Title: Probabilistic programming as a framework for cyber security monitoring
Author: David Robinson, drobin@sandia.gov
Organization: Scalable Analytics and Visualization, Sandia National Laboratories
Focus Area: Extreme Scale Data, Knowledge, and Analytics

Traditional machine learning methods for detecting abnormal patterns or information anomalies rely on first identifying 'normal' and then detecting deviations from this baseline. Development of security monitoring systems, including those based on popular machine learning methods, can be a time consuming, system dependent, and unique to a particular signature. The result is that once a machine learning model is instantiated as a security monitor, there is substantial resistance to update the model while the 'normal' operational environment continues to evolve with new users, new hardware, new code. New methods for monitoring must be flexible either through a natural evolutionary (learning) capability or permit fast re-configuration and re-deployment.

The extreme number and variety of system parameters that need to be monitored in a streaming environment quickly overwhelm traditional statistical and machine learning security monitoring methods. Computational requirements of these algorithms necessitate approximations such as dimension reduction or shortened observation windows. The result is a reduced sensitivity to potential security events. New methods must formally quantify the impact of these approximations on the generation of false-positive and false-negative security event alarms.

Development of metrics and/or features for monitoring is critical. Equally imperative is the ability to identify *which* metrics are the most valuable to providing insight into the state or performance of the system? Probabilistic programming methods have the potential to provide a solid mathematical foundation for monitoring the sensitivity of critical metrics, well beyond machine learning algorithms. Finally, there is the situation of incomplete information. Not all HPC components are monitored or instrumented. Anomalies or intrusions within those areas would be undetectable. Some consideration must be made when enhancing trustworthiness or integrity, to identifying anomalous behavior when true anomaly is either masked or not directly observable.

DARPA has recognized many of the above limitations to processing information in environments such those experienced in an HPC cyber security monitoring. The result was to establish the Probabilistic Programming for Advanced Machine Learning research program. However, research efforts have been underway in this area long before DARPA expressed concern.

In particular, there has been a number of these research efforts related to HPC monitoring and detection of internal/external tampering. Highlights of these preliminary forays include:

- Developed and deployed, using probabilistic programming methods, various Bayesian statistical monitoring techniques that are flexible to reconfiguration, while robust to extreme data and signature dimension. These methods are also applicable when data is incomplete or anomalies are masked.
- Risk-based identification of HPC anomalies involving a probabilistic statement regarding the existence of a fault, an intrusion or tampering. This permits, for each event, the quantification of the likelihood that it is a true concern or a localized excursion from normal operation; providing control over false-positive and false-negative alarms.
- Establishing causal relationships between internal and external events, state of the HPC system, and negative cyber security signatures. Information theoretic and statistical methods have been explored.
- Formal statistical methods for objective integration of information from HPC monitoring with information from external sources; job logs, personal sysadmin observations, and other external humint.
- Optimal rescheduling of jobs to avoid HPC subsystems, either because the job was suspect or the subsystem was corrupted. Job scheduling was also proposed to specifically probe hardware/software elements, to reduce risk of further tampering, and increase confidence in event identification.

The above represent preliminary investigations in HPC system security monitoring, but in some cases have led to the deployment of HPC monitoring algorithms for the intelligence community to improve resiliency to events and increase overall throughput. While the research efforts were successful, they were basic investigations focused on proof of concept and to suggest areas for further, more substantial efforts.

Current research efforts are much broader in scope with the possibility of extreme dimensions for signature investigation. These efforts involve the use of Bayesian statistical methods to aid identification of specific HPC performance signatures and the association of these signatures with system integrity. The goal is to assist in making trade-offs between computational integrity and performance, and provide information for optimal scheduling of jobs.