Deadline-aware Scheduling of Coflows in Datacenter Networks

Ph.D. Advisors : Olivier Brun (http://homepages.laas.fr/brun/) and Balakrishna Prabhu (https://homepages.laas.fr/bala/), LAAS-CNRS, Toulouse, France

Collaboration : Rachid El Azouzi and Francesco De Pellegrini, LIA, Avignon, France.

Keywords: coflows, scheduling, admission control, datacenters

Candidate profile: the position is intended for a highly motivated applicant holding a master degree in Computer Science, Applied Mathematics or a related field. Key requirement is proficiency in mathematical programming, approximation algorithms and operations research; Knowledge in the domain of communication networks and strong programming skills are considered a plus. Written communication skills (English) is a prerequisite. Applying candidates will prove and/or justify the requested knowledge and skills by providing by **June 5th 2022:**

- A full curriculum vitae.
- A motivation letter consistent with the proposed PhD project.
- All documents attesting the requested skills and knowledge.
- Academic records and marks of the two years of Master, or of the two last years of Engineer school.
- One or two recommendation letters are considered a plus.

Required skills: the ideal candidate should have solid background in applied computer science and applied mathematics. Her/his curriculum should prove excellent scholar records in relevant subjects such as, e.g., operations research, control theory, probability theory and applications to queueing theory, performance evaluation, mathematical programming and algorithmic design. Familiarity with networking technologies used in datacenter networks is considered a plus but not mandatory.

Detailed description: The SARA team (http://www.laas.fr/SARA-EN/) at LAAS-CNRS (http://www.laas.fr/), Toulouse, France, is seeking a highly motivated Ph.D. candidate in the field of traffic engineering for future datacenter networks. The candidate will work on the optimization of scheduling of coflows. The activity of the candidate will involve the modelling of the problem, its mathematical formulation for the resulting resources optimization problems and the development of efficient solutions by means of exact methods and approximation algorithms. The focus of this thesis requires a background in mathematical modeling, distributed optimization, performance evaluation and algorithmics applied to computer networks. Interested candidates who want to have further details on the thesis' subject and workplace laboratory, can contact directly Olivier Brun (brun@laas.fr) or Balakrishna Prabhu (balakrishna.prabhu@laas.fr)

Research Domain: Operational Research and Communication Networks

Objective. This Ph.D. position is aimed at the development of advanced algorithms and optimization frameworks to improve the deadline satisfaction rate of coflows generated at runtime. Overall, this PhD position is to address the following key objectives:

- A modelling framework for the minimization of the (weighted) number of late coflows in datacenters.
- Mathematical modelling and Operation Research techniques to provide efficient offline solutions.
- Design of efficient online algorithms for joint coflow admission control and scheduling.

Context and challenges: the focus of this thesis is the development of new mathematical models for the scheduling of time-sensitive coflows in datacenters, where the core issue is the presence of competing coflows, thus requiring optimizing admission control and scheduling decisions. Roughly speaking, at the network level a coflow is any series of flows that have common sets of sources and destinations. At the application level, coflows are generated under various frameworks for distributed computation, e.g., in Hadoop Map-Reduce, where consecutive computation stages are interleaved with communication, creating swarms of flows across server racks within the datacenter traffic. Coflows represent very specific traffic sources requiring to rethink the flow scheduling logic in order to be able to allocate flow routes over the datacenter fabric towards a global optimization objective. At the same time, several technological constraints are to be accounted for with respect to the datacenters' network fabric, e.g., with respect to the network topology and in terms of computing capabilities of the fabric switches. The classic problem of scheduling flows in datacenter networks has been described in several works appeared in the literature [1][2][3]. The customary performance metric for coflow scheduling is the makespan or the weighted coflow completion time (CCT). Minimizing the average CCT is an appropriate goal in order to increase the number of computing jobs dispatched per hour in a datacenter. The weighted CCT minimization has thus been addressed in several works, both in the deterministic [1-4] and stochastic settings [5]. A decade's research on the problem has shed light on its complexity and several algorithmic solutions have been devised. However, the context changes radically in the case of time-sensitive jobs, where the data transfer phase may be subject to strict coflow deadlines. In this thesis, we propose to advance the state of the art by designing online solutions for joint admission control and scheduling of coflows so as to minimize the number of late coflows. The research work will build on our recent results on deadline-aware coflow scheduling [6] and will both seek to improve existing solutions and extend their scope to more general settings. In particular, the non-clairvoyant case when the volumes of coflows are not known precisely in advance is an interesting extension. Another interesting research direction is the joint placement and scheduling problem when the coflows can be sent to one of many egress ports. The problem now allows to also determine the egress ports for the coflows in order to maximize the weighted deadline accepted rate.

Bibliography:

[1] M. Chowdhury, Y. Zhong, and I. Stoica, *"Efficient Coflow Scheduling with Varys",* ACM SIGCOMM, 2014.
[2] Y. Zhao et al. RAPIER: Integrating Routing and Scheduling for Coflow-aware Data Center Networks in Proc. of IEEE INFOCOM April 2015
[3] M. Chowdhury, Y. Zhong, and I. Stoica, "Efficient coflow scheduling with varys," in ACM SIGCOMM Computer Communication Review, vol. 44, no. 4. ACM, 2014, pp. 443–454.
[4] M. Shafiee and J. Ghaderi, "Scheduling coflows in datacenter networks: Improved bound for total weighted completion time," ACM SIGMETRICS 2017.
[5] R. Mao, V. Aggarwal, and M. Chiang, "Stochastic non-preemptive co-flow scheduling with time-indexed relaxation," IEEE INFOCOM, DCPerf Workshop, 2018.
[6] Q.-T. Luu, O. Brun, R. El-Azouzi, F. De Pellegrini, B. J. Prabhu, and C. Richier. DCoflow:Deadline-aware scheduling algorithm for coflows in datacenter networks. To appear in IFIP Networking 2022, May 2022.URL:https://arxiv.org/abs/2205.01229