Energy-aware Design Space Exploration for GPGPUs

Pascal Libuschewski[1], Dominic Siedhoff[2], Frank Weichert[2]

[1] Computer Science 12 – Embedded Systems, TU Dortmund
[2] Computer Science 7 – Computer Graphics, TU Dortmund

Motivation

- General Purpose computing on Graphics Processing Units (GPGPU):
  - HPC systems
  - Cyber Physical Systems
  - Mobile devices
- Energy consumption should/must be considered
- Automatic tools needed for design space exploration of GPUs

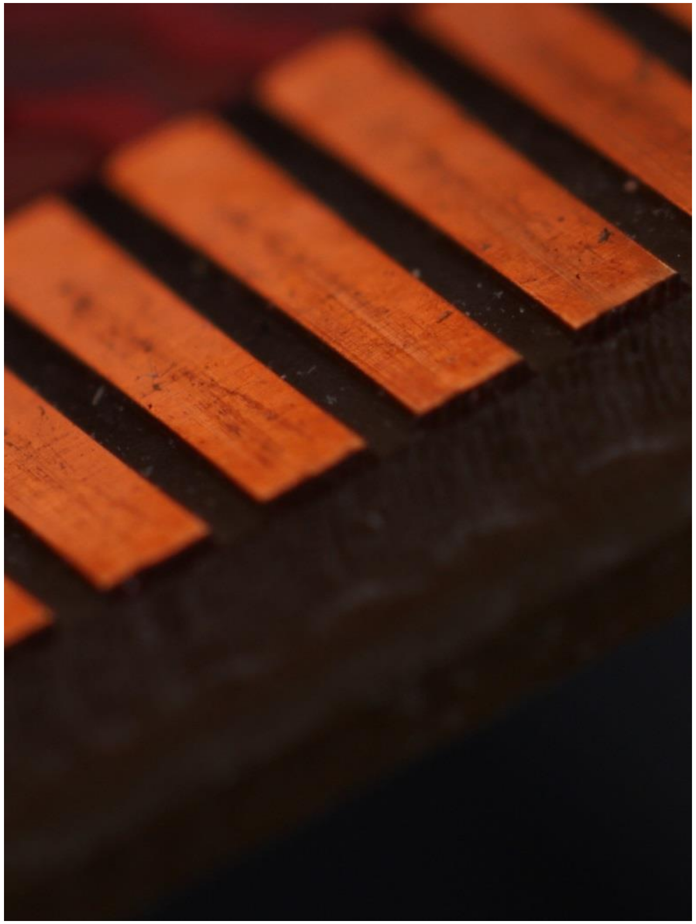

Arising Questions

- Which GPU is:
  - the most energy efficient?
  - the most power efficient?
  - the fastest?
  - ... for a given task?
  - ... for a set of tasks?
- What about real time contstraints?
- How to develop new energy efficient GPUs?

Background

- Application:
  - Image analysis in real time,
  - of physical sensor data,
  - with mobile devices
  - or centralized on a HPC system.

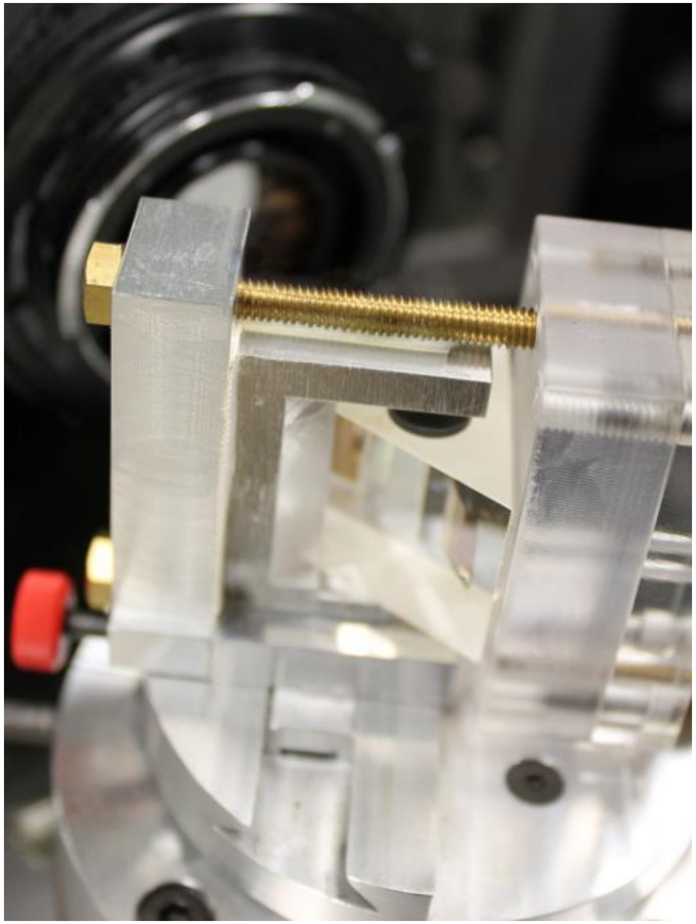

Which is the most energy efficient GPU?

Automatic Design Space Exploration

- Testing different GPUs is both:
  - cost-intensive
  - time-consuming
- Instead: Using a simulator to determine power- or energy-consumption

