Information Leakage in Energy Cyber Physical Systems: A Stochastic Control Perspective

Parv Venkitasubramaniam, *Lehigh University*

Cyber-physical systems are envisioned to transform the way energy systems function, far exceeding the capabilities of systems today. These are typically controlled stochastic systems which rely on joint functioning of information systems and physical components, and have already begun to replace our existing infrastructures for energy generation, distribution, management and control. While the success of cyber physical systems relies on the *power of information exchange* to improve system functioning, it's fallibility lies in the *power of information leakage* to compromise system operation. Despite tremendous advances in cryptography, cyber communication is far from being truly secure, which is a critical impediment to the success of cyber physical systems. Leaked system information can deprive energy systems of desired consumer privacy and can leave them vulnerable to adversarial attacks. For instance, the response to pricing signals in a smart metering system can reveal daily activities in a household. Indeed, the privacy concerns of smart meters has resulted in many successful lawsuits against their deployment, and in many states led to opt-out schemes for consumers, thus undercutting its real benefits. In addition to violation of individual user privacy, leakage of sensitive information can also empower adversaries to cause alarming infrastructural disturbances such as power blackouts to malfunctioning nuclear reactors. If we are to realize the full potential of cyber physical systems, it is critical that we study the design of control policies for general stochastic systems from the perspective of minimizing information leakage.

There are some key challenges that need to be addressed towards realizing the objectives of privacy preserving energy cyber physical systems design. First, privacy vulnerabilities could arise from *implicit information* inferred and not merely explicit data communicated which could be protected using cryptographic mechanisms. Implicit information on a cyber link includes timing of data transmitted, size of packets and paths of data flow; in other words information readily available to eavesdroppers with fairly unsophisticated equipment. Second, stochastic control systems present a unique scenario where actions not only leak current information (through observed outputs) but also generate future information (through state transitions). Consequently, when studying the privacy of stochastic control systems, it is important to formulate a quantitative metric that measures the information leaked from the perspective of an observer who has access to the complete horizon of information. This raises an important caveat; in order to determine the state of the system at a given time, an eavesdropper can use all information inferrable from the **past, present and future** observations. In other words, the value of information is neither causal nor does it diminish over time. To better understand this caveat, consider the example of an electricity distribution system which uses smart meters to gather electric usage data every second as a means to provide the user the benefit of the time varying electricity prices. If such a user were to install a mechanism to exploit price fluctuation by deferring loads in time, then the observed consumption at a particular time contains information about loads generated in the past as well. As a result, the knowledge about demands generated at a particular time will be updated by an adversary upon receiving future observations. Any measure of privacy must necessarily capture this phenomenon and will therefore be *non-causal*.

A rigorous analytical framework is required that expands the science of controlled dynamical systems to include privacy or information leakage as a parameter/performance metric. The expanded framework should yield the twin perspectives of any *secure* system: the policy design perspective of the system operator and the attack design perspective of the adversary. Modifications to controller policy design is rarely free, and it is essential to understand the cost incurred to provide privacy as measured in the primary utility reduction of the controlled system. Some key questions that arise in the study of such a tradeoff are of importance: How does achievable privacy scale with the key parameters of the system? Are there conventional utility metrics that align with privacy with regard to policy design? For an adversary with practical limitations, how best should he invest his available resources, and how does this choice vary with system parameters? How do classical control properties such as observability, reachability and controllability impact the achievable privacy and the design of privacy preserving control policies? If an adversary is capable of active modification of external measurements with the objective of compromising internal privacy, can statistical information about legitimate measurements be used to detect such intrusion?

The grand vision is to build scientifically sound privacy preserving adaptive dynamical energy systems for our future; systems that are decentralized, self configuring and can scale and adapt to unforeseen events and uncertainties across multiple dimensions, including types of energy sources, climactic conditions, network topologies, mobility patterns, and heterogeneity of devices and technologies.