FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Welcome to the FORCES newsletter

Summer 2016

It's summer, but outstanding research continues among my colleagues involved in the FORCES project. This issue of the FORCES newsletter includes articles about transportation and routing games, load allocation in energy systems, and the challenge of using probability with large sets of data. These all touch on the common theme of how different cyber-physical systems interact with economics. I'm also very pleased to include a commentary from FORCES industrial advisory board member Dr. Shaunak Bopardikar at UTRC, who offers an observation regarding the relevance of the work of FORCES researchers to his own company's priorities.

Finally, I want to particularly draw your attention to the FORCES in the News section, which provides a link to a recent "Dear Colleague" letter from the National Science Foundation. Although this letter was published last year, it indicates current NSF funding interests and I'd like to include it here for your review as it has great resonance among the FORCES group.

Thanks very much for taking time to read the FORCES summer 2016 newsletter. As always, if you have feedback, comments, and suggestions I'd appreciate hearing from you.

Sincerely,
S. Shankar Sastry
Professor and Dean of Engineering
University of California, Berkeley

RESEARCH SPOTLIGHT

How Players Learn in Sequential Decision Problems

by Walid Krichene (University of California, Berkeley), Chedly Bourguiba (Ecole Polytechnique), and Alexandre Bayen (University of California, Berkeley)

The last decade has witnessed an unprecedented transformation of mobility, in which motorists have started, on a massive scale, to follow recommendations of apps on their smartphones to obtain the best routes to their destinations. This paradigm, increasingly adopted in the US, has over time altered the mobility patterns in large urban regions of the US. Public agencies are hearing concerns of residents in suburban parts of cities who are worried about increased traffic in their neighborhoods. The stability of these patterns has become an issue as well (i.e., its unpredictability and the inability of public agencies to control it).

As part of the FORCES agenda, our team has undertaken the study of routing games to model the congestion phenomena and dynamics that can be observed in transportation and communication networks. In particular our framework describes situations in which a decision maker routes traffic over the network. A decision maker can be a single driver, or a controller that is responsible for routing a large amount of traffic--for example a company running a routing app that will inform a certain percentage of the drivers that will follow its suggestions. In this framework, it is assumed that each decision maker faces a sequential decision problem--adjusting decisions based on past observations. Each person is assumed to act selfishly (i.e., companies try to provide motorists the best route for them, which is not necessarily optimal from a system-wide or societal perspective) and non-cooperatively (i.e., the routers do not share information and have no incentives to collaborate). We are developing models of online distributed learning, under which such mechanisms are guaranteed to converge to Nash equilibria. In other words, if each decision maker follows an algorithm in a given class that we characterize, the coupled system is well-behaved, in that it is guaranteed to converge to the equilibrium set.

Additionally, we study whether these online learning models can describe the behavior of human decision makers. More precisely, assuming that we have a parameterized family of learning algorithms, and given that we observe a sequence of decisions (e.g., route choices) of a human agent participating in the routing game, can we estimate the model parameters to best fit these observations? We developed a web application that implements the routing game, and interfaced it with Amazon's Mechanical Turk to run the game with a large number of human participants who play remotely and non-cooperatively.