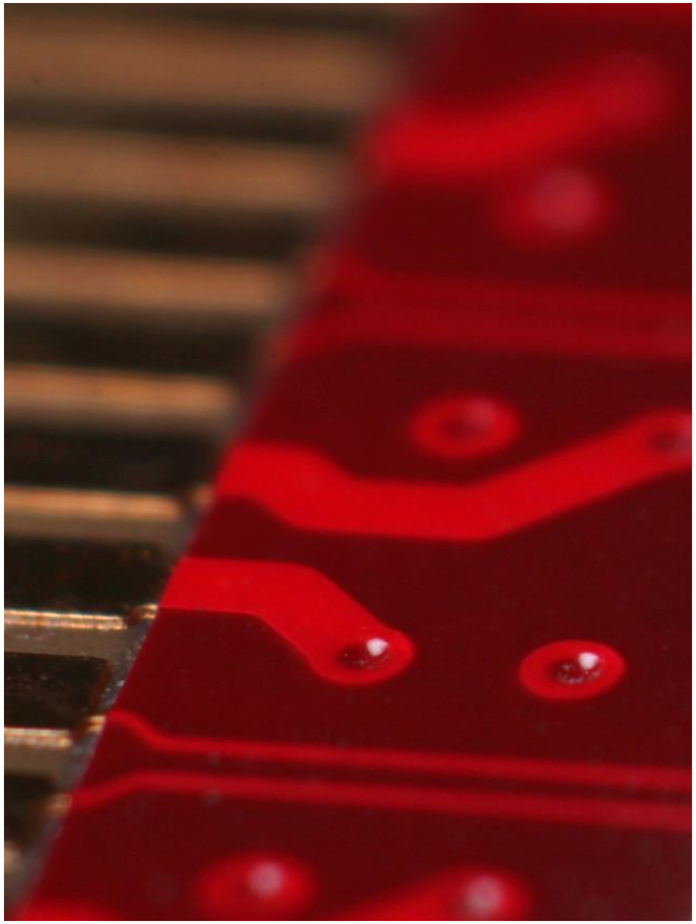

Automatic Design Space Exploration

- Design space exploration framework:
  - Cycle accurate simulator
  - Energy model for GPUs
  - Genetic Algorithm

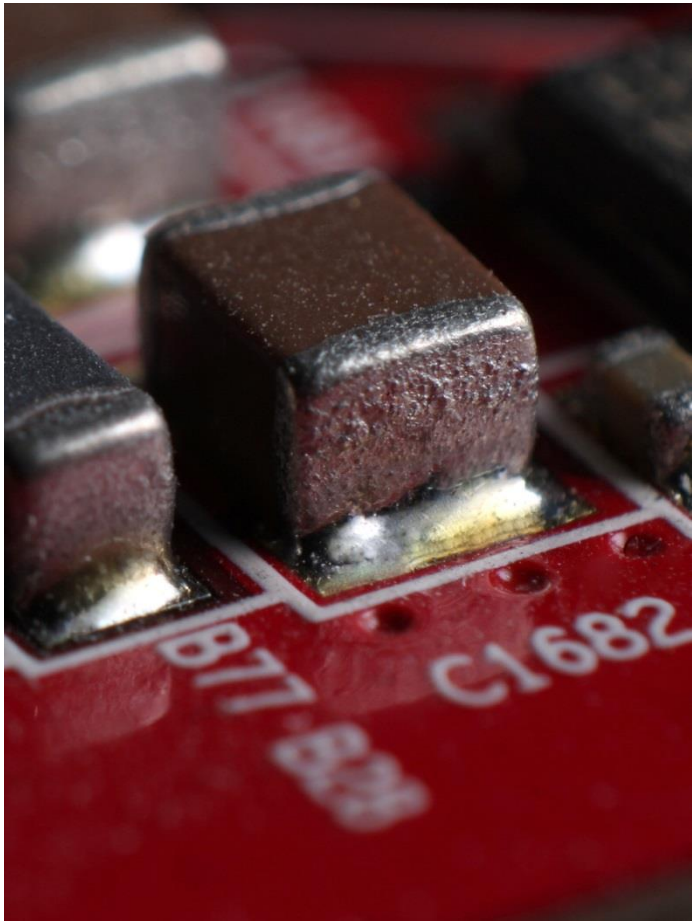

GPGPU-Sim

- GPGPU-Sim [1] is a cycle accurate simulator to simulate GPGPU code

GPGPU-Sim Pipeline

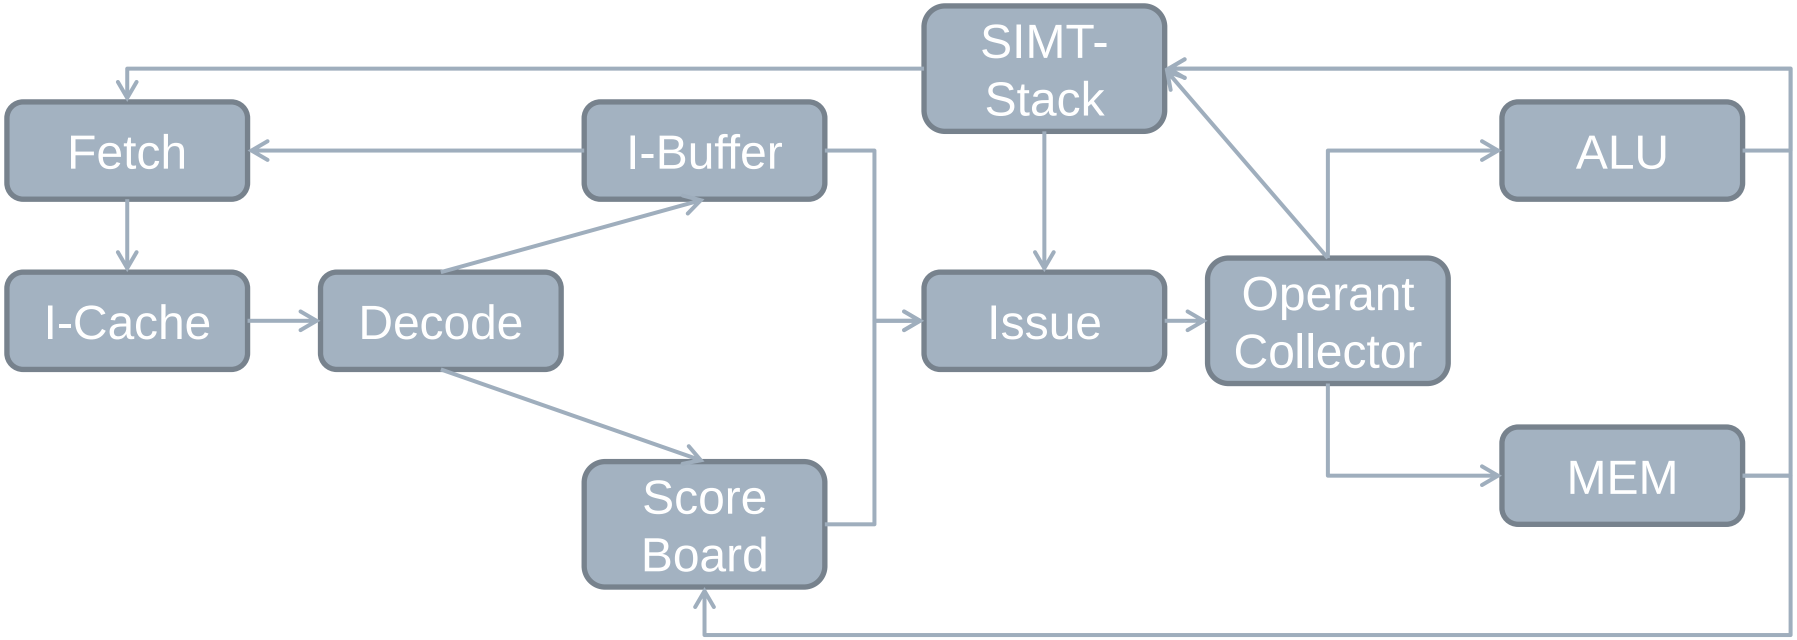

[1] Bakhoda, et al.: Analyzing CUDA workloads using a detailed GPU simulator. ISPASS 2009.

