

A Runtime Controller for OpenCL Applications on Heterogeneous System Architectures*

Cristiana Bolchini
DEIB, Politecnico di Milano
cristiana.bolchini@polimi.it

Simone Libutti
DEIB, Politecnico di Milano
simone.libutti@polimi.it

Stefano Cherubin
DEIB, Politecnico di Milano
stefano.cherubin@polimi.it

Antonio Miele
DEIB, Politecnico di Milano
antonio.miele@polimi.it

Gianluca C. Durelli
DEIB, Politecnico di Milano
gianlucacarlo.durelli@polimi.it

Marco D. Santambrogio
DEIB, Politecnico di Milano
marco.santambrogio@polimi.it

ABSTRACT

Heterogeneous architectures nowadays are becoming very attractive in the embedded and mobile markets thanks to the possibility to exploit the best computational resource to optimize the performance per Watt figure of merit. Unfortunately, deciding the right resource to use and its operating frequency is a difficult problem that depends on the actual conditions in which the system is operating. This work aims at proposing a runtime controller, integrated in Linux Operating System (OS), for optimizing the power efficiency of a running application deciding the system configuration. Our experimental results over a set of applications from the Polybench suite on the Odroid XU3 board show that our controller is able to obtain a power efficiency of more than 90% of the one achievable via offline profiling.

1. INTRODUCTION

Heterogeneous System Architectures (HSAs) [7] are becoming nowadays an attractive solution for achieving an optimal trade-off between performance and power/energy consumption thanks to the availability of different kind of resources, such as Central Processing Units (CPUs), possibly integrating heterogeneous cores, Graphic Processing Units (GPUs), Digital Signal Processors (DSPs) and other kinds of HW accelerators. Examples are the Samsung Exynos 5 Octa [15], hosting an ARM big.LITTLE asymmetric octa-core CPU and an ARM Mali GPU, and the Xilinx Zynq [17], integrating an ARM dual-core CPU and a reconfigurable Field Programmable Gate Array (FPGA) unit.

However this increase in heterogeneity comes at the cost of new issues in *programmability* and *runtime management* of these resources to achieve the pursued performance/power consumption trade-off. In particular, various kind of processing units imply different type of programming languages and models thus introducing new implementation and integration challenges. Nevertheless, this abundance of resources has to be properly managed in the execution of the workload, since each type of processing unit offers a different level of performance/power efficiency to each single application and part of it.

In 2009, Khronos Group, including Apple, ARM, Samsung and many other industrial partners, has defined OpenCL [10], a cross-platform programming model designed around the *Single Instruction Multiple Thread (SIMT)* computational

*EWiLi'16, October 6th, 2016, Pittsburgh, USA. Copyright retained by the authors.

paradigm, to exploit data parallelism on heterogeneous accelerators. OpenCL, that has been implemented as an extension of C/C++ languages, enables the programmability and the usage of a large variety of processing units with a single programming model. However, even if enabling functional portability between different processing units, the OpenCL API still requires the programmer to explicitly select and tune the resources to be used for the execution of the application. This still constitutes a limitation since each application may have different optimal operating points on different platforms, and, also on the same platform, the optimal configuration may also vary on the basis of performance requirements expressed by the user or on the overall working conditions of the board (e.g., low-battery mode). Thus, there is a quest in self-adaptation of OpenCL applications to identify in each working scenario the optimal working point.

In this paper, we present a runtime controller integrated within OpenCL applications running on Linux enabling the monitoring of system status and the automated adaptation of the application itself¹. We also propose a novel policy integrated within this controller allowing the application to self-tuning by acting on the mapping and the Dynamic Voltage and Frequency Scaling (DVFS) of the processing units to optimize the performance/power consumption trade-off. Experimental results presented in the paper show the efficiency of the controller to quickly converge to the optimal solution with less than 10% of error.

The rest of the paper is organized as follows. Section 2 briefly discussed the related work. Then Section 3 introduces the working scenario and states the addressed optimization problem. The implementation of the proposed integrated runtime controller and the decision policy are provided in the Sections 4 and 5, respectively. Then, an experimental evaluation of the approach is provided in the subsequent Section 6, and, finally, Section 7 concludes the paper.

