2015 IGEM Special Allocation Grant Proposal
Cover Sheet

Proposal Number: (Assigned by Board Office)	Amount Requested: $50,989
Name of Institution: Boise State University	Name of Industry Partner: Micron Technology, Inc.
Name of Institution Contact: Dr. Vijay Dialani; Dr. Tim Andersen	Name of Industry Contact: Dr. Richard Murphy
Phone Number: 208-426-5703 E-mail Address: vijaydialani@boisestate.edu	Phone Number: 208-368-3286 E-mail Address: rcmurphy@micron.com

Title of Proposed Project: ARC: Automata processor for Research in Computing

Project Description:

Modern scientific applications collect and process a large amount of data either from distributed sensor networks or simulation-based experiments. Use of heterogeneous computing resources, such as CPU+GPU, to accelerate application performance for compute intensive applications is well understood [1-3]. For accelerating data intensive applications, “memory wall” [5, 6] needs to be scaled. Automata processor (developed by Micron) [7] provides one such alternative architecture that accelerates data intensive applications that have a high density of conditional branching. Initial investigations [8, 9] in the field of bioinformatics have shown promising results, potential applications include pattern detection in graphs and unstructured data streams. To investigate the use of heterogeneous computing architectures this proposal requests HERC support to build a computing server to carry out data intensive research and to compare and contrast alternative computing architectures.

The proposed instrument would integrate high performance Von-Neumann architecture (CPU+GPGPU) with Non-Von-Neumann architectures such as automata processor [7] and PCIe mounted SSD drives into a single system. It is the first time ever that such a combination of Von-Neumann and non-Von-Neumann architecture with automata processor will be created. The proposed infrastructure would provide a novel capability particularly for computer scientists, enabling them to conduct state-of-the-art research in HPC systems, which is currently not feasible with available resources.

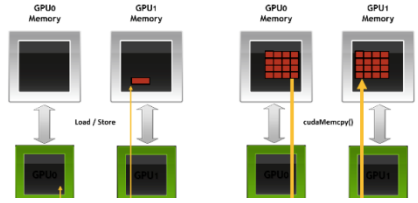

Figure 1: Direct GPU-to-GPU transfer technology from NVIDIA

High performance computing using combined CPU-GPU architecture has been found to be highlyeffective in the fields of information retrieval, visualization, database systems, computational material science and deep learning [4]. Many researchers are looking into ways the CPU-GPU combined architecture could help accelerate the algorithms in these fields. At present, a hybrid-programming model in which the CPU translates the virtual memory space to the GPU specific memory space is the most common usage pattern, but it highly restricts power, programmability, performance, and portability. We intend to use NVMe enabled devices that will allow direct data transfer between the GPU’s and SSDs and will free up the CPU from performing the memory space translations. To support data intensive computation, server will be equipped with 8TB of high-speed HDD drives, in addition to the 640GB of SSD capacity and 64GB local DRAM.

A complementary purpose for developing this machine is to compare and contrast the CPU-GPU combined Von-Neumann processing style against the non-Von Neumann architecture built using the automata processor [7]. Automata processors make it possible to implement codes with high density of conditional branching by using a Non-deterministic Finite Automata (NFA). Use of specialized processors would help avoid state-space explosion and exponential runtime complexity observed in CPU based implementations of NFA. An automata processor is available as an accelerator with a PCIe interface and is controlled and operated by the onboard CPU. It exists as a conglomeration of 6 distinct ranks, where each rank consists of 8 automata processor cores on a single chip. Each core consists of two half cores capable of containing 24K elements, where no connections can be made between the transition elements on half cores. Each automata core processes data at 1Gbps; assuming an 8-bit fundamental unit to represent state, this can be viewed as 128 * 10^6 characters per second for a combined update-latency of $7.81*10^{-9}$ seconds. These processors can be arranged in the pairs of 1-, 2-, 4-, and 8-cores, where grouped cores will receive the update from the same stream at 8Gbps. The proposed server will be equipped with 48 such cores, which could host a combined NFA of 1.17 million states. In addition, the machine will be instrumented to profile the hardware utilization and power consumption.

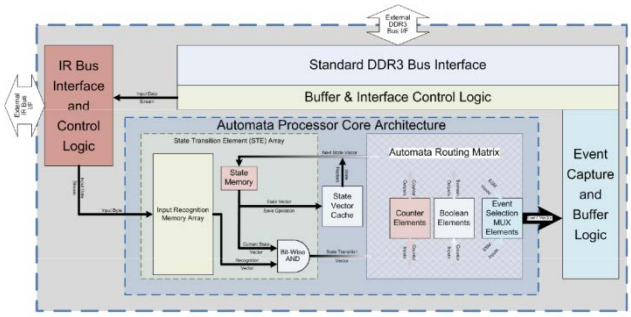

Figure 2: Internal Architecture of Automata Processor

The intended primary users of this cluster will be Boise State University researchers and students working in diverse areas of Computer Science (CS) such as Visualization, Database & Analytics, Machine Learning and Computer Architecture and their collaborators working in the application areas (of CS) such as computational chemistry, computational material science and bio-informatics located at BSU and WSU. These researchers will collaborate to jointly develop algorithms and techniques that predominantly address a fundamental CS research issue and apply them in the context of their discipline to speed up their calculations. Users will utilize the instrument for their externally funded research as well as undertake new research projects that would be made feasible only by the proposed server. This resource will also be made available to a wider research community to support research in applications of automata processors, e.g., researchers at U. Virginia are very keen on using such a resource for applications in bioinformatics and image processing. Making it available to researchers from such universities (as WSU and UVA) will also foster new collaborations and allow Boise State students to network with top-notch researchers from outside the state and further promote research activities within the state of Idaho. If funded, the server will be procured by 15th of June 2015.

Project Start Date: 3/1/15	Project End Date: 6/30/16
Total Project Budget: $50,989	Annual Project Budget: $50,989
Authorized Organization Representative Signatures	Date
Institution:	2/13/15
Partner: Not required	

Project Budget			
A. Personnel Cost (Faculty, Staff, Visiting Professors, Post-Doctoral Associates, Graduate/Undergraduate Students, Other)			
Name/Title	**Salary/Rate of Pay**	**Fringe**	**$ Amount Requested**
Dr. Vijay Dialani, Associate Professor	$64.76/hr for 87 hrs	31%	$7,381
Undergraduate researcher (to be hired)	$13/hr for 20 hrs/wk for 18 wks	7%	$5,008
% of Budget **24%**		**Subtotal**	**$12,389**
B. Equipment (List each item with a cost in excess of $1,000. Lump small items together)			
Item/description			**$ Amount Requested**
Automata Processing Board			$25,000
Server to host Automata Processing (Dell + Quad Core E2625 + Fusion IO SSD + 64 GB RAM + 2TB HDD			$13,600
			$38,600
C. Operating Expenses (Including but not limited to materials and supplies, consultant services, subcontracts, etc.)			
Item/description			**$ Amount Requested**
D. Other Costs			
Item/description			**$ Amount Requested**

E. Total Costs	
F. Total Amount Requested	**$50,989**

Institutional and Other Sector Support	
A. Institutional/Other Sector Dollars	
Source/Description	**Amount**
B. Faculty/Staff Positions (Description)	
C. Capital Equipment (Description)	
D. Facilities & Instrumentation (Description)	