GPUWattch

- GPUWattch [2] is used for the power model

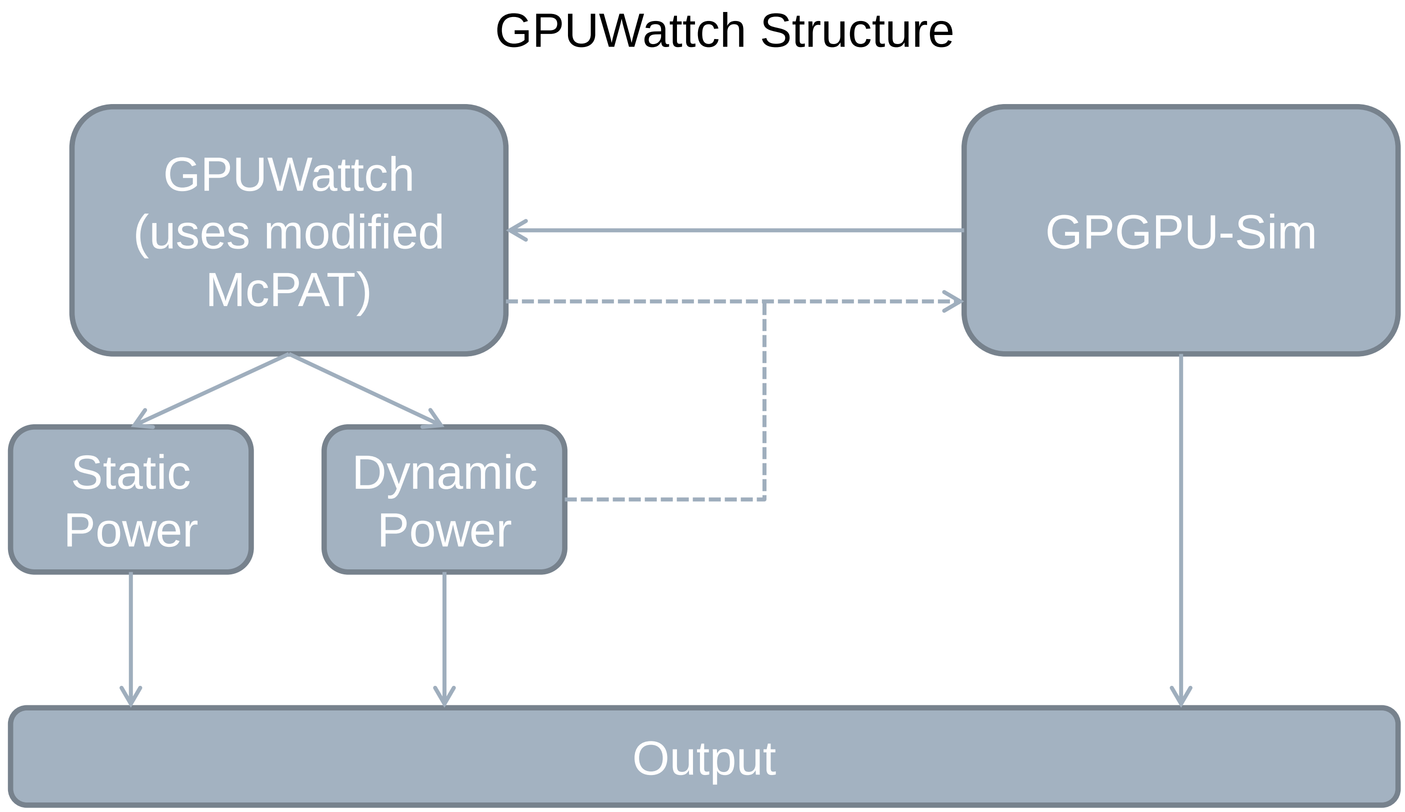

[2] Leng, J. et al.: Gpuwattch: Enabling energy optimizations in gpgpus. International Symposium on Computer Architecture 2013

Genetic Algorithm

- GPU configuration is coded in genes and chromosomes
- Single GPUs are represented by individuals
- All individuals form a population
- Individuals are evolved in an evolutionary process

Chromosomes

- Optimization of ~50 different hardware parameters
- Parameters are identified by name convention and type

```
GPGPU-Sim parameter:
-gpgpu_n_clusters 12

First gene in chromosome:
pop.subpop.0.species.gene-name.0 = -gpgpu_n_clusters
pop.subpop.0.species.gene-type.0 = GPGPU_SIM_TYPE
```