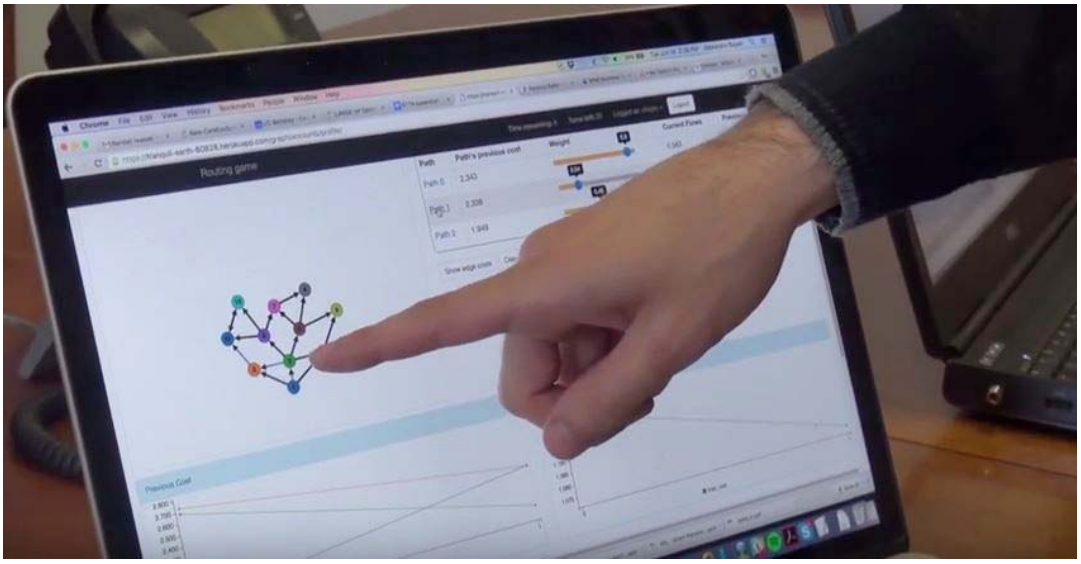

Demonstrating modeling software

We used the data collected from the experiments to give qualitative and quantitative insights on the decision dynamics of human players. In particular, we observed that after an initial exploration phase, the distributed learning tends to converge close to equilibrium, but that occasionally, irrational behavior of players can push the system away from equilibrium for a few iterations. The experiment also shows the expressive power of the online learning model, and that it can be used to accurately predict the evolution of the system over short horizons.

Studying Electrical Load Aggregation

by Ian Hiskens, University of Michigan

Electrical loads, when considered in aggregation, have the potential to provide appreciable demand-side control. With appropriate control strategies, loads can participate in a variety of power system operational services, for example assisting in maintaining supply-demand balance or levelling the output of renewable generation. To do so, however, requires controls that coordinate the behavior of large numbers of small, diverse devices in a manner that causes negligible disruption to consumers.

Loads that operate according to a hysteresis band, such as thermostatic loads, are quite amenable to control strategies that are non-disruptive in the sense that they have no impact on the quality of service to consumers. Load control can be achieved, for example, by adjusting the hysteresis band of each participating load to effect a change in aggregate power demand. Such a control mechanism was originally proposed for thermostatically controlled loads, for example air conditioners, refrigerators and water heaters. Similar ideas have subsequently been extended to other types of loads, including pool pumps and plug-in electric vehicle (PEV) chargers.

Predicting the input-output behavior of a load aggregation is vitally important for designing controls and analyzing system response. Such systems are unusual though, as they consist of many hundreds (possibly thousands) of individual devices, each obeying its local control law, but which in turn is influenced by the aggregating controller. Initially Fokker-Planck models were derived to capture the aggregate dynamics of the population of loads. More recently, Markov decision process (MDP) models have been developed. These models are based on partitioning the hysteresis band into several equal-sized “bins” and keeping track of the propagation of devices from one bin to another. This general idea is illustrated in Figure 1.

