

Thomas Erlebach (Ed.)

LNCS 4235

Combinatorial and Algorithmic Aspects of Networking

Third Workshop, CAAN 2006
Chester, UK, July 2006
Revised Papers



Springer


Thomas Erlebach (Ed.)

Combinatorial and Algorithmic Aspects of Networking

Third Workshop, CAAN 2006

Chester, UK, July 2, 2006

Revised Papers



Volume Editor

Thomas Erlebach
Department of Computer Science
University of Leicester
LE1 7RH, U.K.
E-mail: t.erlebach@mcs.le.ac.uk

Library of Congress Control Number: 2006937396

CR Subject Classification (1998): F.1.1, F.2.1-2, C.2, G.2.1-2, E.1

LNCS Sublibrary: SL 5 – Computer Communication Networks and
Telecommunications

ISSN 0302-9743
ISBN-10 3-540-48822-7 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-48822-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11922377 06/3142 5 4 3 2 1 0

Preface

The Internet, because of its size, decentralized nature, and loosely controlled architecture, provides a hotbed of challenges that are amenable to mathematical analysis and algorithmic techniques. The primary goal of the 3rd Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN 2006) was to bring together mathematicians, theoretical computer scientists and network specialists in this fast-growing area that is an intriguing intersection of computer science, graph theory, game theory, and networks. CAAN 2006 took place on July 2, 2006 in Chester, UK, co-located with the 13th Colloquium on Structural Information and Communication Complexity (SIROCCO 2006). The two previous CAAN workshops were held during August 6-7, 2004 at the Banff International Research Station, Alberta, Canada and on August 14, 2005 in Waterloo, Ontario, Canada.

In response to the call for papers we received 22 submissions. Each submission was reviewed by four referees. Based on the reviews, the Program Committee selected ten papers for presentation at the workshop. The workshop program also featured an invited talk by David Peleg. This volume contains the contributed papers and an abstract of the invited talk.

We would like to thank the Organizing Committee of SIROCCO 2006, in particular Christoph Ambühl, Catherine Atherton, Leszek Gąsieniec and Prudence Wong, for all the organizational help that made it easy for us to arrange CAAN together with SIROCCO. Furthermore, we are grateful to Andrei Voronkov for providing the EasyChair conference system, which we used to manage the electronic submissions, the review process, and the electronic program committee meeting. It simplified our task significantly. Finally, we thank the invited speaker, the authors of the contributed papers, and all participants of CAAN 2006 for helping to make the workshop a success.

August 2006

Thomas Erlebach
Program Chair
CAAN 2006

Organization

Steering Committee

Andrei Broder	Yahoo! Inc., USA
Angèle Hamel	Wilfrid Laurier University, Canada
Srinivasan Keshav	University of Waterloo, Canada
Alejandro López-Ortiz	University of Waterloo, Canada
Rajeev Motwani	Stanford University, USA
Ian Munro	University of Waterloo, Canada

Program Committee

Christoph Ambühl	University of Liverpool
Holger Bast	MPI Saarbrücken
Gruia Calinescu	Illinois Institute of Technology
Andrea Clementi	University of Rome “Tor Vergata”
Colin Cooper	King’s College London
Xiaotie Deng	City University of Hong Kong
Thomas Erlebach (Chair)	University of Leicester
Angèle Hamel	Wilfrid Laurier University
Samir Khuller	University of Maryland
Stavros Kolliopoulos	University of Athens
Danny Krizanc	Wesleyan University
Stefano Leonardi	University of Rome “La Sapienza”
Alejandro López-Ortiz	University of Waterloo
Christian Scheideler	TU München
Christian Schindelhauer	University of Freiburg
Angelika Steger	ETH Zürich
Csaba Tóth	Massachusetts Institute of Technology
Eli Upfal	Brown University

Additional Referees

Luca Becchetti	Yoo-Ah Kim	Maurizio Pizzonia
Rene Beier	Elias Koutsoupias	Guido Proietti
Tian-Ming Bu	Debapriyo Majumdar	Qi Qi
Hung Chim	Vangelis Markakis	Anuj Rawat
Miriam Di Ianni	Russell Martin	Gianluca Rossi
Angelo Fanelli	Julian Mestre	Wei Sun
Alexei Fishkin	Alfredo Navarra	Ingmar Weber
Leszek Gąsieniec	Francesco Pasquale	Prudence Wong
Min Jiang	Andrzej Pelc	Anders Yeo

