PRiME
run smarter, live longer
Research Brief
Power-efficient, Reliable, Many-core Embedded systems www.prime-project.org

Energy Efficient Thermal Management of Embedded Systems

***PRiME** is a five year EPSRC funded research programme (2013-2018), in which four UK universities address the challenges of **power consumption** and **reliability** of future high-performance embedded systems utilising many-core processors.*

Hierarchical Run-time Management for Energy-Efficient Thermal Management

A major challenge for modern multicore embedded systems is decreasing lifetime reliability due to elevated operating temperatures caused by high power densities. This leads to an acceleration of device wear-out. **Thermal management** is, therefore, emerging as a major design consideration for mobile computing, alongside energy consumption which translates directly to battery life. Effective management of the energy consumption and temperature requires dynamic control of the hardware (e.g. voltage-frequency scaling of the multi core CPU) and the software (e.g. allocation of application threads), in response to an executing application's performance requirement, its cross-layer interactions with the operating system/hardware and the resulting thermal profile.

PRiME has investigated three main thermal parameters – peak temperature, average temperature and thermal cycling. Two control levers are used – CPU voltage-frequency and thread allocation, managed using the operating system APIs ***cpufreq*** and ***affinity***, respectively. The underlying control algorithm for managing energy and temperature is based on reinforcement learning principles, specifically Q-learning. The overall run-time approach is implemented in a hierarchical manner:

- CPU voltage-frequency is controlled from the lower hierarchy at a shorter interval to manage average temperature, peak temperature and energy consumption.
- Affinity is controlled from the upper hierarchy at longer intervals to manage thermal cycling.

Demonstrator: nVidia Jetson

The adaptive and hierarchical approach has been validated by implementing a run-time manager for embedded Linux running on an nVidia Jetson development board. This platform integrates a Tegra K1 System-on-Chip (SoC) with quad-core ARM A15 CPU and 192 GPU cores (see figure 1).

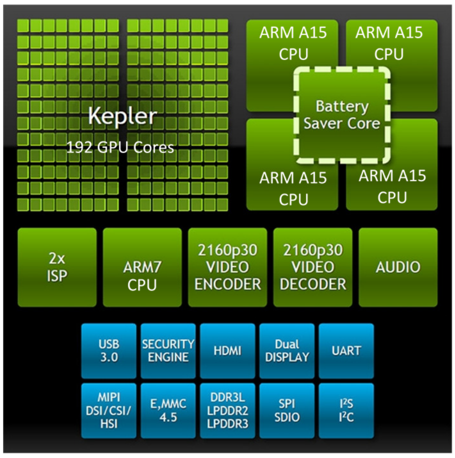

Figure 1. nVidia Tegra K1 block diagram

The SoC supports 22 frequency levels and integrates a thermal sensor for accurate on-chip temperature measurement.

The work to date has primarily focussed on energy and thermal management of the multicore ARM CPUs, but in future, this approach will also be extended to the GPU cores.

The PRiME demonstrator uses "*FFmpeg*" application software to illustrate thermal run-time management. *FFmpeg* is the most common video decoding application in mobile and embedded platforms. In the demonstrator, the embedded Linux plays a 24 frames per second (fps) 1080p video using the *FFmpeg* application.

Application hook-ups and embedded Linux modifications are illustrated in figure 2.