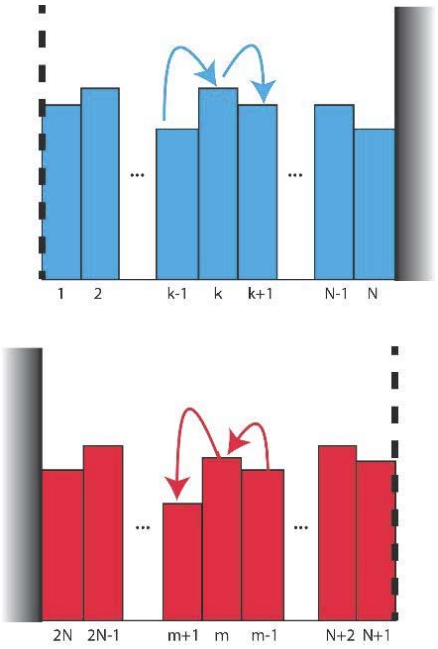

Figure 1: Markov decision process model

Standard MDP models are applicable when the control signal changes slowly relative to the time constants of the individual loads. If that is not the case, the required model adaptation results in a bin model that has a hybrid (switched) dynamical system structure. Consequently, such load aggregations may exhibit rich forms of behavior. As an example, a load aggregation was subjected to a triangular-wave input and the period of that input was varied. Figures 2-4 show the total power of the aggregate load for input periods of 30.8, 24.4 and 15.6 minutes, respectively. It can be seen that as the period reduces the load control system undergoes period-adding bifurcations, i.e. the output power shown in Figure 2 has the same period as the input whereas the output period in Figure 3 is double the input and triple in Figure 4. These bifurcations appear to be a consequence of synchronizing phenomena, though ongoing work is seeking to develop a more complete understanding of the relevant nonlinearities. Furthermore, for other values of the input period, the output power exhibits non-periodic, apparently chaotic, behavior. This certainly complicates control of the load aggregation.

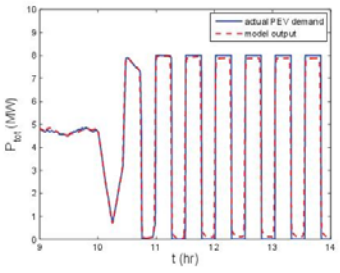

Figure 2: Period = 30.8 min.

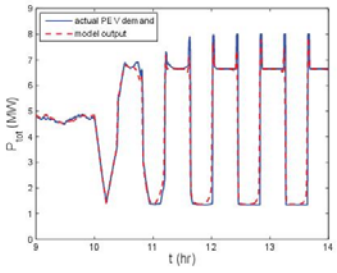

Figure 3: Period = 24.4 min.

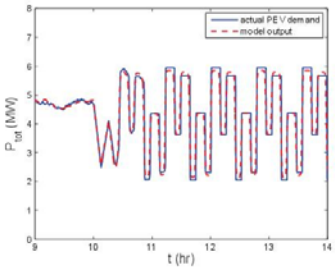

Figure 4: Period = 15.6 min.

Work on the modelling and control of load aggregations is ongoing. One aspect of this work is exploring so-called “transactive energy control” where controls are enacted through price signals. Initial investigations indicate that transactive control could easily result in synchronization of loads, leading to highly oscillatory patterns of electrical demand. Such an outcome is undesirable, as it could potentially destabilize the electricity supply system. Further work is required to develop control mechanisms that mitigate this synchronization behavior.

Load aggregation offers enormous benefits by facilitating the involvement of many small loads in the operation of large-scale power systems. Much work remains in the development of analytical models and the use of those models for the design of robust control strategies.

A Scalable Approach to Dynamic Cyber-security in the Presence of Non-probabilistic Uncertainty

by Erik Miehling, Demos Teneketzis, and Mohammad Rasouli, University of Michigan

Central to all cyber-physical systems and smart cities is a cyber component, making cyber-security an important part of ensuring our systems operate as intended. When modeling such security problems, one needs to have knowledge of the attacker's strategy and the ability to assign appropriate values to parameters in the problem. Specifically, when using probabilities to model uncertainties in the problem, one needs to either ensure that the defined probabilities are accurate or that there is enough data in order to learn these probabilities over time. Unfortunately, this data is often unavailable, resulting in an inaccurate model. Furthermore, the attacker's strategy is private information.

As a result, we propose a minmax approach to defending the system — that is, taking actions that minimize the worst-case damage possible. This modeling choice circumvents the need to define probabilities. Instead it considers all possible events, no matter how unlikely. Decision problems in this domain require one to maintain a record of all possible system trajectories that are consistent with the available information. As one can imagine, these decision problems are computationally difficult, often exploding in complexity as the system grows in size.