Table of Contents

Invited Lecture

Recent Advances on Approximation Algorithms for Minimum Energy Range Assignment Problems in Ad-Hoc Wireless Networks	1
<i>David Peleg</i>	

Contributed Papers

The Price of Anarchy in Selfish Multicast Routing	5
<i>Andreas Baltz, Sandro Esquivel, Lasse Kliemann, Anand Srivastav</i>	
Designing a Truthful Mechanism for a Spanning Arborescence Bicriteria Problem	19
<i>Davide Bilò, Luciano Gualà, Guido Proietti</i>	
On the Topologies of Local Minimum Spanning Trees	31
<i>P.F. Cortese, G. Di Battista, F. Frati, L. Grilli, K.A. Lehmann, G. Liotta, M. Patrignani, I.G. Tollis, F. Trotta</i>	
Distributed Routing in Tree Networks with Few Landmarks	45
<i>Ioannis Z. Emiris, Euripides Markou, Aris Pagourtzis</i>	
Scheduling of a Smart Antenna: Capacitated Coloring of Unit Circular-Arc Graphs	58
<i>Guy Even, Shimon Shahar</i>	
On Minimizing the Number of ADMs - Tight Bounds for an Algorithm Without Preprocessing	72
<i>Michele Flammini, Mordechai Shalom, Shmuel Zaks</i>	
Tolerance Based Contract-or-Patch Heuristic for the Asymmetric TSP	86
<i>Boris Goldengorin, Gerold Jäger, Paul Molitor</i>	
Acyclic Type-of-Relationship Problems on the Internet	98
<i>Sven Kosub, Moritz G. Maaß, Hanjo Täubig</i>	

Minimum-Energy Broadcasting in Wireless Networks
in the d -Dimensional Euclidean Space (The $\alpha \leq d$ Case) 112
Andrzej Lingas, Mia Persson, Martin Wahlen

Optimal Gossiping with Unit Size Messages in Known Topology
Radio Networks 125
Fredrik Manne, Qin Xin

Author Index 135

Recent Advances on Approximation Algorithms for Minimum Energy Range Assignment Problems in Ad-Hoc Wireless Networks

David Peleg*

Department of Computer Science and Applied Mathematics,
The Weizmann Institute of Science, Rehovot 76100, Israel
`david.peleg@weizmann.ac.il`

Ad-hoc wireless networks have no wired infrastructure. Instead, they consist of a collection of radio stations $S = \{1, 2, \dots, n\}$ deployed in a given region and connected by wireless links. Each station is assigned a transmission range, and a station t can correctly receive the transmission of another station s if and only if t is within the range of s . The overall range assignment, $r : S \rightarrow R^+$, determines a (directed) transmission graph G_r . The transmission range of a station depends on the energy invested by the station. In particular, the power P_s required by a station s to correctly transmit data to another station t must satisfy the inequality $P_s \geq \text{dist}(s, t)^\alpha$, where $\text{dist}(s, t)$ is the Euclidean distance between s and t and $\alpha \geq 1$ is the *distance-power gradient*. The value of α may vary from 1 to more than 6 depending on the environment conditions at the location of the network (see [16]).

In order to allow an ad-hoc network to carry out certain basic communication paradigms, a fundamental design problem that needs to be solved is to calculate a transmission range assignment r such that (a) the corresponding transmission graph G_r satisfies a given connectivity property Π , and (b) the overall energy required to deploy the range assignment r is minimized. For any desired graph property Π , the resulting problem is denoted MIN-RANGE(Π).

We focus on two basic types of communication paradigms.

- *Broadcast* is a task initiated by a source station which has to disseminate a message to all stations in the wireless network. This task constitutes one of the main activities in real life multi-hop wireless networks [10].
- *Routing* is a task initiated by a source station which transmits a message intended to one particular destination station in the network.

To facilitate these two paradigms, the underlying transmission graph G_r is required to satisfy one of the following two properties, respectively.

B: Given a set of stations and a specific source station s , G_r has to contain a directed spanning tree rooted at s .

SC: Given a set of stations, G_r has to be *strongly connected*, i.e., contain a directed path from every station to every other station.

* Supported in part by a grant from the Israel Ministry of Science and Technology.

This characterization of the properties does not restrict the number of hops the communication might require. For quality of service purposes, it may be desirable to impose a bound h on the maximum number of hops in any communication path. This yields the following two variants.

