This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered and offers several algorithms and a lightweight co-run performance and power predictive model for computing the performance bounds of the optimal co- schedules and finding appropriate schedules. Expand

This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered and offers several algorithms and a lightweight co-run performance and power predictive model for computing the performance bounds of the optimal co- schedules and finding appropriate schedules. Expand

Heterogeneous processors with architecturally different devices (CPU and GPU) integrated on the same die provide good performance and energy efficiency for wide range of workloads. However, they also create challenges and opportunities in terms of scheduling workloads on the appropriate device. Current scheduling practices mainly use the characteristics of kernel workloads to ???

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

This subsection discusses three current scheduling strate-gies for CPU+GPU co-scheduling and Figure 2 shows the overview of these strategies. 2.2.1 Static Scheduling The traditional scheduling strategy is static scheduling. It sets the ratio of performance between the CPU and GPU statically and partitions workload according to it at the be-

This paper provides detailed infrared imaging results that show the impact of mapping decisions on the thermal and power profiles of CPU+GPU processors and proposes techniques to characterize the OpenCL kernel workloads during run-time and map them on appropriate device under time-varying physical and CPU load conditions. Heterogeneous ???

Youngmoon Lee, Kang G. Shin, and Hoon Sung Chwa. 2019. Thermal-aware scheduling for integrated CPUs-GPU platforms. Qi Zhu, Bo Wu, Xipeng Shen, Li Shen, and Zhiying Wang. 2017. Co-run scheduling with power cap on integrated CPU-GPU systems. An Evaluation Framework for Dynamic Thermal Management Strategies in 3D MultiProcessor ???

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

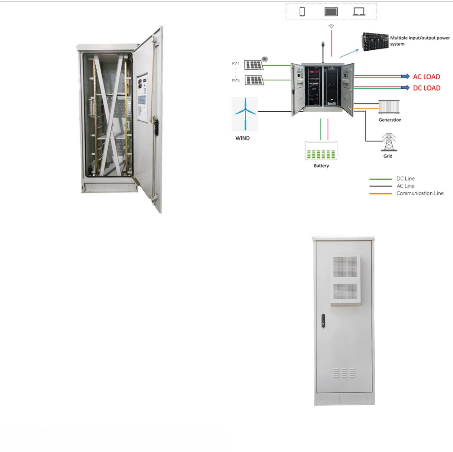

This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered and offers several algorithms and a lightweight co-run performance and power predictive model for computing the performance bounds of the optimal co- schedules and finding appropriate schedules. Expand

As modern embedded systems like cars need high-power integrated CPUs???GPU SoCs for various real-time applications such as lane or pedestrian detection, they face greater thermal problems than be- called CPU???GPU co-scheduling, for CPUs and GPU. It deter-mines which tasks to schedule on CPUs by considering the task running on its

DOI: 10.1109/IPDPS.2017.124 Corpus ID: 23878317; Co-Run Scheduling with Power Cap on Integrated CPU-GPU Systems @article{Zhu2017CoRunSW, title={Co-Run Scheduling with Power Cap on Integrated CPU-GPU Systems}, author={Qingnhua Zhu and Bo Wu and Xipeng Shen and Li Shen and Zhiying Wang}, journal={2017 IEEE International Parallel and Distributed ???

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

Fused CPU-GPU architectures integrate a CPU and general-purpose GPU on a single die. Recent fused architectures even share the last level cache (LLC) between CPU and GPU. This enables hardware-supported byte-level coherency. Thus, CPU and GPU can execute computational kernels collaboratively, but novel methods to co-schedule work are required. ???

Our experimental results across a large set of over 70 mobile games show that Co-Cap improves energy per frame by 10.6% and 10.0% (23.1% and 19.1% in CPU dominant applications) on average and achieves minimal frames per second (FPS) loss by 0.5% and 0.7% (1.3% and 1.7% in CPU dominant applications) on average in training- and deployment sets

This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered. It reveals the performance degradations caused by the co-run contentions at the levels of both memory and power. It then examines the problem of using job co-scheduling to alleviate the degradations in this less understood scenario.