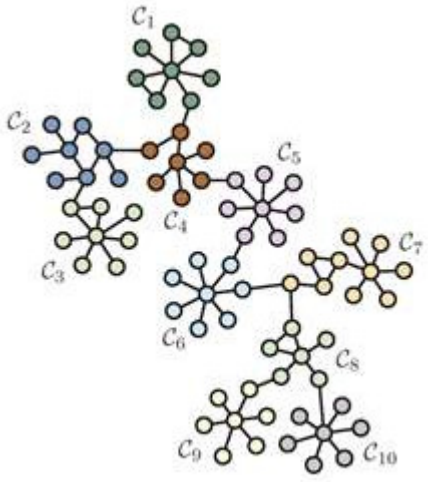

Influence graph

Since maintaining a record of all possible system trajectories is prohibitively complex, we seek methods that are more scalable. We define an influence graph which serves to describe the functional dependencies between variables and allows us to decompose the large system into many, smaller sub-systems. Each sub-system is operated by a local controller which communicates necessary variables with its neighboring sub-systems. The resulting approach allows local controllers to compute local (defense) policies, allowing us to tackle realistically-sized problems in the cyber-security domain.

On Privacy in Cyber-Physical Systems

by Shaunak D. Bopardikar, Staff Research Scientist, Systems Department, United Technologies Research Center, Inc.

A prototypical privacy problem arises when 1) one or many owners of datasets would like to pool their data together in order to compute a certain function on the aggregate and that, 2) the revelation of certain private parameters in their data will lead to an improvement in the outcome to the owners (utility). A distinguishing feature of this problem from more traditional approaches in the database security community is the connection of Cyber-physical Systems (CPS), which implies that the datasets are not static. Additional features such as the dataset being directly related to an end-product compound the problem further in the sense that corporations now need to be concerned about how to protect the functioning of their products against resource-rich, semi-honest adversaries trying to extract proprietary information.

In this context, research of some of the NSF FORCES project members on topics such as energy disaggregation suggests that private parameters can be learned with high confidence from data arising out of typical CPS. Therefore, practical tools which will allow the characterization or visualization of complexity-utility-privacy tradeoffs may be valuable from an end-user as well as from a product design point of view.

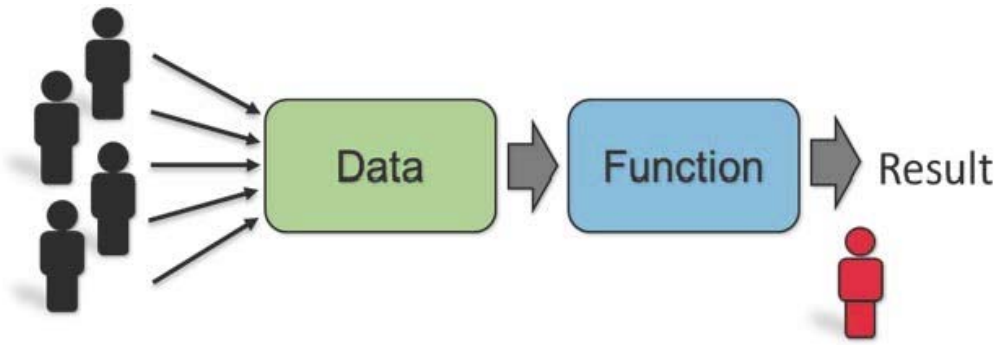

Data from multiple users can endanger privacy

Upcoming Events

CPS PI Meeting
Fall 2016
October 31 - November 1
Location TBD

FORCES In the News

NEW CPS TESTBED AT VANDERBILT UNIVERSITY

Vanderbilt University has developed the Resilient Cyber-Physical Systems (RCPS) testbed, a new research tool to evaluate the resilience of CPS algorithms and software in a realistic environment. It consists of a cluster of 32 embedded computing boards that run a real-time operating system and the CPS software, a programmable network switch for emulating a network among the cluster's nodes, and a high-performance, real-time physics simulator that is connected to the embedded boards via a second network switch. An additional server node is used as the development machine for the cluster. Arbitrary

scenarios involving faults and cyber effects can be emulated on the platform and the performance of realistic CPS software in stressful situations can be evaluated.

"DEAR COLLEAGUE" LETTER FROM NSF

A letter from the National Science Foundation included information regarding current funding prioririte that have great relevance to the FORCES community.