$B[h]$: Given a set of stations and a specific source station s , G_r has to contain a directed spanning tree of depth at most h rooted at s .

$SC[h]$: Given a set of stations, G_r has to contain a directed path of at most h hops from every station to every other station.

We now review some known results on these problems. For broadcast, observe that if $\alpha = 1$, then the $\text{MIN-RANGE}(B)$ problem is solvable in polynomial time. Moreover, in the 1-dimensional case (i.e., when the stations are placed on a line), the problem is solvable in polynomial time for any $\alpha \geq 1$ [6]. For $d \geq 2$ and $\alpha > 1$, however, $\text{MIN-RANGE}(B)$ is NP-hard [5]. In [2,5] it is shown that whenever $\alpha \geq d$, the algorithm proposed in [10] based on constructing a minimum weight spanning tree achieves constant approximation. It is not known whether the problem admits a polynomial time approximation scheme.

Efficient solutions were given for $\text{MIN-RANGE}(B[h])$ when h is constant. In particular, a polynomial-time algorithm for $\text{MIN-RANGE}(B[h])$ for $h = 2$, based on a nontrivial dynamic program, is given in [1]. Moreover, the problem is given a polynomial-time approximation scheme for any fixed constant $h > 1$. For $\epsilon > 0$, the scheme has time complexity $O(n^\mu)$ where $\mu = O((\alpha 2^\alpha h^\alpha / \epsilon)^{\alpha^h})$.

For arbitrary h , an $O(hn^4)$ -time exact algorithm for $\text{MIN-RANGE}(B[h])$ on trees is presented in [18]. In addition they present a probabilistic $O(\log n \log \log n)$ approximation algorithm for $\text{MIN-RANGE}(B[h])$ on a general metric space. These results are improved to a ratio of $O(\log n)$ in [3,12] independently. The existence of a polynomial time approximation scheme (or even a polynomial time constant ratio approximation algorithm) for arbitrary h is not known.

Turning to the strong connectivity property SC , we first remark that in the one-dimensional case, i.e., when the stations are located on the real line, the problem is polynomial. An $O(n^4)$ -time algorithm for this problem is described in [13]. When the stations are spread in d -dimensional space ($d > 1$), finding an optimal solution for $\text{MIN-RANGE}(SC)$ is NP-hard [9,13], and moreover, it is APX-hard for $d \geq 3$ [13]. On the positive side, the problem has a 2-approximation algorithm based on constructing a minimum spanning tree [13].

Finally, consider the bounded-hop strong connectivity requirement $SC[h]$. It is known that $\text{MIN-RANGE}(SC[h])$ is NP-hard on general metric spaces for constant h [12]. For the 1-dimensional case where the stations of S are spread on the line, an $O(hn^3)$ -time 2-approximation algorithm for $\alpha = 2$ and any $h > 0$ is described in [7]. In higher dimensions, lower and upper bounds are shown in [8] on the optimal cost for any 2-dimensional instances with distance power gradient $\alpha \geq 1$, where h is an arbitrary constant. It is also shown therein that when S is a family of *well-spread* instances (namely, the locations in S are suitably distributed), the $\text{MIN-RANGE}(SC[h])$ problem on S admits a polynomial time approximation algorithm with constant ratio, i.e., $\text{MIN-RANGE}(SC[h])$ is

in APX. Additionally, it is shown that the $\text{MIN-RANGE}(SC[h])$ problem with a uniform instance probability is in the class Av-APX.

For arbitrary h , a polynomial time approximation algorithm of ratio $O(n^2)$ for $\text{MIN-RANGE}(SC[h])$ on trees and a randomized polynomial time approximation algorithm of ratio $O(n^2 \log n \log \log n)$ for the problem on general metric spaces were presented in [18]. This was improved in [3] to an approximation algorithm of ratio $O(\log n)$.