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

Another related study introduces a runtime framework for scheduling each of multiple users'' OpenCL tasks to its optimal device, either a GPU or a CPU on a CPU-GPU system (Wen and O''Boyle 2017). The runtime framework uses a performance prediction model based on machine learning at runtime to select optimal devices.

but they do not guarantee predictable GPU timing on integrated CPU-GPU based platforms because they do not consider the problem of the shared memory bandwidth contention between the CPU and the GPU. Integrated GPU based platforms have recently gained much attention in the real-time systems community. In [2], [12],

Zhu et al. worked on the combination of job scheduling and power capping for integrated CPU-GPU systems [25] Zhu, Q., et al.: Co-run scheduling with power cap on integrated cpu-gpu systems. In: IPDPS. pp. 967???977 (2017) [26] Zhuravlev, S., et al.: Addressing shared resource contention in multicore processors via scheduling. In: ASPLOS. pp

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

Co-Run Scheduling with Power Cap on Integrated CPU-GPU Systems. CPU-GPU??????????? ??<<???????????-???????????????Co??(C)???? ???????,????? 1/4 ????????????Powered by NICT??? Publisher site ???

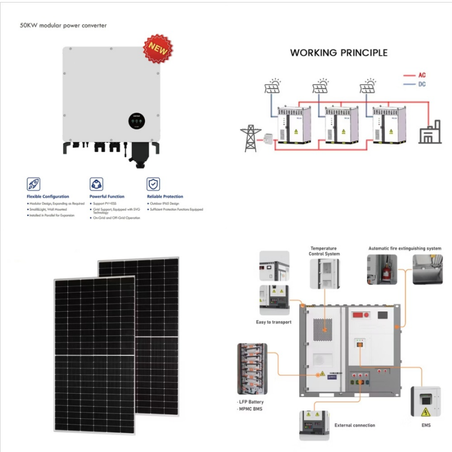

Since not all the programs can benefit from integrated architectures, we build an automatic decision-tree-based model to help application developers predict the co-running performance for a given

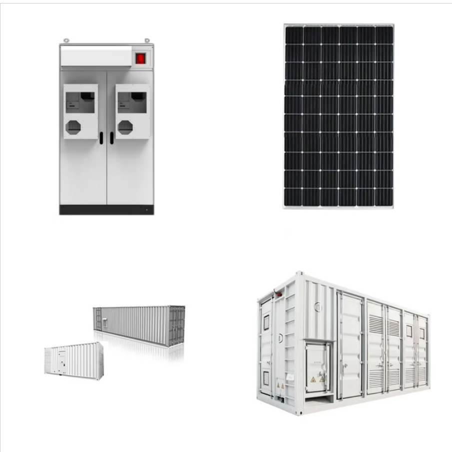

Abstract: This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered. It reveals the performance degradations caused by the co-run contentions at the levels of both memory and power. It then examines the problem of using job co-scheduling to alleviate the degradations in this less ???

CO-RUN SCHEDULING WITH POWER CAP ON INTEGRATED CPU-GPU SYSTEMS

This work investigates the performance implications of independently co-running CPU and GPU programs on these platforms, and produces a list of novel insights, including the important roles of operating system (OS) context switching and power management in determining the program performance. Recent years have witnessed a processor ???

In this section, rst we highlight the thermal and power e ects of scheduling for both CPU-only (e.g., SPEC work-loads [17]) and CPU/GPU applications (e.g. OpenCL ker-nels). Then, we discuss two key challenges and opportunities for e cient scheduling of kernels on CPU+GPU processors. 2.1 Thermal & Power Effects of Scheduling 1.

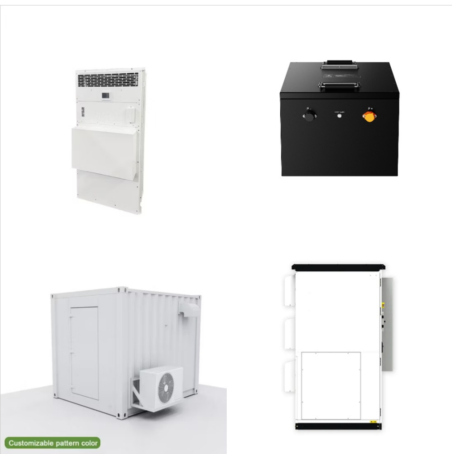

Abstract: This paper presents the first systematic study on co-scheduling independent jobs on integrated CPU-GPU systems with power caps considered. It reveals the performance degradations caused by the co-run contentions at the levels of both memory and power. It then examines the problem of using job co-scheduling to alleviate the degradations in ???