Abstract

The widespread adoption of advanced metering infrastructures based on Internet of Things (IoT) could enable the development of Energy Management and Information Systems (EMIS) capable to leverage useful knowledge extracted from building related data. This dissertation focuses on a specific category of EMIS technologies called Automatic System Optimization (ASO). The purpose of ASO tools is to actively manage the control strategies responsible for the operations of building energy systems with the aim of enhancing energy usage. Among building subsystems, Heating Ventilation and Air Conditioning (HVAC) systems are rated among the most energy-intensive end-uses. The non-linear and stochastic nature of these systems makes the definition of robust and effective control strategies particularly challenging. In the current paradigm of smart buildings, building managers and owners can leverage ASO tools to automatically optimize the performance of their systems. However, the management of HVAC systems is mainly based on classical approaches characterized by different drawbacks including a reactive approach, lack of an optimization process and impossibility to handle multiple objectives at the same time. To overcome these limitations the application of advanced control strategies based on predictive and adaptive approaches represents a promising direction. In this dissertation four different applications of deep reinforcement learning based control strategies were conceived and tested. Deep Reinforcement Learning (DRL) is a model-free approach in which a control agent leveraging deep neural networks directly learns an optimal policy from interacting with the controlled environment. The developed applications were carried out in a co-simulation environment combining Python and EnergyPlus specifically developed in the context of this dissertation. Each application was designed to address different challenges and questions related to the application of DRL controllers to HVAC systems. In the first application, DRL is implemented to control the supply water temperature setpoint to terminal units of a heating system. The performance of the agent is evaluated against a reference controller that implements a combination of rule-based and climatic-based logics. As a result, when the set of variables are adequately selected a heating energy saving ranging between 5% and 12% is obtained with an enhanced indoor temperature control with both static and dynamic deployment. In the second application a DRL agent was trained employing a data-driven model of the building dynamics. The trained agent was statically deployed on a calibrated Eplus model of the building to evaluate its performance. The agent was conceived to control the supply water temperature setpoint of the heating system of an office building achieving a reduction in the energy consumption of 18% while improving indoor temperature control of 5% with respect to a baseline rule-based controller. In the third application, was investigated the potentialities of DRL strategies for the management of integrated energy systems in buildings with on-site electricity generation and storage technologies. The controller is tested considering various configurations of battery energy storage system capacities, and thermal energy storage sizes. Results show that the proposed control strategy leads to a reduction of operational energy costs respect to a rule-based controller ranging from 39.5% and 84.3% among different configurations. The last application introduces a comparison between an online and offline DRL with a Model Predictive Control (MPC) architecture for energy management of a cold-water buffer tank linking an office building and a chiller subject to time-varying energy prices, with the objective of minimizing operating costs. Simulation results showed that the online-trained DRL agent, while requiring an initial 4 weeks adjustment period achieving a relatively poor performance (160% higher cost), it converged to a control policy almost as effective as the model-based strategies (3.6% higher cost in the last month). Findings and outcomes of the present research study are discussed providing a robust reasoning about the application of DRL control strategies to HVAC systems. Eventually, a wide overview on the lessons learned throughout this research study is proposed to outline the future opportunities and barriers to the adoption of advanced control strategies in the energy and building sector.