Finally, a polynomial time constant factor approximation algorithm for $\text{MIN-RANGE}(SC[h])$ on general metrics is given in [12]. The approximation ratio of the algorithm is $(1/(\sqrt[h]{2}-1))^\alpha (1+3^\alpha) (3^{\alpha+1})^{h-2}$. This is done by first considering a new variant of the classical *uncapacitated facility location* problem UFL [4,15,17], named the *hierarchical facility location* problem on metric powers, $\text{HFL}_\alpha[h]$. This problem involves a set F of locations that may open a facility, subsets D_1, D_2, \dots, D_{h-1} of locations that may open an intermediate transmission station and a set D_h of locations of clients. Each client in D_h must be serviced by an open transmission station in D_{h-1} and every open transmission station in D_l must be serviced by an open transmission station on the next lower level, D_{l-1} . An open transmission station on the first level, D_1 must be serviced by an open facility. A cost c_{ij} is associated with assigning a station j on level $l \geq 1$ to a station i on level $l-1$. Also, for $i \in F$, a cost f_i is associated with opening a facility at location i . It is required to find a feasible assignment that minimizes the total cost.

A polynomial time constant ratio approximation algorithm is then established for the $\text{HFL}_\alpha[h]$ problem, by solving a linear relaxation of the corresponding integer linear program and then using the filtering and rounding technique of [14,17]. Finally, the $\text{MIN-RANGE}(SC[h])$ problem is reduced to $\text{HFL}_\alpha[h]$, so that the constant approximation algorithm for $\text{HFL}_\alpha[h]$ yields also an approximate solution for $\text{MIN-RANGE}(SC[h])$.

References

1. C. Ambühl, A.E.F. Clementi, M. Di Ianni, N. Lev-Tov, A. Monti, D. Peleg, G. Rossi and R. Silvestri. Efficient Algorithms for Low-Energy Bounded-Hop Broadcast in Ad-Hoc Wireless Networks. In *Proc. 21st Symp. on Theoretical aspects of Computer Science*, pages 418–427, 2004.
2. G. Călinescu, X.Y. Li, O. Frieder, and P.J. Wan. Minimum-Energy Broadcast Routing in Static Ad Hoc Wireless Networks. In *Proc. 20th INFOCOM*, pages 1162–1171, 2001.
3. J. Chlebikova, D. Ye and H. Zhang. Assign Ranges in General Ad-Hoc Networks. In *Proc. 1st Conf. on Algorithmic Applications in Management*, Xian, China, pages 411–421, 2005.
4. F.A. Chudak. Improved approximation algorithm for uncapacitated facility location problem. In *Proc. 6th Conf. on Integer Programming and Combinatorial Optimization*, pages 180–194, 1998.
5. A.E.F. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the Complexity of Computing Minimum Energy Consumption Broadcast Subgraphs. In *Proc. 18th Symp. on Theoretical Aspects of Computer Science*, pages 121–131, 2001.

6. A.E.F. Clementi, M. Di Ianni, and R. Silvestri. The Minimum Broadcast Range Assignment Problem on Linear Multi-Hop Wireless Networks. *Theoretical Computer Science* 299, (2003), 751–761.
7. A. Clementi, A. Ferreira, P. Penna, S. Perennes, and R. Silvestri. The minimum range assignment problem on linear radio networks. In *Proc. 8th European Symp. on Algorithms*, pages 143–154. 2000.
8. A. Clementi, A. Ferreira, P. Penna, S. Perennes, and R. Silvestri. The power range assignment problem in radio networks on the plane. In *Proc. 17th Symp. on Theoretical Aspects of Computer Science*, pages 651–660, 2000.
9. A.E.F. Clementi, P. Penna, and R. Silvestri. Hardness results for the power range assignment problem in packet radio networks. In *Proc. 2nd Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 197–208, 1999.
10. A. Ephremides, G.D. Nguyen, and J.E. Wieselthier. On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks. In *Proc. 19th INFOCOM*, pages 585–594, 2000.
11. J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proc. 35th ACM Symp. on Theory of Computing*, pages 448–455, 2003.
12. E. Kantor and D. Peleg. Approximate Hierarchical Facility Location and Applications to the Shallow Steiner Tree and Range Assignment Problems. *Proc. 6th Conf. on Algorithms and Complexity*, pages 211–222, 2006.
13. L. M. Kiousis, E. Kranakis, D. Krizanc, and A. Pelc. Power Consumption in Packet Radio Networks. *Theoretical Computer Science* 243, (2000), 289–305.
14. J.H. Lin and J.S. Vitter. ϵ -approximations with small packing constraint violation. In *Proc. 24th ACM Symp. on Theory of Computing*, pages 771–782, 1992.
15. M. Mahdian, Y. Ye, and J. Zhang. A 1.52-approximation algorithm for the uncapacitated facility location problem. In *Proc. 5th Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 229–242, 2002.
16. K. Pahlavan and A. Levesque. *Wireless information networks*. Wiley-Interscience, 1995.
17. B.D. Shmoys, E. Tardos, and Aardal K. Approximation algorithms for facility location problems. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 265–274, 1997.
18. D. Ye and H. Zhang. The range assignment problem in static ad-hoc networks on metric spaces. In *Proc. 11th Colloq. on Structural Information and Communication Complexity*, pages 291–302, 2004.

