COMPUTER AND CONTROL ENGINEERING

MUR DM 117/Aruba - Next Generation Software for Datacenter Networking

Funded By	MINISTERO DELL'UNIVERSITA' E DELLA RICERCA [P.iva/CF:97429780584] Politecnico di TORINO [P.iva/CF:00518460019] ARUBA S.P.A. [P.iva/CF:01573850516]
Supervisor	RISSO FULVIO GIOVANNI OTTAVIO - fulvio.risso@polito.it
Contact	Marco Mangiulli
Context of the research activity	This project aims at investigating the problem of efficient software-based network services in medium-large datacenters, in particular with respect to (a) novel kernel-based paradigms (e.g., eBPF, DPDK), (b) their potential acceleration with special-purpose hardware (e.g. Smartnic). Furthermore, the integration of special-purpose hardware in future generations of server-type processors will be investigated, with the aim of accelerating and making the aforementioned services faster and more energy efficient. Progetto finanziato nell'ambito del PNRR - DM 117/2023 - CUP E14D23002020004
	Research objectives The research goal of improved efficiency of software-oriented data centers is expected to be achieved by pursuing the following objectives: • Evaluate the possible performance improvements when leveraging multiple technologies for packet processing, such as DPDK, AF_XDP, and eBFP/XDP. • Assess the possible efficiency improvements when running network processing software on SmartNICs. • Evaluate the network functions currently required by modern datacenter orchestrators, looking for potential optimization opportunities. • Explore novel parallelization approaches that can allow to scale horizontally (multiple CPU cores exploited in parallel) and vertically (split the software across multiple software modules executed in sequence, on different servers). • All the above objectives will require the design of dedicated processing algorithms, as well as proper tools for validation. All the above studies will be validated by means of a set of network services already running in current datacenters, executed in realistic conditions.

Objectives	Outline of the research work plan • First year: the Ph.D. student will review the state of the art regarding software-based network processing, with a focus on technologies whose support was recently introduced in the Linux operating system. A conference publication is expected to be produced based on the results of the review. Afterward, the candidate will explore the performance of technologies such as DPDK, AF_XDP, and eBFP/XDP, providing the ground for possible optimizations. • Second year: the Ph.D. will develop a model for performance prediction for DPDK, AF_XDP, and eBFP/XDP software (which is expected to be published in a paper), then he/she will leverage the above model to determine the best technology to be used in each running condition, based on static (e.g., software configuration) and dynamic (actual traffic to be processed) data. This is expected to originate a third paper. • Third year: the Ph.D. will pursue two directions. First, analyze the possible speed-up of SmartNICs when applied to the field of investigation. Second, the extension of the achieved results in a small data center scenario, such as the ones operated by datacenter operators, in order to leverage the power of cloud-native functions. The above activities are expected to be published in two separate papers. Expected target publications Top conferences: • USENIX Symposium on Operating Systems Design and Implementation (OSDI) • USENIX Symposium on Networked Systems Design and Implementation (NSDI) • International Conference on Computer Communications (INFOCOM) • IEEE conference on Network Softwarization (Netsoft) Journals: • IEEE/ACM Transactions on Networking • IEEE Transactions on Computers • ACM Transactions on Computer Systems (TOCS) Magazines: • IEEE Computer, IEEE Networks

Skills and competencies for the development of the activity	The ideal candidate has good knowledge and experience in cloud computing and networking. Availability for spending periods abroad would be preferred for a more profitable investigation of the research topic.