

BrainScaleS

Resource Management and Monitoring


HBP Code Jam #8

2017-09-14 | Christian Mauch | Kirchhoff-Institute for Physics, Heidelberg University

- Resource Management  
How to manage fair exclusive hardware access for multiple users?
- Monitoring  
How to operate and maintain a computing platform?


## BrainScaleS – Single Chip (*HICANN*)

- Emulates AdEx neuron model
- Up to 512 analog neuron circuits, 112,640 synapses
- Accelerated time-scale  $10,000 \times$  biology
- Configurable parameters (neurons, ...)
- Flexible network topology




# BrainScaleS Wafer Module

- Wafer-scale integration
- 384 HICANN Chips
- $O(200K)$  neurons,  $O(43M)$  synapses
- On-Wafer digital-spike routing  
 $O(TEvents/s)$



# BrainScaleS Wafer Module



# BrainScaleS

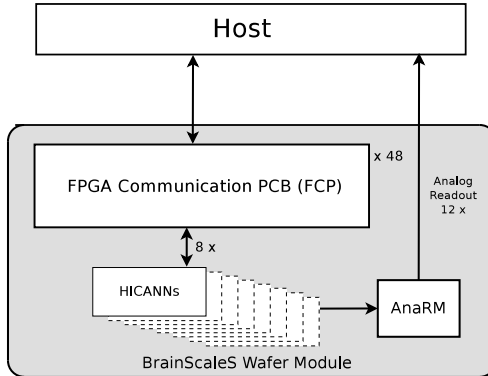


## Whole installation (NMPM1)

- 20 wafer modules
- 20 host machines
- 10 GBit/s uplink per module
- 40 GBit/s network backbone

# Resource Management

How to manage fair and exclusive hardware access for multiple users?



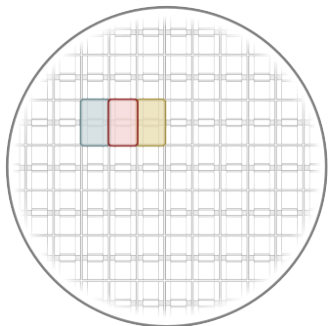
# Slurm Workload Manager

- Widely used job scheduler in HPC  
↳ good community support
- Handles exclusive resource allocation
- Provides fair resource distribution
- Supports highly customizable plugins






## Slurm for BrainScaleS



- Use Slurm licenses for each FPGA and analog readout board
- `srun -L W20F6,W20F7,W20F8, B201456 run_experiment`

# Slurm for BrainScaleS



- Allow access to Hardware only with valid licences  
↳ block access to other hardware parts on system level (iptables, etc.)

## Custom BrainScaleS Slurm Plugin

- Problem: very tedious licence request string for multiple resources
- e.g. `srun -L W20F0,W20F1,W20F2,W20F3,W20F4,W20F5,W20F6,  
W20F7,W20F8,W20F9,W20F10,W20F11,W20F12,B201456,  
B201457,B201458,B201459,B201461 run_experiment`

## Custom BrainScaleS Slurm Plugin

- Problem: very tedious licence request string for multiple resources
- e.g. `srun -L W20F0,W20F1,W20F2,W20F3,W20F4,W20F5,W20F6,W20F7,W20F8,W20F9,W20F10,W20F11,W20F12,B201456,B201457,B201458,B201459,B201461 run_experiment`
- Solution: Custom Slurm job submit plugin
- Use hardware database & coordinate translation  
`srun --wafer 20 --fpga 0,1,2,6 run_experiment`  
`srun --wafer 33 --hicann-with-aout 357 run_experiment`  
`srun --wafer 20,21 run_experiment`

# HBP Collaboratory - Neuromorphic Computing Platform

- Hardware access for collaboratory users
- REST API (web GUI and Python library)
- Platform jobs are replicated and synchronized to local BrainScaleS Slurm jobs

Workspace Job Manager

ID	Status	Platform	Code	Submitted on	Submitted by
Q2888	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:49:04	Jérémie Kuffen
Q2885	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:45:04	Jérémie Kuffen
Q2884	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:42:25	Jérémie Kuffen
Q2883	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:40:39	Arthur Erni-Antonioli
Q2882	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:37:14	Jérémie Kuffen
Q2881	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:34:51	Jérémie Kuffen
Q2880	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:31:05	Jérémie Kuffen
Q2859	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:29:16	Arthur Erni-Antonioli
Q2858	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:24:03	Arthur Erni-Antonioli
Q2857	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:22:28	Jérémie Kuffen
Q2856	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:20:47	Jérémie Kuffen
Q2855	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:19:21	Jens Egholm Pedersen
Q2854	Completed	BrainScaleS	Task 2.1   HMC 2017 Geneva, BrainScaleS tutorial ...	2017-09-13 17:15:08	Jens Egholm Pedersen

# Monitoring

How to operate and maintain a computing platform?

- Reliability, reproducibility, efficiency
- Myriad of components → myriad of error sources
- Over 16000 metrics for 20 wafer modules monitored
- Only few system administrators
  - ↳ Automated logging and alerting

# Monitoring - What we used before

- Ganglia Monitoring System
  - Main application is cluster monitoring
  - Rather archaic
  - Visualization not very customizable
- Graphite
  - Logging of time series data
  - Quick and easy lookup of data
  - Dashboard rather laborious



# Monitoring - Grafana + Elastic Search

## Grafana

- "Modern version" of graphite
- Easy and highly customizable dashboards
- Template support for dashboards
- Custom plugin support

## Elastic Search Plugin

- Provides search engine on text based data
- Log events like power up or errors in log files
- Combine time series data with events
- Compare actual vs target state





# Monitoring - Example



# Monitoring - Example



# Outlook

- Provide system status to users via neuromorphic platform interface
- Annotate platform jobs with monitoring metadata
- Splitting experiment setup (map & route) and execution phase to enable automated resource allocation
- Automated execution of tests in case of user experiment failures
- Provide support of federated operation (Unicore) in collaboration with the HPC platform (neural network emulation vs. analysis)