The Price of Anarchy in Selfish Multicast Routing

(Extended Abstract)

Andreas Baltz*, Sandro Esquivel, Lasse Kliemann*, and Anand Srivastav

Institut für Informatik, CAU Kiel

Christian-Albrechts-Platz 4

24118 Kiel

{aba, sae, lki, asr}@numerik.uni-kiel.de

Abstract. We study the price of anarchy for selfish multicast routing games in directed multigraphs with latency functions on the edges, extending the known theory for the unicast situation, and exhibiting new phenomena not present in the unicast model. In the multicast model we have N commodities (or player classes), where for each $i = 1, \dots, N$, a flow from a source s_i to a finite number of terminals $t_i^1, \dots, t_i^{k_i}$ has to be routed such that every terminal t_i^j receives flow $n_i \in \mathbb{R}_{\geq 0}$.

One of the significant results of this paper are upper and lower bounds on the price of anarchy for edge latencies being polynomials of degree at most p with non-negative coefficients. We show an upper bound of $(p+1) \cdot \frac{\nu^{p+1}}{\nu^*}$ in some variants of multicast routing. We also prove a lower bound of ν^p , so we have upper and lower bounds that are tight up to a factor of $(p+1)\nu$. Here, ν and ν^* are network and strategy dependent parameters reflecting the maximum/minimum consumption of the network. Both are 1 in the unicast case. Our lower bound of ν^p , where in the general situation we have $\nu > 1$, shows an exponential increase compared to the Roughgarden bound of $O(p/\ln p)$ for the unicast model. This exhibits the contrast to the unicast case, where we have Roughgarden's (2002) result that the price of anarchy is independent of the network topology. To our knowledge this paper is the first thorough study of the price of anarchy in the multicast scenario. The approach may lead to further research extending game-theoretic network analysis to models used in applications.

1 Introduction

Multicast routing in communication networks is a natural and practically relevant extension of the so far quite well studied unicast routing. Among the applications of multicast routing are the transmission of music, movies, conferences, or any other popular content, that is requested by several customers at a time. A formal description of our multicast routing model needs many technical definitions. We keep the introduction on a more informal level and refer the reader to Sect. 2 for all necessary details.

* Supported by Deutsche Forschungsgemeinschaft.

Problem Formulation. An instance of selfish multicast routing consists of a directed multigraph $G = (V, E)$, where the edges are also called links, a set of N player classes, called commodities, where commodity i is characterized by a source s_i and terminals (or sinks) $t_i^1, \dots, t_i^{k_i}$, and a (flow) demand of $n_i \in \mathbb{R}_{\geq 0}$. The links are each equipped with a latency function $l_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. For commodity i , a set $S = \{P_1, \dots, P_{k_i}\}$ where P_j is an s_i - t_i^j -path, is called a strategy. The task is to realize for every commodity i , a flow in the network from s_i to all terminals $t_i^1, \dots, t_i^{k_i}$, satisfying the demand n_i for every terminal. We think of the demand as being under control of infinitely many players, each controlling a negligible fraction and selfishly trying to find the fastest route for it. This game-theoretic model is known as the Wardrop model.

In the unicast model, $k_i = 1$ for all i , so we have a collection of single source/single sink commodities, and every strategy S consists of one path only. In the multicast case there are two different ways to route the flow $f(S)$ assigned to a strategy S for commodity i : either, we route $f(S)$ on each path, which is the usual notion of flows satisfying the Kirchhoff conservation law (here shortly called *conservation flow*), or we allow multiple duplication of flow at certain nodes: a link which serves several, say r , terminals in a strategy, i.e., a link contained in r paths of that strategy, only needs to transmit the data once, not r times. That is because the data can later be duplicated to serve all terminals. In this way, the congestion on the links can be reduced. We call such a flow *duplication flow*.

