

**I** ILLINOIS

CSL | Coordinated  
Science Lab

COLLEGE OF ENGINEERING

Subho S. Banerjee, Mohamed el-Hadedy, Ching Y. Tan,  
Steve Lumetta, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer

# **Symphony**: Probabilistic Graphical Models for Scheduling Heterogeneous Processors

[csl.illinois.edu](http://csl.illinois.edu)



# Overview

Medical Knowledge



Continuous Monitoring



\*omics Data



Medical Devices



Patient Records



- Timely Diagnosis
- Personalized Drugs
- Model Drug Response
- New Biological Insight

# Overview

Medical  
Knowledge

Continuous Monitoring

## Summary of Results

Top-down approach to building the CompGen machine

- Static analysis of genomic analyses algorithms

- Hardware Acceleration

- Scheduling strategies to deal with heterogeneous hardware

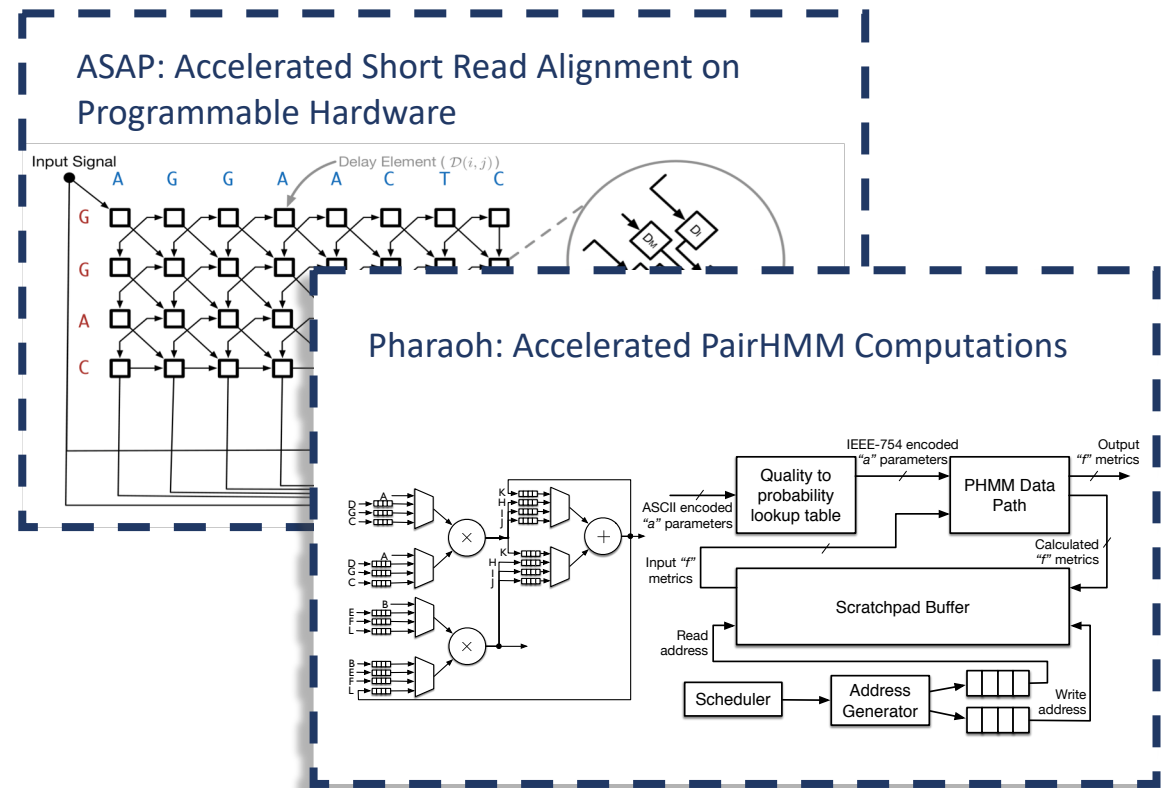
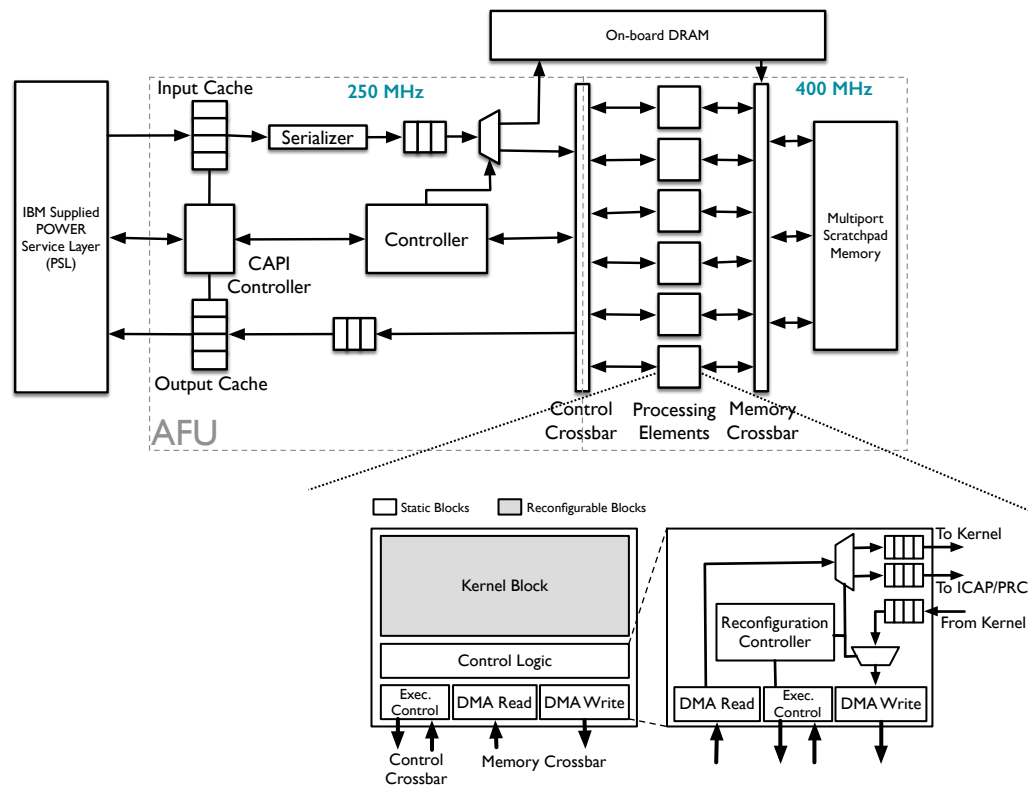
“Variant Calling and Genotyping” Workflow as the driver

