Title : Enhanced Delta-tolling: Traffic Optimization via Policy Gradient Reinforcement Learning

Authors : Hamid Mirzaei Buini, Guni Sharon, Stephen Boyles, Tony Givargis and Peter Stone
Presenter: Guni Sharon

Abstract : The prospect of widespread deployment of autonomous vehicles invites the reimagining of the multiagent systems protocols that govern traffic flow in our cities. One such possibility is the introduction of micro-tolling for fine-grained traffic flow optimization. In the micro-tolling paradigm, different toll values are assigned to different links within a congestable traffic network. Self-interested agents then select minimal cost routes, where cost is a function of the travel time and tolls paid. A centralized system manager sets toll values with the objective of inducing a user equilibrium that maximizes the total utility over all agents. A recently proposed algorithm for computing such tolls, denoted Delta-tolling, was shown to yield 32% reduction in total travel time in a simulated traffic scenario compared to when there are no tolls. Delta-tolling includes two global parameters: beta which is a proportionality parameter, and R which influences the rate of change of toll values across all links. This talk will introduce a generalization of Delta-tolling which accounts for different beta and R values on each link in the network. While this enhanced Delta-tolling algorithm requires setting significantly more parameters, we show that they can be tuned effectively via policy gradient reinforcement learning. Experimental results from several traffic scenarios indicate that Enhanced Delta-tolling reduces total travel time by up to 33% compared to the original Delta-tolling algorithm, and by up to 52% compared to not tolling.