Evolution Process

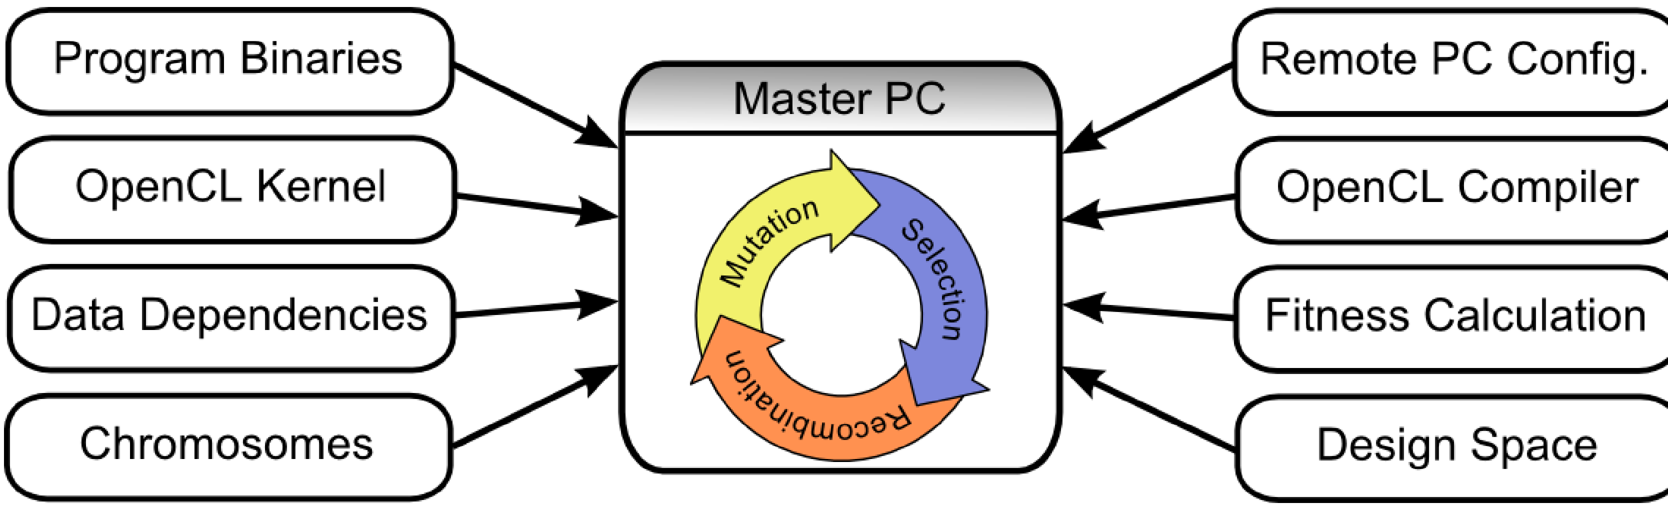

Evolution Process

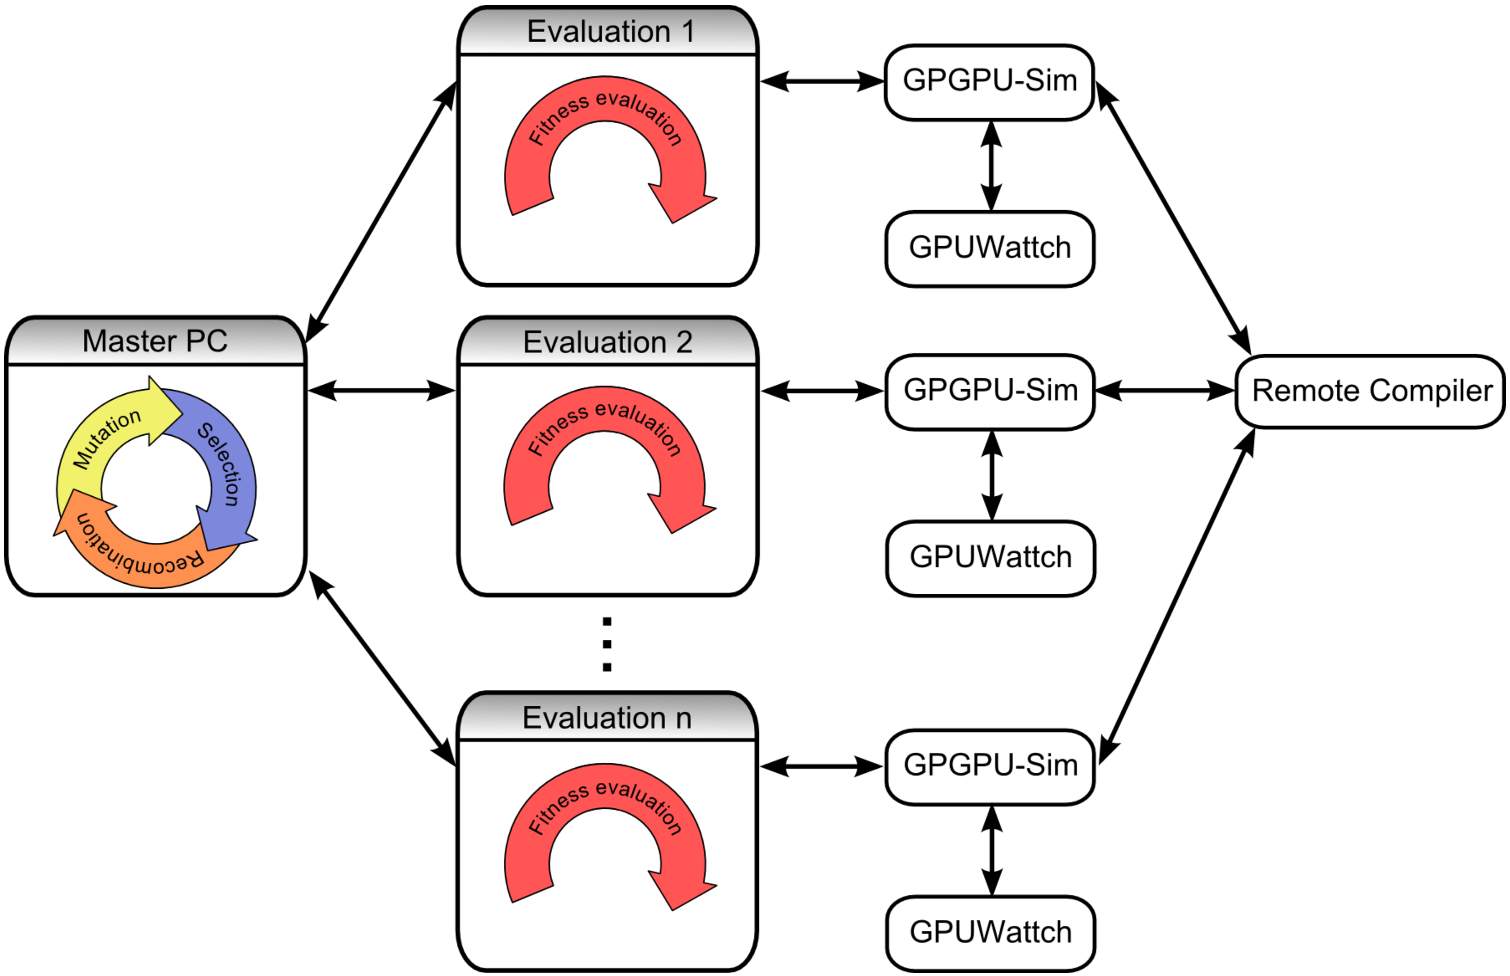

Evaluation Setup Example

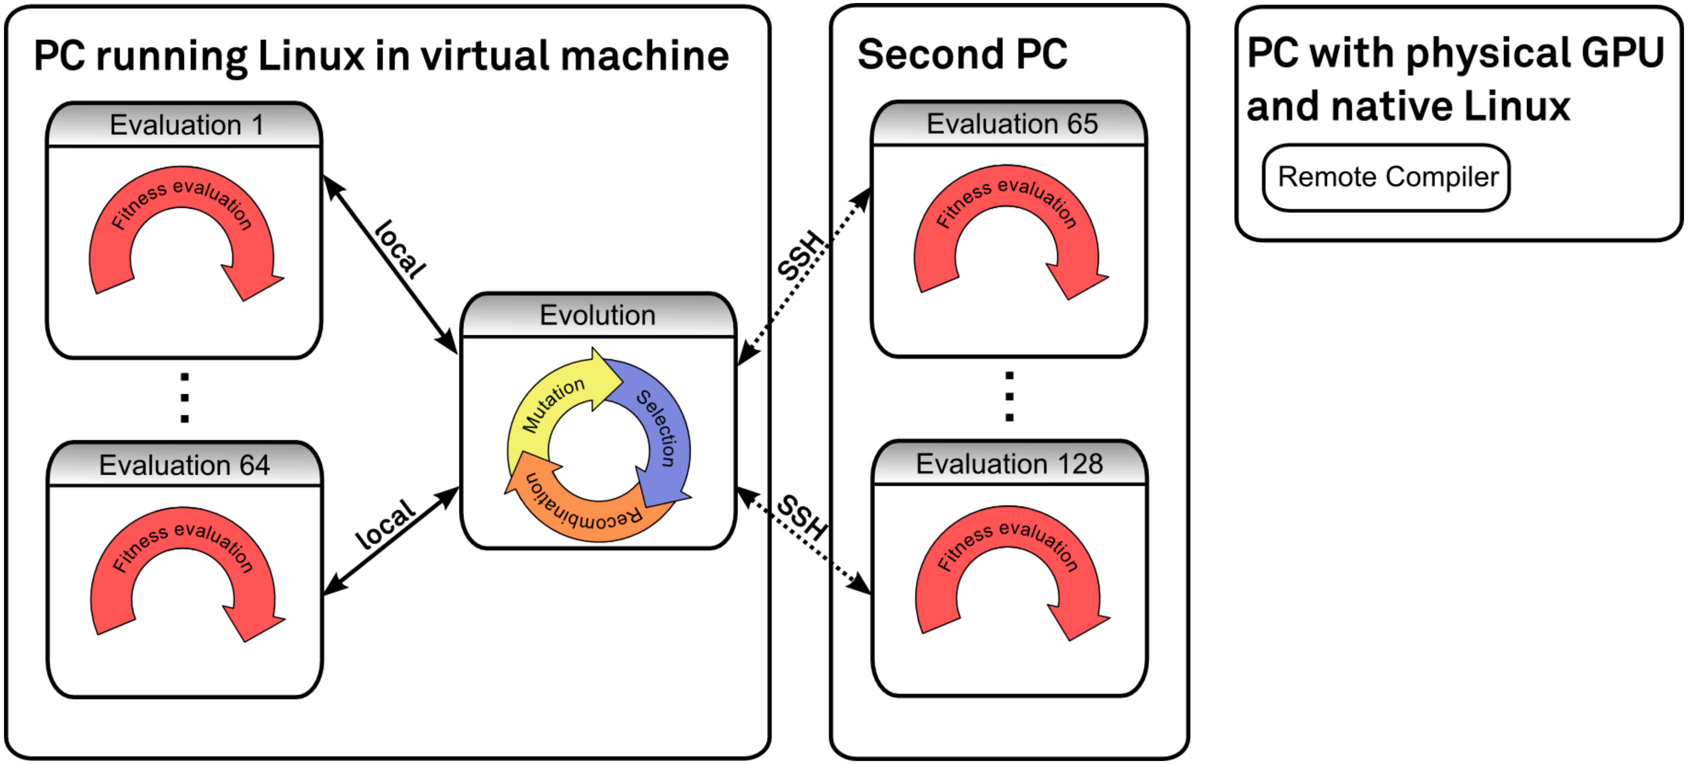

Evaluation Setup Example

- 19 cluster nodes
- Each with 64 AMD Opteron cores
