Accepted Manuscript

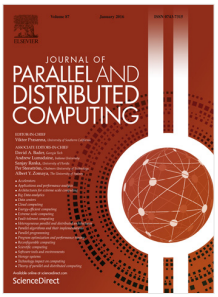

Approximate and incremental network function placement

Tamás Lukovszki, Matthias Rost, Stefan Schmid

PII: S0743-7315(18)30427-1
DOI: https://doi.org/10.1016/j.jpdc.2018.06.006
Reference: YJPDC 3898

To appear in: *J. Parallel Distrib. Comput.*

Received date : 20 June 2017
Revised date : 13 November 2017
Accepted date : 13 June 2018

Please cite this article as: T. Lukovszki, M. Rost, S. Schmid, Approximate and incremental network function placement, *J. Parallel Distrib. Comput.* (2018), https://doi.org/10.1016/j.jpdc.2018.06.006

Approximate and Incremental Network Function Placement

Tamás Lukovszki[a], Matthias Rost[b], Stefan Schmid[c]

[a] *Eötvös Loránd University, Hungary*
[b] *Technische Universität Berlin, Germany*
[c] *Aalborg University, Denmark*

Abstract

The virtualization of modern computer networks introduces interesting new opportunities for a more flexible placement of network functions and middleboxes (firewalls, proxies, traffic optimizers, virtual switches, etc.). This paper studies approximation algorithms for the incremental deployment of a minimum number of middleboxes, such that capacity constraints at the middleboxes and length constraints on the communication routes are respected.

Based on a new, purely combinatorial and rigorous proof of submodularity, we obtain our main result: a deterministic greedy approximation algorithm which only employs augmenting paths to serve future communication pairs. Hence, our algorithm does not require any changes to the locations of existing middleboxes or the preemption of previously served communication pairs when additional middleboxes are deployed. It is hence particularly attractive for incremental deployments. We prove that the achieved polynomial-time approximation bound is optimal, unless $P = NP$ holds.

This paper also initiates the study of a weighted problem variant, in which entire groups of nodes need to communicate via a middlebox, possibly at different rates. We present an LP relaxation and randomized rounding algorithm for this problem, leveraging an interesting connection to scheduling.

We complement our formal results with a simulation study of a large set of synthetically generated instances. Our results indicate that the presented algorithms yield near-optimal solutions in practice.

Email addresses: lukovszki@inf.elte.hu (Tamás Lukovszki), mrost@inet.tu-berlin.de (Matthias Rost), schmiste@cs.aau.dk (Stefan Schmid)