Ph.D. in Information Technology
Thesis Defenses

January 13th, 2025
At 11:00 p.m.
Sala Conferenze Emilio Gatti, Edificio 20

Nicola DI CICCO – XXXVII Cycle

MACHINE LEARNING AS A NETWORK MANAGEMENT PRIMITIVE: FROM END-TO-END OPTIMIZATION TO ATOMIC NETWORK FUNCTIONS
Supervisor: Prof. Massimo Tornatore

Abstract:

Machine Learning (ML) is rapidly becoming the jack-of-all-trades of the network management stack. Thanks to its purely data-driven nature, ML can be leveraged for developing a broad set of network management functions tailored to the specific task at hand. Unfortunately, we are still far from harnessing the full potential of ML for network management. Contemporary literature in this field has made tremendous progress in identifying algorithms and design principles for producing well-trained models from historical data. These contributions, though fundamental, cover only one aspect of the whole ML model lifecycle. Deploying, using, and maintaining ML-based network management solutions poses nontrivial research challenges, orthogonal with respect to training a model, that is presently not yet thoroughly addressed. This hinders a widespread adoption of ML by network operators, who instead prefer relying upon classical, battle-tested network management solutions. To take a step toward solving this fundamental problem, this Thesis identifies five fundamental challenges in ML for network management: 1) Generalizability to data beyond training, 2) Adaptability to dynamic network environments, 3) Reliability, in terms of providing theoretical performance guarantees after model deployment, 4) Data efficiency, for minimizing the amount of labor required for training and updating models, and 5) Performance, i.e., ML-based solutions must tangibly improve over conventional methods to be worth considering. We address these challenges through multiple representative network management applications: online and offline network optimization with Reinforcement Learning, focusing on generalizability and performance; ML-based hardware fault classification in microwave networks, focusing on data efficiency and reliability; Continual in-network ML, focusing on data efficiency and adaptability; and intent-based networking with Large Language Models, focusing on performance and data-efficiency. We quantitatively validate the practical effectiveness of our proposed solutions through extensive comparisons against the state of the art, and by leveraging novel real-world datasets, which we make publicly available.

Qiaolun ZHANG – XXXVII Cycle

GREEN, RESILIENT, AND SECURE NEXT-GENERATION OPTICAL NETWORKS
Supervisor: Prof. Massimo Tornatore

Abstract:

Optical networks represent the essential backbone for various communication systems, such as long-haul, metro, and data center networks. As future optical networks will have to handle a dramatic increase in the volume of highly sensitive data, it is crucial for future optical networks to energy-efficiently accommodate increasing traffic demands as well as to be resilient against network failures and secure against attackers. Specifically, as current communication systems already contribute 2% to 3% of global energy consumption, and with this figure expected to rise due to increasing traffic demands in future optical networks, it is crucial to develop sustainable solutions to reduce energy consumption. Moreover, ensuring network resiliency during failures, such as those caused by natural disasters, is vital for maintaining uninterrupted services. At the same time, quantum computers present significant security challenges to current cryptosystems, making it essential to develop effective countermeasures against potential quantum attacks. This thesis aims to design resource allocation algorithms to improve the energy efficiency, resiliency, and security of future optical networks, and the main contributions can be summarized into three parts as follows.

1. Proposal to Improve Energy Efficiency based on Novel Optical-Transmission Technologies: Recent advances in optical-transmission technologies bring new possibilities for enhancing energy efficiency for future optical networks. I have designed resource allocation algorithms for two emerging optical-transmission approaches, namely pluggable optics and power profile monitoring (PPM). Regarding pluggable optics, I have comprehensively investigated the energy efficiency of different "IP over Wavelength Division Multiplexing" (IPoWDM) network architectures with ZR/ZR+, and evaluated the advantages of ZR/ZR+ compared to long-haul muxponders. Specifically, I have proposed a pragmatic power consumption model for optical nodes and designed an auxiliary-graph-based approach to quantify the power consumption of different IPoWDM network architectures. Regarding PPM, I have proposed and investigated the novel optimized monitoring placement (OMP) problem for PPM to minimize the monitoring power consumption. Then, I have quantitatively compared the power consumption of PPM vs. the traditional monitoring technique, Optical Time-Domain Reflectometer (OTDR), with both an Integer Linear Programming (ILP) model for small-scale scenarios and an efficient heuristic algorithm for large-scale scenarios. These works provide guidelines for the deployment of ZR/ZR+ and PPM to improve the energy efficiency of future optical networks.
2. Proposal to Enhance Network Resiliency with Proactive and Reactive Approaches: A key challenge in network virtualization for future optical networks is to improve network resiliency against possible failures and enable swift recovery after failures. I have investigated proactive solutions to improve network resiliency against double-link failures and reactive solutions to provide swift network recovery under massive failures. Specifically, to provide network resiliency against any possible double-link failures in the network, I have investigated capacity sharing among existing virtual networks to reduce the additional resources to enhance network resiliency. More specifically, I have explored two strategies to enhance resiliency through capacity sharing, one employing a dedicated slice to improve resiliency and the other without such a slice. To solve this problem, I have formulated an ILP model and designed a scalable local-search-based heuristic algorithm. Regarding swift recovery, I have introduced a novel problem, called progressive slice recovery (PSR) with guaranteed reachability among network nodes (network connectivity) and reachability of data centers (content connectivity). Specifically, I have investigated to determine a recovery sequence for

slices to swiftly recover the slices and, in the meantime, consider accelerating the provision of network services with connectivity constraints. To solve the PSR problem, I have developed an ILP model and a two-phase progressive slice recovery algorithm based on column generation and deterministic rounding. The numerical results confirmed that network resiliency can be guaranteed with less additional resources through proactive approaches and shorter recovery time using reactive approaches.

3. Proposal to Improve Network Security with Quantum Technologies: Quantum computers pose emerging security threats to future optical networks, making it critical to explore novel algorithms and architectures with quantum technologies. I have investigated novel resource allocation algorithms and architectures for QKD and quantum networks, enhancing their ability to securely distribute keys among nodes with guaranteed information-theoretic security. First, I have introduced a novel problem of Routing, Channel, Key-rate and Time-slot Assignment (RCKTA) to allocate quantum resources for QKD networks. This problem considers incorporating the advantages of trusted relays to increase the key rate and optical bypassing to reduce the unnecessary QKD modules (i.e., transceivers for QKD networks). To solve the RCKTA problem, I have devised a Mixed Integer Linear Programming (MILP) model and a near-optimal heuristic algorithm, achieving a substantial increase in acceptance ratio. Next, I have designed a novel QKD net-work architecture integrating Information-Centric Networking (ICN) to address the challenges of limited secret key rates and diverse service requirements. This architecture can accelerate key distribution with in-network key caching of ICN and meet diverse service requirements by analyzing the requirements based on semantic information. Lastly, I have investigated link configuration for fidelity-constrained routing and purification (LC-FCRP) in quantum networks, developing a MILP model and a shortest-path-based fidelity determination algorithm to optimize the fidelity of distributed entanglements. These approaches significantly increase the performance of QKD networks and quantum networks to guarantee the security of future optical networks.

PhD Committee

Prof. Francesco Musumeci, **Politecnico di Milano**

Prof. Alessio Giorgetti, **Università di Pisa**

Prof. Fen Zhou, **University of Avignon**