- Each with 265GB RAM
- Evaluation can run e.g. on 2 nodes

Evaluation Setup Example (2)

- Evaluation process can run on heterogeneous machines
- Virtually no limits

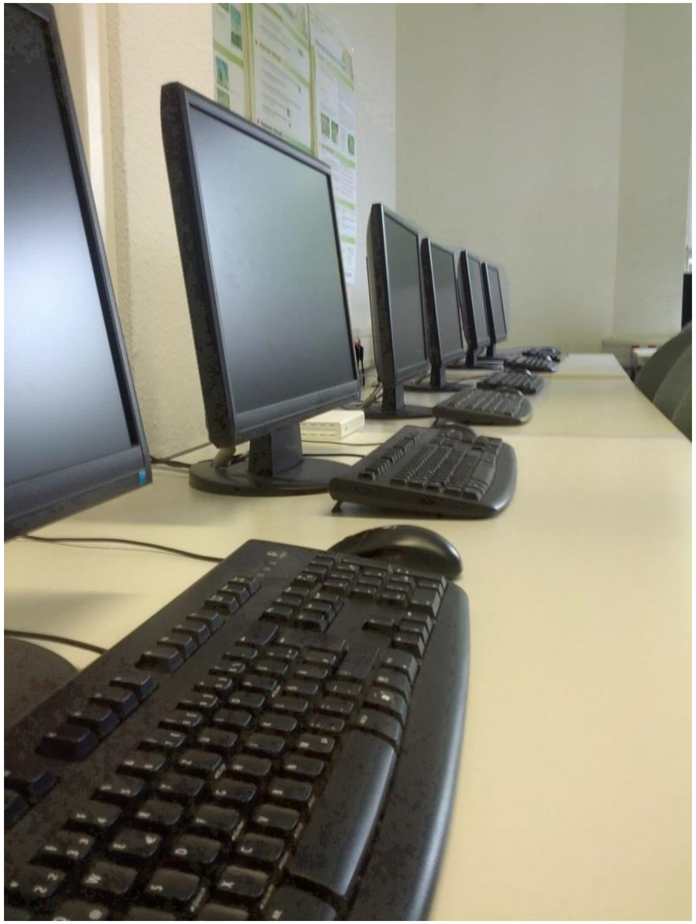

Evaluation Setup Example (2)