The cost of a flow is defined by $SC(f) = \sum_{S \in \mathfrak{S}} l_S(f) f(S)$, where \mathfrak{S} is the set of all strategies of all commodities, $f(S) \leq n_i$ is the portion of the demand that by the decision of the selfish players has been allocated to strategy S , and $l_S(f)$ is the strategy latency for S . We study four different definitions for l_S , which all coincide in the unicast case. Together with the two types of flows (conservation and duplication), we thus have 8 variants of multicast. The price of anarchy for a multicast instance \mathcal{I} is $\rho(\mathcal{I}) = \sup_f \frac{SC(f)}{SC(f^*)}$, where f ranges over all Nash equilibria and f^* is an optimal flow. A Nash equilibrium is a flow in which no player (meaning: no portion of the flow, however small) has an incentive to unilaterally deviate from his current strategy.

Previous and Related Work. By the pioneering work of Roughgarden [1] and Roughgarden and Tardos [2] we know that $\rho(\mathcal{I})$ in the unicast model for latency functions being polynomials of degree p , is bounded from above by $O(\frac{p}{\ln p})$ (and is $\frac{4}{3}$ for $p = 1$). As already an example of a 2-parallel links network has a price of anarchy of $\frac{4}{3}$ (for $p = 1$), the surprising conclusion is that it is *independent* of the network topology from a worst-case point of view [1, Sec. 3.4].

Our Results. A solid foundation for the analysis of multicast routing games is given in Sect. 2. In Sect. 2.1 we introduce a concise model, and in Sect. 2.2 we show, using results on variational inequalities, the existence of Nash equilibria.

As a main result, we show in Sect. 3 that the price of anarchy in multicast routing may depend heavily on the network topology and the strategies. Certain

edges of the graph may be utilized under certain strategies more than others, although the players on those strategies are not charged for this. On the other hand, some strategies may depend highly on some edges but only contribute a small amount to their utilization. See Remark 2 for a more detailed discussion of this. To capture the effects of this phenomenon, which does not occur in unicast routing, we introduce for each edge and strategy an integer called the *consumption*. Moreover we introduce two new invariants for a graph G and a set of strategies \mathfrak{S} , which we call *maximum (resp. minimum) consumption number*, $\nu = \nu(G, \mathfrak{S})$ resp. $\nu^* = \nu^*(G, \mathfrak{S})$. We have $\nu = 1 = \nu^*$ in unicast routing. We show in Sect. 3.3 that in two variants of multicast for polynomial latency functions of degree p (where we will always assume non-negative coefficients), the price of anarchy is at most $(p+1)^{\frac{\nu^{p+1}}{\nu^*}}$ and in Sect. 3.2 provide a lower bound for one of these variants of ν^p (with $\nu^* = 1$). So, we have here a gap of $(p+1)\nu$.

We then present (also in Sect. 3.2) a multicast instance with price of anarchy at least ν^p . As in general $\nu > 1$, the ν^p bound is exponentially larger than the corresponding unicast bound of $O(\frac{p}{\ln p})$. This is surprising (and disappointing from the point of view of a company running the network). For instances using the advantages of duplication flows in order to de-load high-latency links, the cost of the global optimum decreases drastically, but unfortunately, due to selfish behavior, the users grab (greedily) certain links without a look ahead and block them out, so that the cost of the Nash equilibrium still stands high. For other definitions of strategy latency, in Sect. 4 we are able to prove that results from non-atomic congestion games, i.e., bounds of the form $O(\frac{p}{\ln p})$, carry over.

Open Problems. A couple of interesting open problems arise from this paper. For example, can the $(p+1)\nu$ factor gap between upper and lower bound be closed? Can exponentially high prices of anarchy be reduced by taxation schemes? It would also be interesting to consider polynomial time algorithms for the computation of equilibria.

An ambitious task would be to study multicast for information flows with duplication *and* coding facilities of the network. Such networks are the state-of-the-art in today's engineering designs. Our work can be considered as a first step in this direction.

2 Basics of Multicast Routing

2.1 Model and Instances

An instance of selfish multicast routing consists of the following.

- A directed multigraph $G = (V, E)$. The edges are also called *links*.
- A set of N *player classes* (or *user classes*). Sometimes, player classes are also called *commodities*. Each player class is characterized by a demand n_i and a vector of vertices $(s_i; t_i^1, \dots, t_i^{k_i})$, where s_i is the source, and the $t_i^1, \dots, t_i^{k_i}$ are the terminals.

