



Power Budget-Aware Distributed Energy Minimization of Many-Core Systems

Contributors: Southampton and Newcastle Lead: RA Shafik and S Yang (Southampton)

Aim: Scalable energy minimization through decentralised concurrency and DVFS controls, guided by the power budget annotation

Key Messages:

This demonstrator illustrates the **PRiME** approach to distributed energy minimization of many-core systems. The approach uses power budget annotation in the application to constrain the max. power consumption. For a given power budget, the energy is reduced through multiple runtime agents:

- Each application thread is managed by a **power trading agent (PTA)**, which manages inter-thread power/performance trade-offs
- Each processor core is managed by a **DVFS controller agent (DTA)**, which reduces the energy based on the power budget from the PTAs.

Economic trading is used between PTAs for power/performance trade-offs, and model predictive control algorithm is used to control concurrency (core allocations) and DVFS to reduce energy minimization within the given power budget, while achieving the best possible application performance.

Collaborative Working:

This demonstrator uses Newcastle's power model, which guides the concurrency control and power/energy estimation

How the demonstrator works:

The demonstrator runs on an Intel Xeon E5-2630, hosting 12 cores and a 32GB shared memory. Each core operates at a minimum frequency, $f=1.2$ GHz (at $V_{dd} = 0.98V$) and a max. $f=2.6$ GHz (at $V_{dd} = 1.35V$), in total thirteen frequency levels. A PARSEC raytrace application is executed on a Linux 2.6.32 OS as a case study. The PTAs are implemented using modified *p_thread* wrapper implementation, while DTAs are implemented as *Linux* power governors.

