#### MARS Framework

#### A Framework for resource management



#### A Framework for resource management



# Example: reflective task mapping



## MARS usage

• Crating a resource manager with two policies:

```
class Simple_Map_Policy : public Policy_Manager {
    void setup() {
        registerPolicy (new Simple_Map_Policy()), 150 );
        registerPolicy (new Simple_DVFS_Policy(), 50 );
    }
};
```

## MARS interface to user policies

sense - returns sensed data for the latest execution period.

E.g.: for a policy that executes every 50ms sense<SEN\_INSTR\_TOTAL>(task) returns the total number of instructions in the last 50ms

• *actuate* - sets an actuation knob

E.g.: changing a task mapping: actuate<ACT\_TASK\_MAP>(task,core.id);

- senself returns predicted sensed data for the next execution period.
- tryActuate "What if I perform an actuation ?"
  - tryActuate affects senself results but it does not actually sets physical actuation values

# Simple\_Map\_Policy implementation

• If a task performance drops below stablished goal, searches for a new task mapping



### **Reflection process**



## MARS deployment



## Implementation on Linux



## Offline implementation



## Offline simulation flow



#### Case studies

- Energy efficient task mapping for HMPs
- Managing chip wear-out on mobile platforms
- Design space exploration of HMPs

## Case study: energy efficient task mapping

- SPARTA: Runtime Task Allocation for Energy Efficient Heterogeneous Many-cores
  - Light-weight QoS-aware task mapping
    - maximize the energy efficiency given performance constraints
  - Orthogonal to other policies
    - Performance/Power predictive models take in to account the behavior of the Linux's **DVFS governor** and **CFS policy**.

Muck, T., Sarma, S., Dutt, N. (2015). Run-DMC: Runtime dynamic heterogeneous multicore performance and power estimation for energy efficiency. CODES+ISSS '15

Donyanavard, B., Mück, T., Sarma, S., Dutt, N. (2016). **SPARTA: Runtime Task Allocation for Energy Efficient Heterogeneous Many-cores**. CODES+ISSS '16





cores: Huge Big Medium Little tasks:  $t_2$   $t_1$   $t_3$  idle

- Due to workloads not being 100% cpu bound (i.e. interactive apps)
- Memory latency or throughput bound workloads

• Opportunity to move task to a slower core and save energy





## SPARTA on MARS framework



Performance/Power model predicts IPS and power given new mapping and frequencies

- Offline trained model predicts per-task IPS/Power
- Scheduler model predicts resource utilization

List scheduling heuristic to explore new mappings

- **tryActuate/senself** to find which cores meet QoS with highest energy efficiency
- Map tasks in order of energy efficiency