Data Predictive Control - Bridging Machine Learning with Controls for Volatile Energy Markets

Rahul Mangharam, University of Pennsylvania

Machine Learning and Control Systems are two foundational but disjoint communities. In this talk I will explore how we can bridge the worlds using data-driven approaches to synthesize control-oriented models to solve challenging problems in Energy and Industrial processes. Decisions on how to best optimize energy systems operations are becoming ever so complex and conflicting, that model-based predictive control (MPC) algorithms must play an important role. However, a key factor prohibiting the widespread adoption of MPC in buildings, is the cost, time, and effort associated with learning first-principles based dynamical models of the underlying physical system. This talk introduces new approaches for implementing finite-time receding horizon control using control-oriented data-driven models. We call this approach Data Predictive Control (DPC) and apply it to price volatility problems in Energy Markets and Industrial Process Control plants.

In December 2014, the average price of wholesale electricity in the PJM market surged from $25/MWh to $2680/MWh - an 83x increase in 5-15mins. Demand response (DR) is becoming increasingly important as the volatility on the grid continues to increase. Current DR approaches are predominantly manual and rule-based or involve deriving first principles based models which are extremely cost and time prohibitive to build. We consider the problem of data-driven end-user DR for large buildings which involves predicting the demand response baseline, evaluating fixed rule based DR strategies and synthesizing DR control actions. We provide a model-based control with regression trees algorithm, which allows us to perform closed-loop control for DR strategy synthesis for commercial buildings. Our data-driven control synthesis algorithm outperforms rule-based DR by 17% for a large DoE commercial reference building and leads to a curtailment of up to 380 kW and over $45,000 in savings for just one season.

Our methods have been integrated into a tool called DR-Advisor, which acts as a recommender system for the building's facilities manager and provides interpretable control actions to meet the desired load curtailment while maintaining operations and maximizing the economic reward. Built upon DR-Advisor is IAX, an Interactive Energy Analytics engine - think of it as a Siri for querying buildings' energy use. We are developing IAX to procedurally generate energy dashboards for open-ended questions. This contributes to a second fundamental goal of making Machine Learning models interpretable so solutions have traceability and can be verified to be stable and safe.

More info at: http://mlab.seas.upenn.edu/projectsites/dr-advisor/