ABSTRACT

This thesis addresses two challenging problems. The first part is focused on developing algorithms for solving multi-time scale optimization models and their application to the electrification of chemical process systems. The second part focuses on constrained Markov decision processes (CMDPs), with an emphasis on the development of reinforcement learning environments for operations research problems and the evaluation of safe reinforcement learning algorithms on these environments.

Multi-time scale optimization provides a mathematical framework for coordinating decision-making across multiple temporal layers, ranging from long-term planning to short-term operational control. Such models are particularly relevant in process systems and energy applications, where strategic and operational decisions are strongly coupled. However, solving these models remains computationally challenging due to their large scale and complex structure. In Part I of this thesis, we propose a multi-time scale mixed-integer linear programming (MILP) model for the integrated design and operation of electrified chemical process systems, capturing decisions related to plant location, renewable energy integration, transmission infrastructure, monthly transportation, and hourly scheduling. To address the computational challenges of this formulation, we develop a K-means clustering-based aggregation-disaggregation matheuristic tailored for solving the problem. In addition, we propose the Parametric Autotuning Multi-time Scale Optimization (PAMSO) algorithm, a more general and scalable approach for generating high-quality feasible solutions for multi-time scale optimization problems.

Part II of this thesis addresses constrained Markov decision processes and safe reinforcement learning for operations research applications. Reinforcement learning has shown significant promise in solving sequential decision-making

problems; however, most existing approaches do not explicitly account for constraints, limiting their applicability in real-world systems. CMDPs provide a principled framework for incorporating such constraints, but remain challenging to solve in complex, structured environments. In this part, we develop SafeOR-Gym, a benchmark suite consisting of nine operations research environments that capture the combinatorial structure, long planning horizons, and constraint-rich nature of practical decision-making problems. These environments are integrated within a CMDP framework and implemented using the OmniSafe library. We evaluate several state-of-the-art safe reinforcement learning algorithms across these environments and demonstrate that, while existing methods perform adequately on simpler tasks, they face significant challenges in handling highly constrained and structured problems. These findings highlight key limitations of current approaches and motivate the development of new methodologies at the intersection of reinforcement learning and operations research.