	Baseline Runtime	CompGen Accelerated Runtime **	Speedup
Blue Waters – Single Node (CPU)	59 hr	<b>28 hr</b>	2.1x
IBM Power 8 – Single Node (CPU + GPU + FPGA)	36 hr	<b>42 min</b>	84x, 51x
Blue Waters – 10 Nodes (CPU)	-	<b>2 hr</b>	29.5x

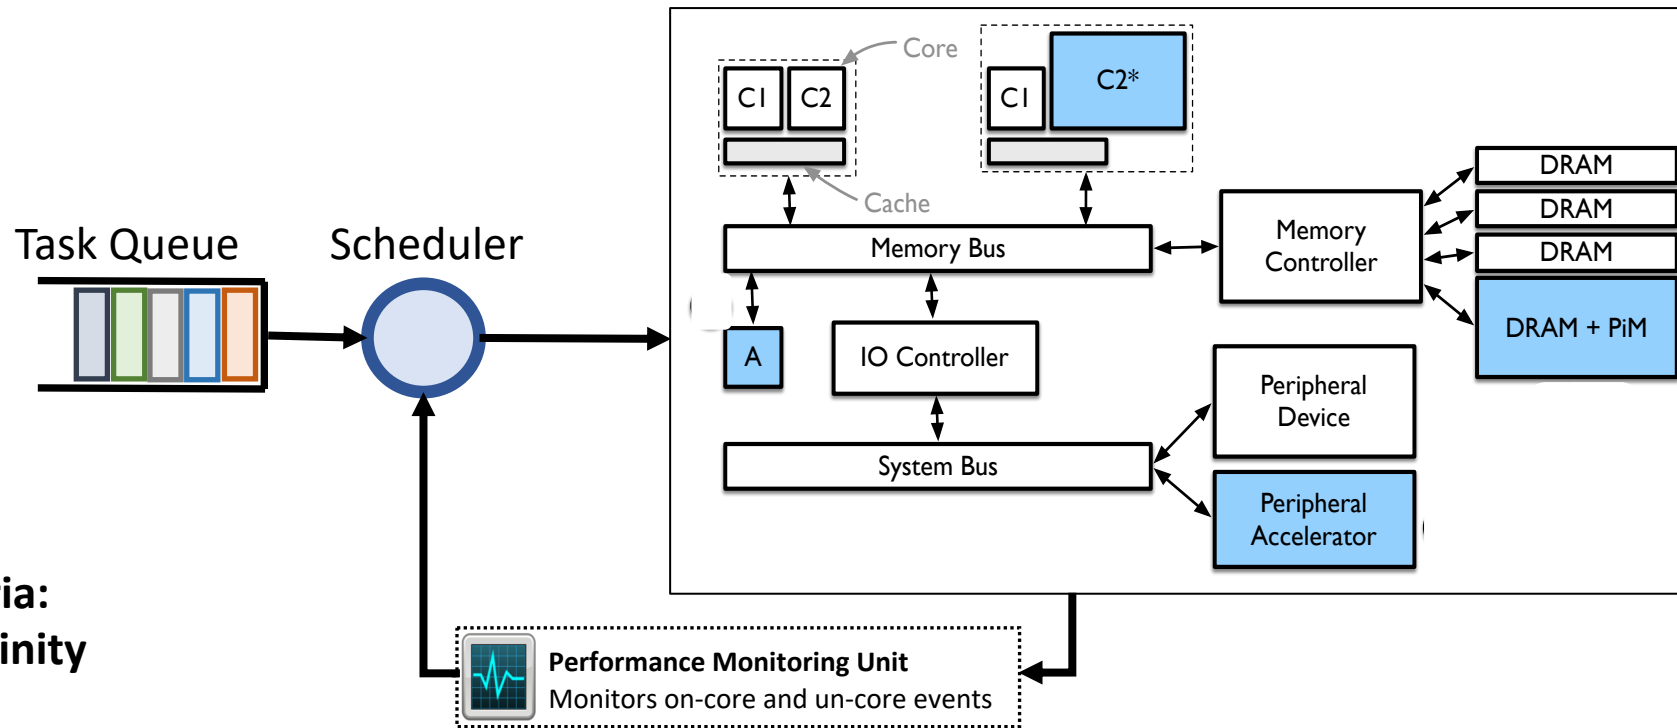
# TCGA: The Computational Genomics Accelerator

A reconfigurable *many-kernel* coprocessor for computational genomics applications

- Common computational kernels across applications
- Memory centric computation: IBM CAPI, Hazard free scheduling, runtime reconfiguration
- Specialized Compute: Algorithmic Approximations, Delay based computation