PRiME
run smarter, live longer

EPSRC
Engineering and Physical Sciences Research Council

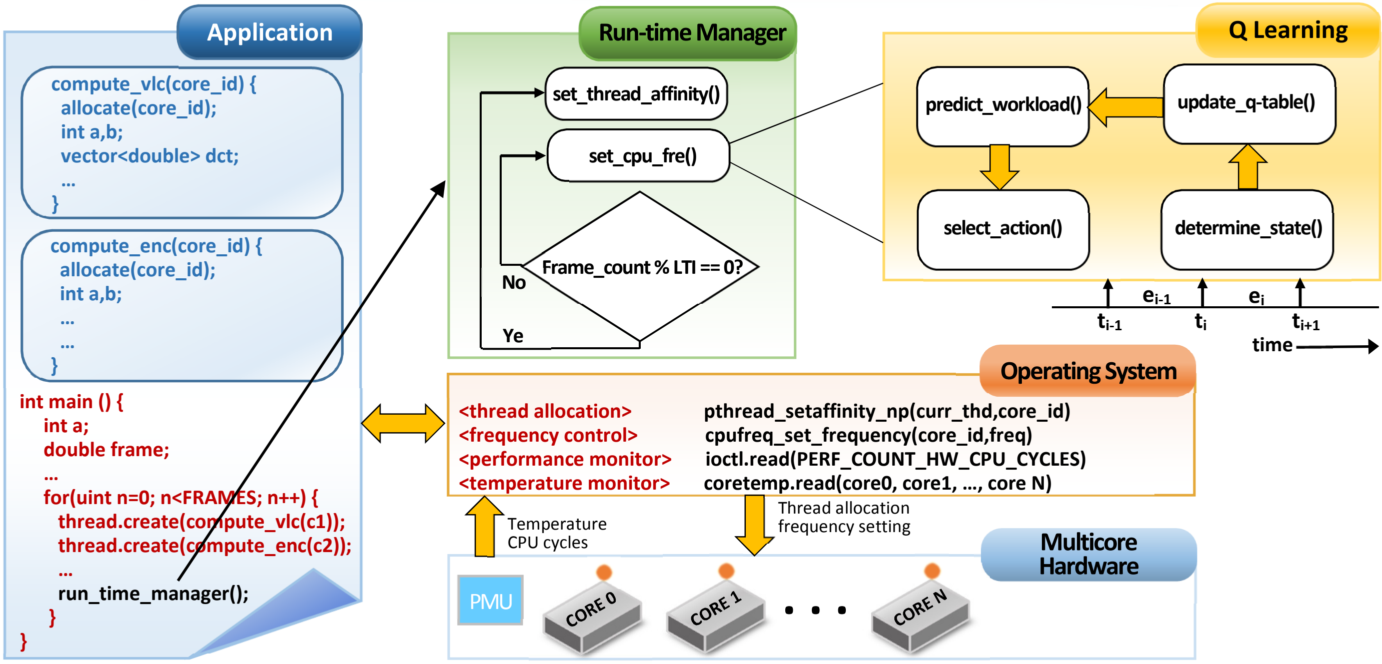

Figure 2. Application hook-ups and embedded linux modifications

Results

Compared to the Linux **ondemand** governor, the following are the key results:

1. thermal improvements:
   - **14 degC** in average temperature
   - **16 degC** in peak temperature
   - **54%** in thermal cycling
2. energy reduction of average **15%**
3. scalability

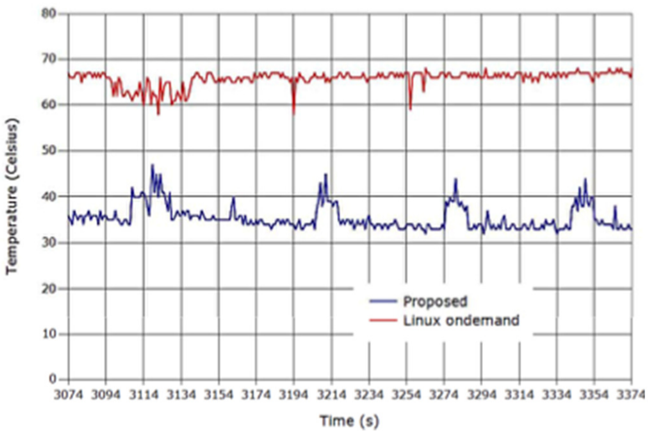

Figure 3. Thermal performance – PRiME vs. linux ondemand

Figure 3 shows the thermal behaviour (blue) compared with Linux's **ondemand** governor (red). The overhead of the hierarchical approach is shown below.

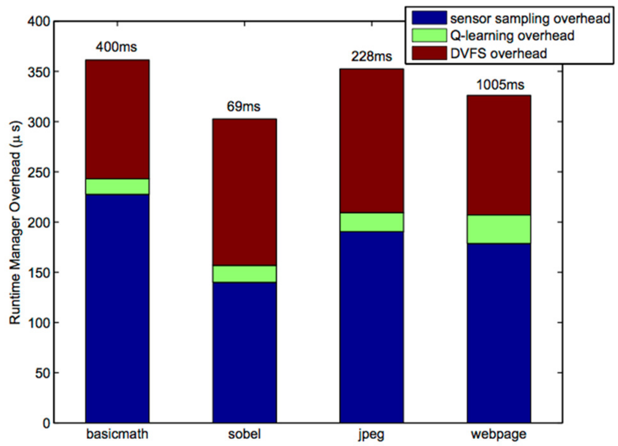

Future Work

Ongoing work will extend this thermal management approach to GPUs and other heterogeneous elements

More information

Visit the PRiME programme web site, including the opportunity to sign-up for programme updates, or contact the PRiME Impact Manager:

Gerry Scott
Email: gerry.scott@soton.ac.uk
Tel: +44 (0)23 8059 2749

www.prime-project.org

Imperial College London | Newcastle University | MANCHESTER 1824 The University of Manchester | UNIVERSITY OF Southampton