- The demand n_i is supposed to be routed from s_i to each of the terminals $t_i^1, \dots, t_i^{k_i}$. We think of the demand as being under control of infinitely many players, each of them controlling a negligible amount of it. This is the well-known Wardrop model (see, e.g., [1, Sec. 2.2]), which will become clearer when we define flows and Nash equilibria below.

- Each link $e \in E$ in the graph is equipped with a latency function $l_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. We always assume each latency function to be non-decreasing and *standard* [1]. This means that it is differentiable and $\xi \mapsto l_e(\xi)\xi$ is convex.

If an amount ξ of traffic is to be routed through the link e , each unit of flow will take $l_e(\xi)$ time to traverse e . Hence we have a total latency of $l_e(\xi)\xi$ on that link.

- For $i \in [N] = \{1, \dots, N\}$, we call a set of paths $S := \{P_1, \dots, P_{k_i}\}$ where P_j is a path connecting s_i with t_i^j for $j \in [k_i]$, a *strategy*. Note that for unicast routing $k_i = 1$ for all i . The set of all strategies we wish to allow for player class i is denoted by \mathfrak{S}_i . We assume¹ that $\mathfrak{S}_i \cap \mathfrak{S}_j = \emptyset$ for all $i, j \in [N]$. Let $\mathfrak{S} := \bigcup_{i \in [N]} \mathfrak{S}_i$.

- An *action distribution* (according to [2]), simply called *flow*, is a map $f : \mathfrak{S} \rightarrow \mathbb{R}_{\geq 0}$ such that all the demands are met, i.e., $\sum_{S \in \mathfrak{S}_i} f(S) = n_i \quad \forall i \in [N]$. A flow can be understood as a partition of each of the real intervals $[0, n_i]$. Each of these intervals represents the continuum of infinitely many players of the corresponding player class. The quantity $f(S)$ gives, for each $S \in \mathfrak{S}_i$, the portion of demand that by the decision of the players from that class is routed according to that particular strategy S .

As described in the introduction, the routing of a flow in the multicast model can be done in two different ways: we can route the demand with flows in the usual sense (conservation flows) or with flows allowing duplication (duplication flows).

- Let $e \in E$ and $S \in \mathfrak{S}$. We define the *consumption* of e under S as $c(e, S) := |\{P \in S; e \in P\}|$, i.e., the consumption is the number of paths in S traversing e , or in other words, the number of terminals served via e in this strategy.

- The *congestion* f_e of a link e with respect to a flow f is the amount of traffic that link e has to process. The total latency of a link e hence is $l_e(f_e)f_e$. Each instance defines the congestion in one of the following ways, depending on whether we have conservation flows or duplication flows.

$$f_e := \begin{cases} \sum_{S \in \mathfrak{S}(e)} c(e, S) f(S) & \text{conservation flows} \\ \sum_{S \in \mathfrak{S}(e)} f(S) & \text{duplication flows} \end{cases} \quad (1)$$

Here, $\mathfrak{S}(e)$ denotes the set of all strategies that contain a path which in turn contains e .

- We denote by $l_S(f)$ the so-called latency of strategy S with respect to a flow f . In unicast $l_S(f)$ is simply the sum of the latencies in the single path of which

¹ Otherwise we have to treat \mathfrak{S} as a multiset.

the strategy S consists². Let $S = \{P_1, \dots, P_{k_i}\} \in \mathfrak{S}$. As in [1], the latency of a path P under f is defined by

$$l_P(f) := \sum_{e \in P} l_e(f_e).$$

By $E(S)$ we denote the union of the edges in all the paths in S . Note that we consider $E(S)$ *not* as a multiset, so edges do not appear multiple times even if they lie in several paths.

We introduce the following four definitions of the latency of a strategy S .

$$\begin{aligned} l_S^{\text{edges}}(f) &:= \sum_{e \in E(S)} l_e(f_e) , & l_S^{\text{paths}}(f) &:= \sum_{P \in S} l_P(f) \\ l_S^{\text{paths avg}}(f) &:= \frac{1}{|S|} l_S^{\text{paths}}(f) , & l_S^{\text{max}}(f) &:= \max_{P \in S} l_P(f) \end{aligned} \quad (2)$$

An instance of selfish multicast routing includes one of these strategy latency functions.