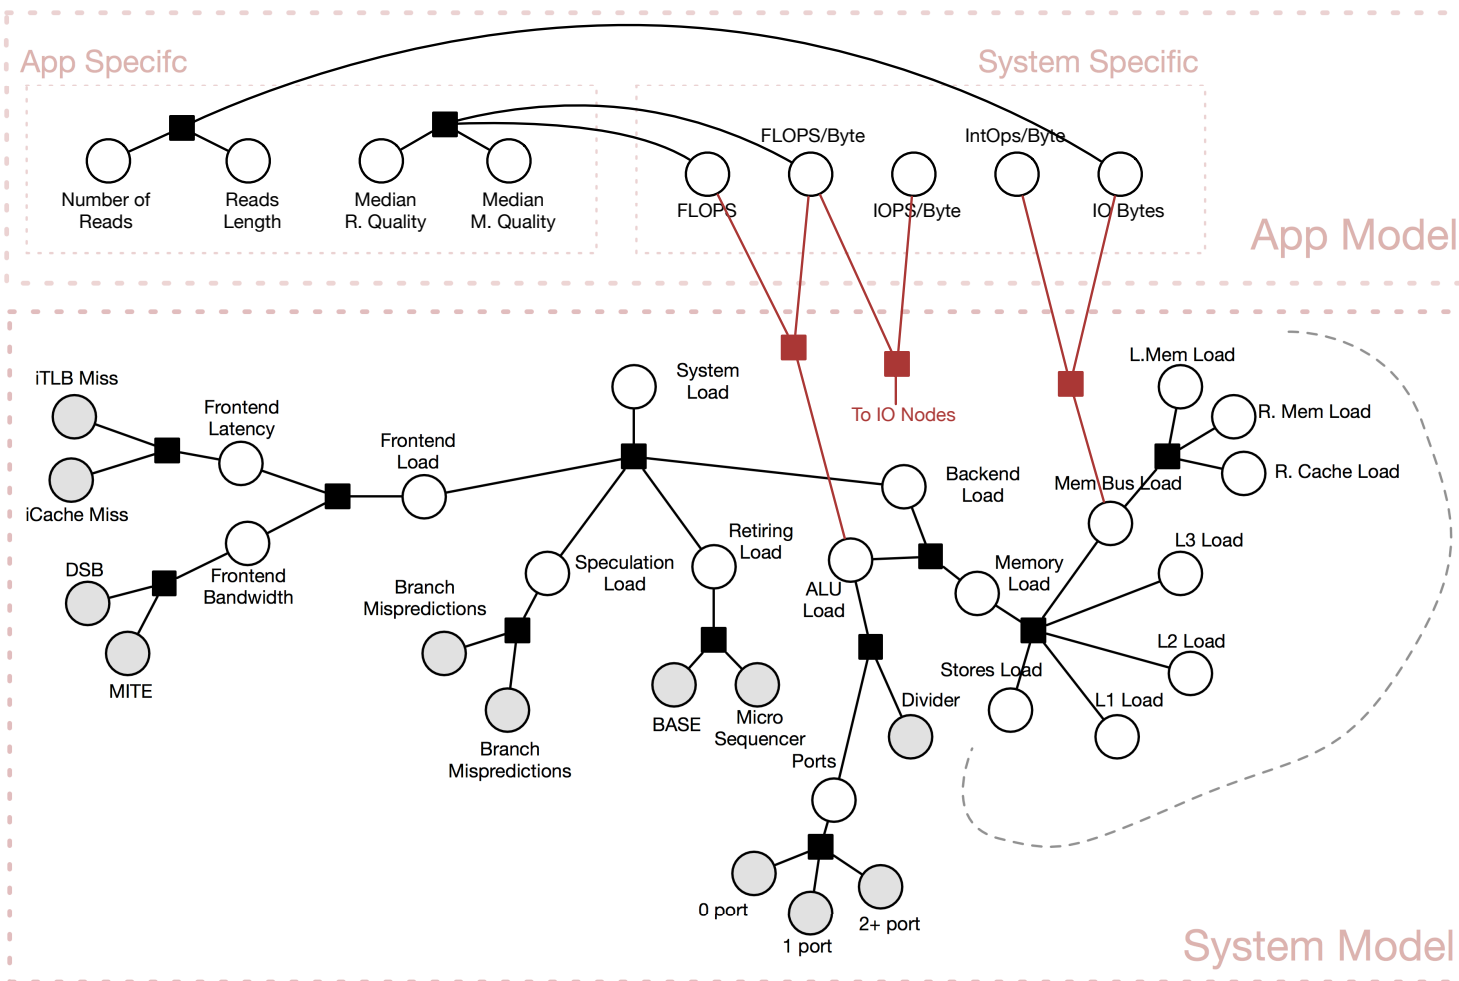
# Symphony: Scheduling across heterogeneous systems and clusters



**Scheduling Criteria:**

- Processor Affinity
- Data Locality
- Shared Resource Contention

- Using probabilistic inference and reasoning to
  - Inexact measurements
  - Time synchronization issues
  - Model processor architectures (Black box)



Task

Scheduling Criteria:

- Processor Affinity
- Data Locality
- Shared Resource C

# Questions?

- Email authors at [ssbaner2@illinois.edu](mailto:ssbaner2@illinois.edu)