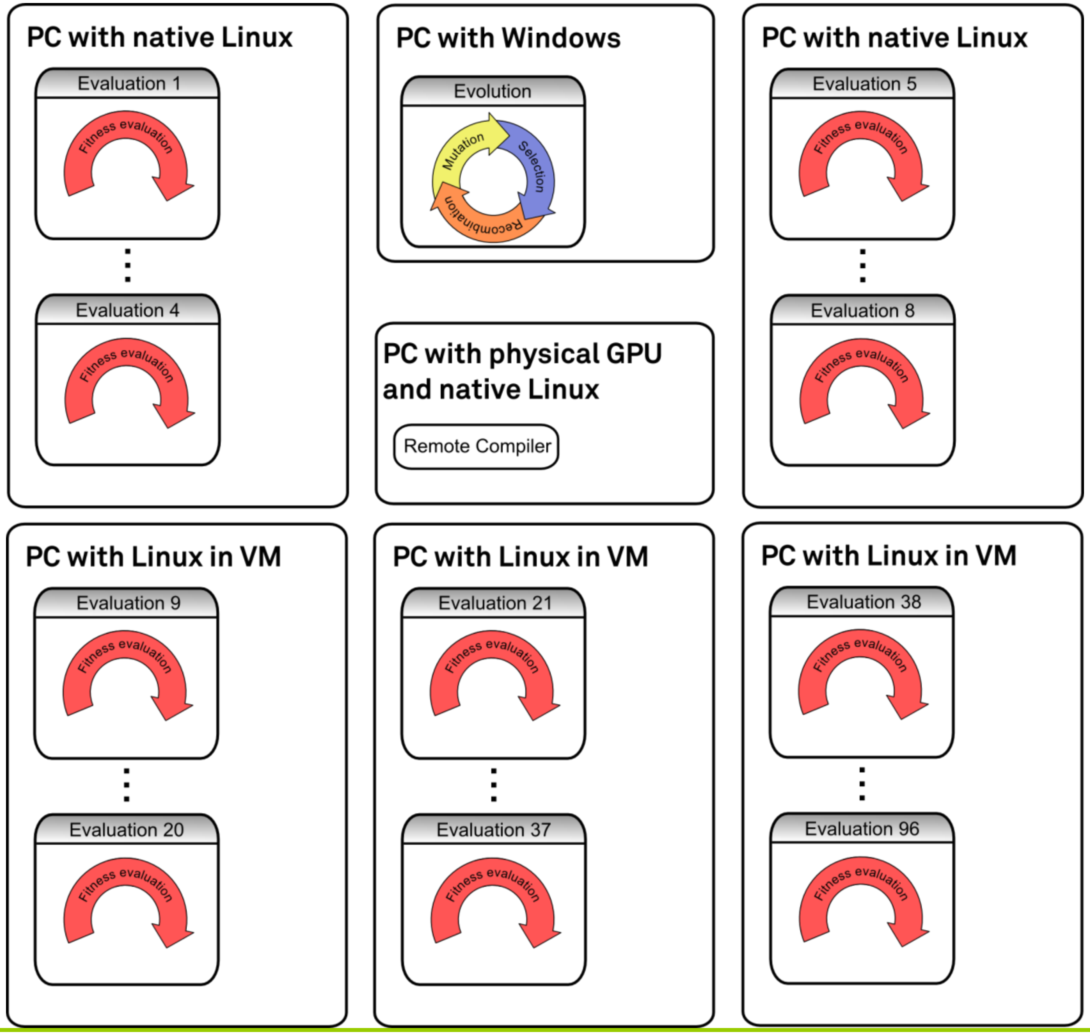

Design Space

Architecture	# SMs	Core clock	DRAM clock	# Registers
GF-100-108	1-15	590-810 MHz	1600-4000 MHz	16-64k

```
pop.subpop.0.species.min-gene.0 = 1
pop.subpop.0.species.max-gene.0 = 15

pop.subpop.0.species.min-gene.1 = 590
pop.subpop.0.species.max-gene.1 = 810

...
```

Design Space

- Unreasonable design space:

  380 billion samples

- Would result in non-existing GPUs
- Not every MHz sample is needed

Design Space

Architecture	# SMs	Core clock	DRAM clock	# Registers
GF-100	11-15	610-780 MHz	3200-4000 MHz	32-64k
GF-104	6-7	650-675 MHz	3400-3600 MHz	32-64k
GF-106	3-4	590-790 MHz	1800-4000 MHz	16-32k
GF-108	1-2	700-810 MHz	1600-1800 MHz	16-32k

```
#GF-100
pop.subpop.0.species.design-space.0.min.0 = 11
pop.subpop.0.species.design-space.0.max.0 = 15
pop.subpop.0.species.design-space.0.sampling.0 = 1

pop.subpop.0.species.design-space.0.min.1 = 610
pop.subpop.0.species.design-space.0.max.1 = 780
pop.subpop.0.species.design-space.0.sampling.1 = 10
```

Design Space

Architecture	# SMs	Core clock	DRAM clock	# Registers
GTX 480	15	806 MHz	4568 MHz	32k
GTX 480	15	700 MHz	3696 MHz	32k
GTX 480	15	607 MHz	3348 MHz	32k
...				

Design space can be modeled as fine-granulated as needed.

