Real-time Causal Internet Log Analytics by HW/SW/Projection Co-design

Bita Darvish Rouhani, Mohammad Ghasemzadeh, Farinaz Koushanfar
University of California San Diego
{bita, mghasemzadeh, farinaz}@ucsd.edu

1 Abstract

In this demo, we will present the first system that is able to perform real-time multi-dimensional analysis of time-series Internet log data. The causal and autoregressive nature of complex Internet logs challenges building the models in an online manner and limits the automated real-time reasoning from the vast amount of stored logs. Our system is built upon streaming-based Probability Density Function (PDF) approximation in the context of causal Bayesian graphical models using FPGA. The latent variables of the corresponding probabilistic model in our system are iteratively updated in line with the data arrival to account for the data dynamics. We devise an accompanying API that can be leveraged for automated end-to-end prototyping of real-time augmented graph analysis of Internet traffic data used to detect anomaly (e.g., DoS attack) in the pertinent network.

2 Global Flow

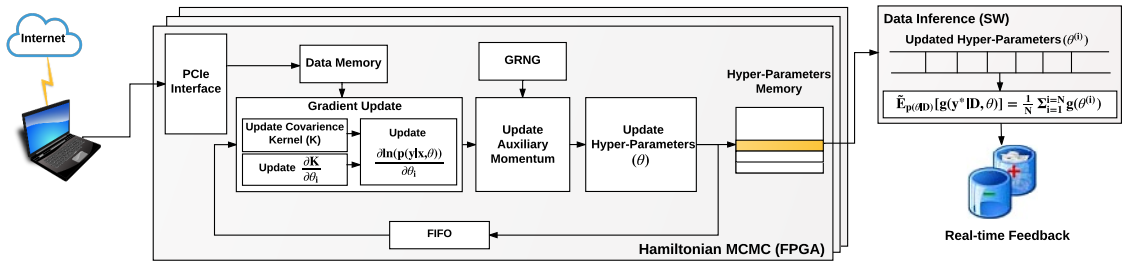

Figure 1: Global flow of proposed demo.

Figure 1 demonstrates the global flow of the proposed framework. Our framework is devised based on a HW/SW/projection co-design approach. It takes the stream of Internet log data as its input and adaptively updates the corresponding random variables and their associated PDFs to comply with the newly arrived data samples. The updated hyper-parameters are used to perform a particular user-defined data inference task (e.g., real-time probabilistic classification). We provide the first implementation of Hamiltonian Markov Chain Mont Carlo (H_MCMC) on FPGA that can efficiently sample from the steady state probability distribution at scales while considering the correlation between the observed data. Computing the gradient-based update per H_MCMC iteration involves a variety of operations with a complex data flow. We modify the conventional Hamiltonian MCMC routine to adapt the underlying data flow for FPGA acceleration.

Our framework, for the first time, addresses the problem of *real-time streaming-based* PDF approximation in the context of *causal* Bayesian models. Our approach is both *dynamic* and *robust*. The dynamism is to adapt to the instantaneous changes in the distribution of the data, while the robustness is to not skew the distribution to the outliers. Note that the existing FPGA realizations of Bayesian networks have been mainly developed based on the assumption that data samples are independently and identically drawn from a certain distribution. As such, unlike our framework, they cannot effectively capture dynamic data correlation in streaming applications. Another unique aspect of our framework is that it operates on streaming data and builds the model in real-time as data evolves over time.

HW settings: In our prototype, we use Xilinx Virtex UltraScale FPGA VCU108 as the primary hardware accelerator. An Intel core-i5 laptop with $8\ GB$ memory running on the Windows OS at $2.40\ GHz$ is used as the general purpose processor hosting the FPGA. We leverage PCIe port to interconnect the host and FPGA platform. All computations are performed using single-precision floating-point operations. Vivado HLS 2016.4 is used to synthesize and simulate our MCMC units. The FPGA platform is programmed with a speed grade of -2 and works at $100\ MHz$ clock frequency.