.

Table 2: Experimental results achieved with the approximation algorithm presented in Section 5.2.

Date	Total # of paths	# of valid paths
18 Apr 2001	3423422	3360926 (98.17%)
4 Feb 2002	4988100	4919310 (98.62%)
6 Apr 2002	6356435	6204849 (97.62%)

References

- [1] S. Agarwal. Data used in [14]. <http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy/>, 2002.
- [2] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation. Combinatorial Optimization Problems and their Approximability Properties*. Springer, Berlin, 1999.
- [3] S. Benson. Semidefinite programming solver DSDP 4.5. <http://www-unix.mcs.anl.gov/~benson/>, 2002.
- [4] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. In *Proceedings of INFOCOM'02*, June 2002.
- [5] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. Technical Report RT-DIA-73-2002, Dipartimento di Informatica e Automazione, Università di Roma Tre, 2002.
- [6] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of ACM SIGCOMM'99*, 1999.
- [7] U. Feige and M. Goemans. Approximating the value of two proper proof systems, with applications to MAX 2SAT and MAX DICUT. In *Proceedings of the Third Israel Symposium on Theory of Computing and Systems*, pages 182–189, 1995.
- [8] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9:733–745, 2000.

-
- [9] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995.
 - [10] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
 - [11] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
 - [12] M. Lewin, D. Livnat, and U. Zwick. Improved rounding techniques for the MAX 2-SAT and MAX DI-CUT problems. In *Integer Programming and Combinatorial Optimization (IPCO)*, LNCS 2337, pages 67–82, 2002.
 - [13] A. Medina, I. Matta, and J. Byers. On the origin of power laws in Internet topologies. *ACM Computer Communication Review*, 30(2):18–28, April 2000.
 - [14] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proceedings of INFOCOM'02*, June 2002.
 - [15] D. Vukadinović, P. Huang, and T. Erlebach. On the spectrum and structure of Internet topology graphs. In *Innovative Internet Computing Systems*, LNCS 2346, pages 83–95, June 2002.