Evaluation

- A subset of the Rodinia [3] benchmark was used for evaluation
- Average power P_{avg} is provided by GPUWattch
- A Geforce GTX 480 was taken for reference

[3] Che, S., et al.: Rodinia: A benchmark suite for heterogeneous Computing. IISWC 2009

Power Savings

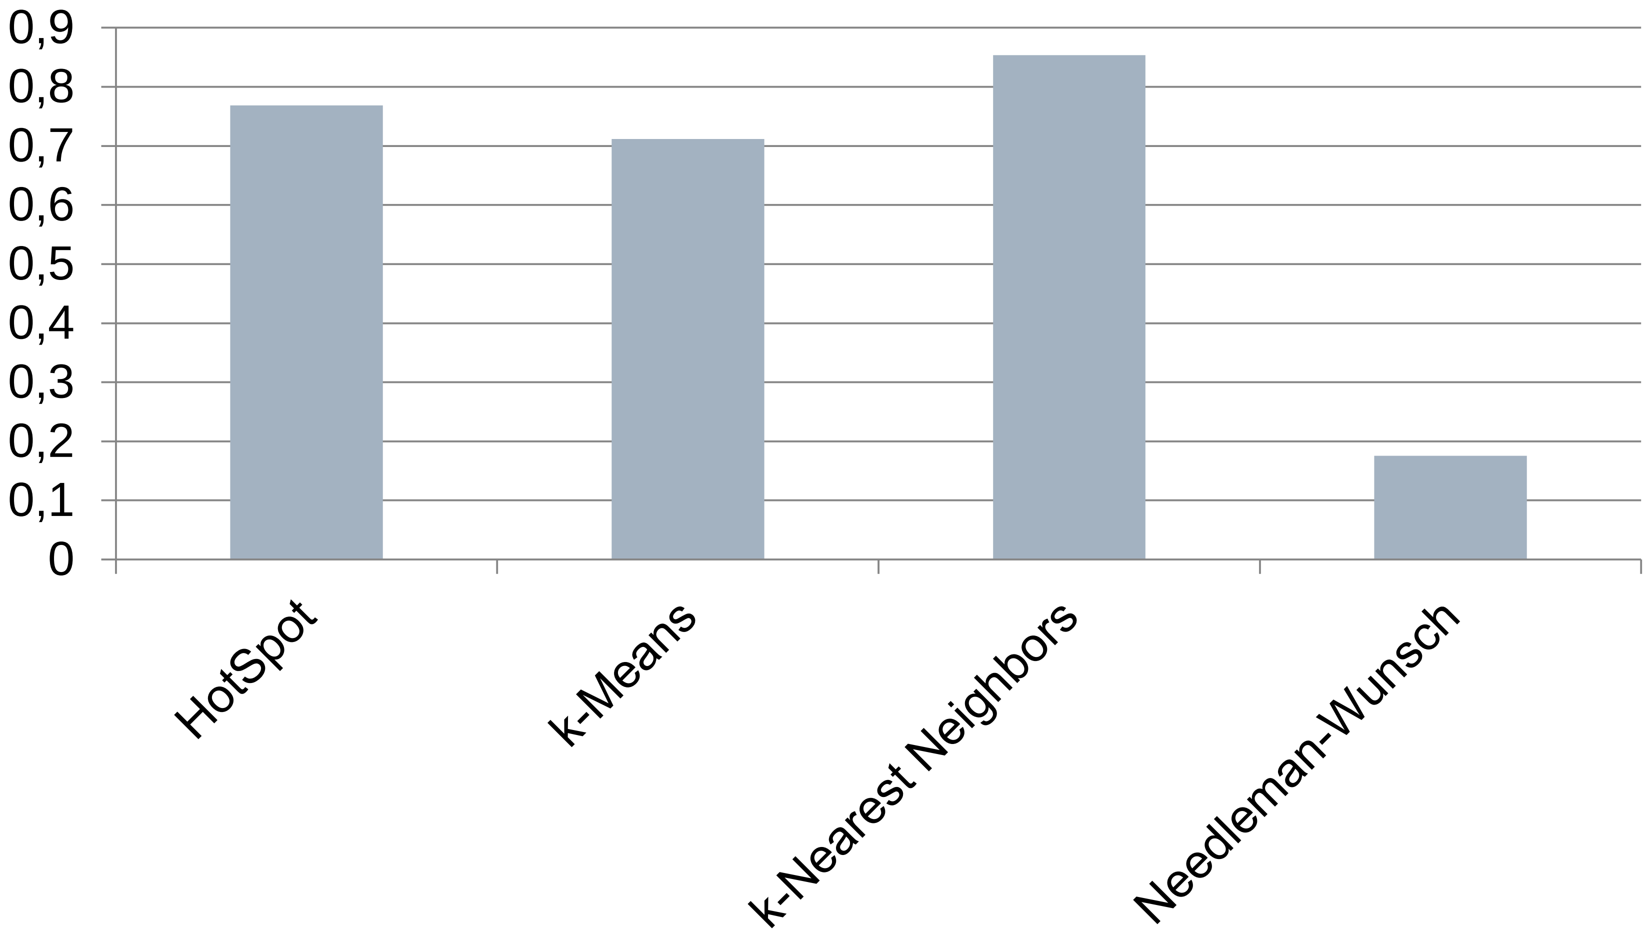

Evaluation – Energy Savings

- Number of needed cycles C provided by GPGPU-Sim
- Average power P_{avg} provided by GPUWattch
- Number of cycles per second C_s