The latency of a strategy $l_S(f)$ is the latency that all players experience who choose strategy S . It can hence be thought of as a kind of equivalent to what is known as utility or payoff function in other game-theoretic settings.

Remark 1. 1. For unicast routing, all four definitions in (2) coincide.

2. It is easy to see that $l_S^{\text{paths}}(f) = \sum_{e \in E(S)} c(e, S) l_e(f_e)$.

2.2 Nash Equilibria, Social Cost, Price of Anarchy

A flow f is called a *Nash equilibrium* (sometimes abbreviated *NE*), if

$$f(S_1) > 0 \implies l_{S_1}(f) \leq l_{S_2}(f) \quad \forall S_1, S_2 \in \mathfrak{S}_i \quad \forall i \in [N] . \quad (3)$$

Hence, in a Nash equilibrium, only minimum-latency strategies are used, since then no player has an incentive to choose a different strategy (provided the rest of the players keep their current decision). If each l_S is continuous (which will be the case during all our studies), then the game admits at least one Nash equilibrium. This follows from the characterization of Nash equilibria as the solutions to a certain variational inequality (see Thm. 1 and the discussion after that).

We define the *social cost* of a flow f as

$$\text{SC}(f) := \sum_{S \in \mathfrak{S}} l_S(f) f(S) .$$

The social cost captures the overall performance of the system for a given flow f . We will always assume that our instances admit a flow f^* with minimum social cost and that $\text{SC}(f^*) > 0$. Existence is guaranteed if all l_S are continuous

² The maximum over all links in the path has also been studied [3].

(which will be the case in our studies), because the set of flows is compact. For an instance \mathcal{I} of selfish multicast routing with optimal flow f^* , define the price of anarchy by

$$\rho(\mathcal{I}) := \sup_{f \text{ is NE}} \frac{SC(f)}{SC(f^*)} .$$

Nash equilibria have a very simple structure, as seen in the following proposition. The proof for this is straightforward.

Proposition 1. *Let f be a Nash equilibrium. Then, for every $i \in [N]$, there exists a real number $l_i(f)$ such that $l_S(f) = l_i(f)$ for all $S \in \mathfrak{S}_i$, whenever $f(S) > 0$, and no strategy in \mathfrak{S}_i has latency less than $l_i(f)$.*

Corollary 1. *Let f be a Nash equilibrium. Then $SC(f) = \sum_{i \in [N]} l_i(f) n_i$.*

We now aim for further characterizations of Nash equilibria. Let f, \tilde{f} be flows. Define $SC^f(\tilde{f}) := \sum_{S \in \mathfrak{S}} l_S(f) \tilde{f}(S)$. The first part of the following theorem is well-known for the unicast case, see, e.g., [1, Lem. 3.3.7] or [4] and the references therein. The whole theorem also holds in a more general context than multicast routing, for it (and its proof) does not require the notion of congestion.

Theorem 1. 1. *Let f be a flow. Then f is a Nash equilibrium if and only if we have*

$$SC^f(\tilde{f}) \geq SC(f) \quad \text{for all flows } \tilde{f} . \quad (4)$$

2. *Let each l_S be continuous. Then the multicast game admits at least one Nash equilibrium.*

Proof. We refer the reader to the full version of this paper for the proof of 1). For 2) note that (4) is equivalent to $\sum_{S \in \mathfrak{S}} l_S(f) (\tilde{f}(S) - f(S)) \geq 0$ for all flows \tilde{f} . This is a well-studied variational inequality. It has been shown in [5] with deep results from the index theory of vector fields that it admits at least one solution. \square

Note that all strategy latencies from (2) are continuous, because we only consider standard latency functions and because the congestion is a continuous mapping. Hence, all our multicast games admit at least one Nash equilibrium.

In the rest of the paper we investigate the price of anarchy for conservation flows resp. duplication flows and the four strategy latencies from (2). These are 8 cases. In Sect. 3 we show for 5 of them that the price of anarchy depends on the network topology, while in two other cases it does not (Sect. 4).

3 Price of Anarchy Dependent on the Network Topology

For a directed graph G and a set of strategies \mathfrak{S} we define the *minimum and maximum consumption number* as

$$\nu^*(G, \mathfrak{S}) := \min_{S \in \mathfrak{S}} \min_{e \in E(S)} c(e, S) , \quad \nu(G, \mathfrak{S}) := \max_{S \in \mathfrak{S}} \max_{e \in E(S)} c(e, S) .$$