$$E = \hat{E} = \int P dt = \frac{P_{avg} C}{C_s}$$

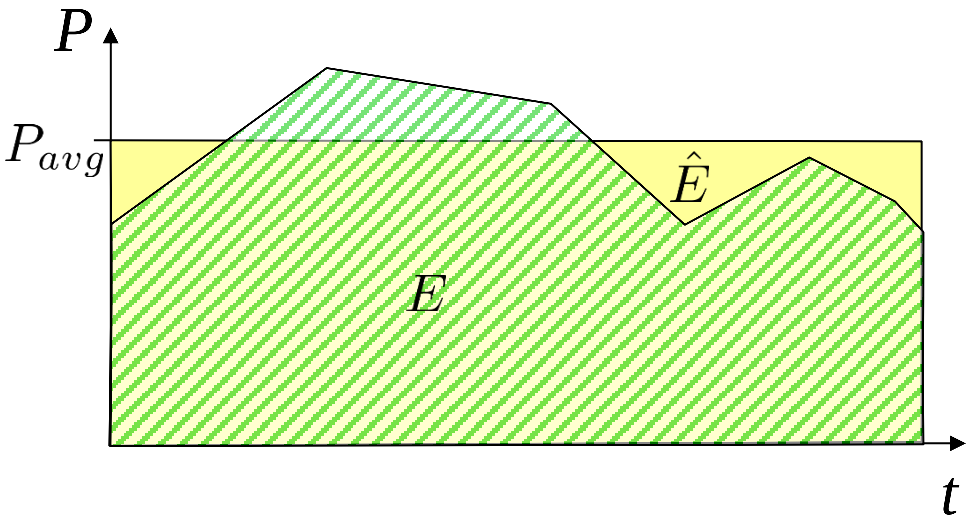

Energy Savings

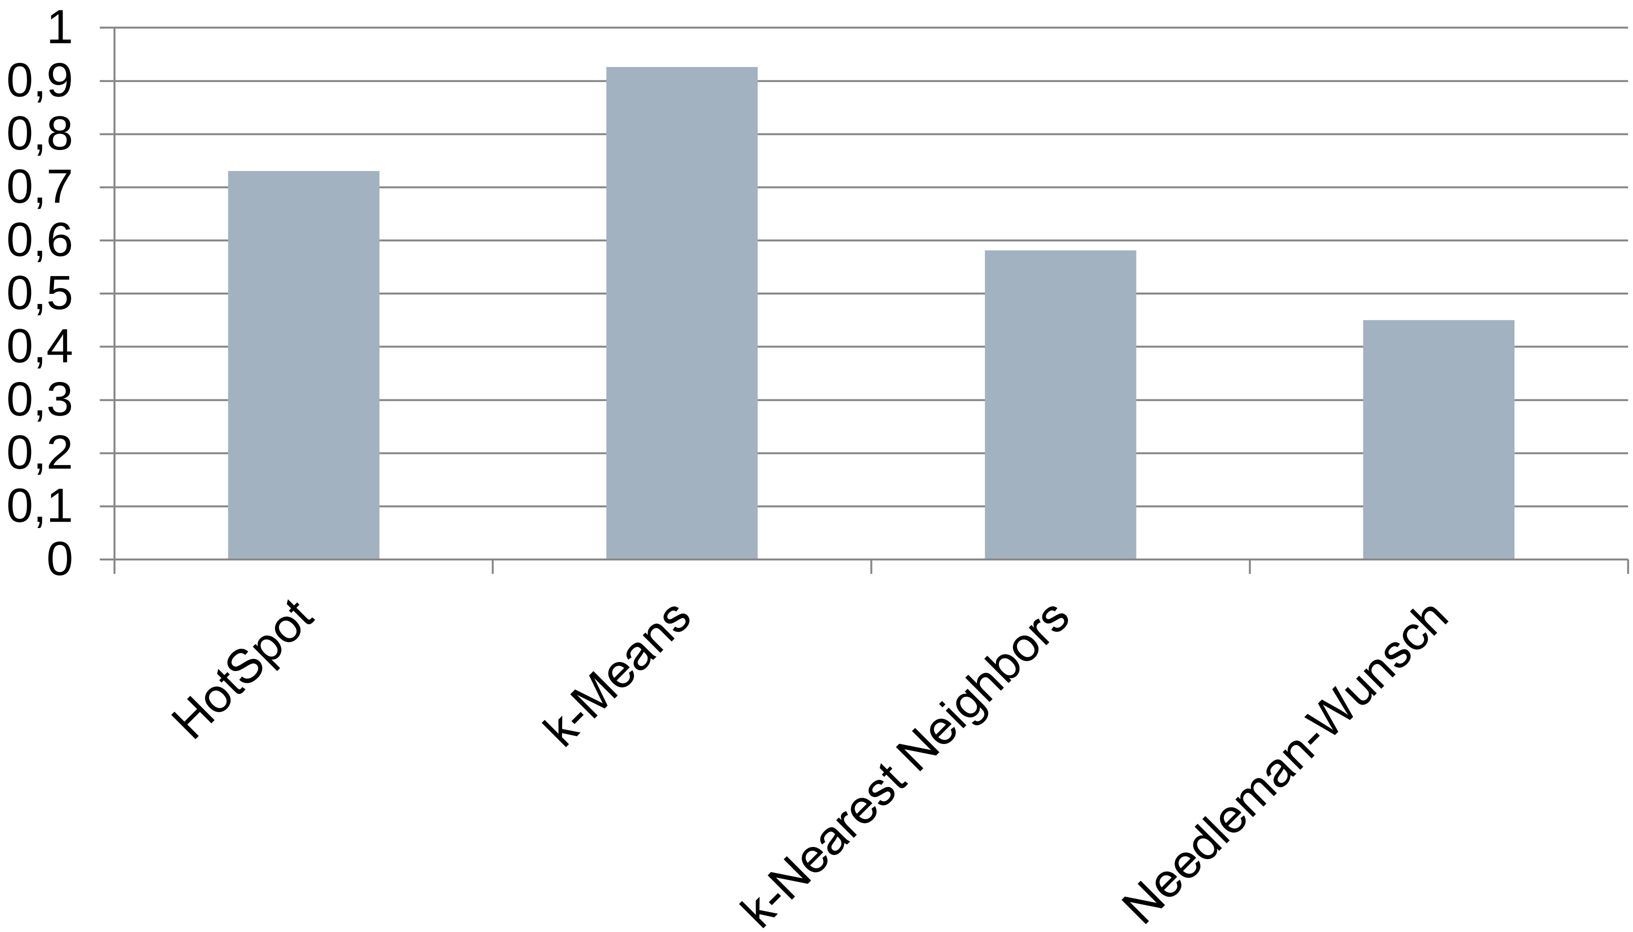

Energy Saving Configurations

Benchmark	Architecture	# SMs	Core clock	DRAM clock	# Registers
HotSpot	GF-106	3	790 MHz	1400 MHz	18k
k-Means	GF-108	2	710 MHz	820 MHz	32k
k-NN	GF-100	14	780 MHz	1600 MHz	32k
Needleman	GF-100	11	610 MHz	2000 MHz	32k
(Reference)	GF-100	15	700 MHz	1400 MHz	32k

Future Work

- Verification with our measurement system
- Multiobjective optimization
- Hardware/software co-design

Summary

- Automatic design space exploration for GPUs
- Different objectives possible
- Design space as desired

- Applications:
  - Best GPU for HPC systems,
  - for mobile devices,
  - or for Cyber Physical Systems.
  - Development of new GPUs