2. RELATED WORK

Many OpenCL runtime supports have been defined by vendors for their designed processing units; examples are Intel for last generations of Pentium, Xeon and HD Graphics units, NVidia and ARM for GPU devices, and Xilinx for FPGAs. Moreover, other open source runtime supports, such as [3, 9, 8], have been designed to overcome the unavailability of commercial solutions especially for some type of

¹The source code is publicly available at <https://bitbucket.org/necst/opencl-cgroups-library-release>




Figure 1: Architecture of the Samsung Exynos 5422.

CPUs. Finally, OpenCL ICD loaders (e.g., [16]) have been also implemented to dynamically discover and use at the same time in the same application various runtime supports in computing systems containing devices from different vendors. When considering the mobile and embedded scenario, the main limitation of these OpenCL solutions is the lack of an advanced support to the widely-used ARM big.LITTLE device. In fact, ARM does not provide any runtime for the CPUs [1], while open source solutions handle such a multi-core as a single device and spawn threads indistinctly on all the cores. Thus, the presence of two highly-different clusters that could be used separately is neglected.

Approaches for tuning and optimizing OpenCL applications on HSAs have been recently investigated by a number of works, such as [13, 14, 11]. Their main idea is to perform a design space exploration to identify the best solution by acting on the workgroup tuning and task mapping [13], or workload partitioning among CPU and GPU [14], or by using Domain Specific Languages and source-to-source compilation of the customized OpenCL [11]. Since all these works are based on design-time activities, they require a specific design optimization for each considered architectural platform and, nevertheless, do not feature runtime controllers able to adapt to changes in the working conditions.

Further works (e.g., [2, 12]) have proposed runtime controllers to perform dynamic resource management; their goal is to optimize the trade-off between performance and power/energy consumption by adapting to the currently running workload and related execution requirements specified by the user. Unfortunately, none of such frameworks support OpenCL applications.

3. PROBLEM DEFINITION

This section presents the working scenario considered in this paper, consisting of the target platform and the class of executed applications. Finally, we formulate the optimization problem addressed in the runtime controller we propose.

3.1 Target Architecture and Applications

In this work we consider an HSA as the Samsung Exynos 5422 [15]. As depicted in Figure 1, this chip features an ARM A15 quad-core cluster (called *big*) and an ARM A7 one (called *LITTLE*). The big cluster, suited for high performance, can run at frequencies in the ranges from 200 to 2000 MHz, while the LITTLE one, suited for low power, at frequencies in the 200-1300 MHz range. Moreover, the architecture contains an ARM MALI GPU which frequency can be configured in the 177-600 MHz range. DVFS can be used to change the frequency at runtime with a *per-cluster* gran-

ularity. Finally, the chip is provided with power monitoring sensors. Both sensors and actuators are accessible through standard interfaces provided by the loaded Linux OS. As a final note, the solution is valid for any alternative HSAs with a similar architecture running Linux OS.

Regarding the target applications, a set of computational kernels that is of great interest in the context of embedded and mobile systems is the family of polyhedral applications. Polyhedral applications comprise algorithms for video processing, filters and algebraic transformations which are at the basis of control, infotainment and augmented reality software. These applications are characterized by a computational intensive kernel continuously executed in a loop on incoming data, such as the frames in a video to be processed. In this work we considered the OpenCL implementation of Polybench [4] benchmark suite as a representative set of polyhedral applications.

3.2 Problem Formulation

The addressed problem is to identify for a single given application running in our system which is the best operating point in terms of Performance per Watt. In particular, we want to identify at runtime, without previous profiling information, which is the best processing unit to use and the related frequency level.

In a more formal way, let us consider the controlled application A running on the target architecture containing three different processing units $P = \{BIG, LITTLE, GPU\}$; each unit is characterized by a set of possible frequencies (in MHz):

$$\begin{aligned} f_{BIG} &= \{200, 300, \dots, 1900, 2000\} \\ f_{LITTLE} &= \{200, 300, \dots, 1200, 1300\} \\ f_{GPU} &= \{177, 266, 350, 420, 480, 543, 600\} \end{aligned}$$

The running application will be characterized for each processing unit and supported frequency level, by two direct metrics the throughput, Thr , and the overall power consumption W , and a derived metric called the power efficiency EFF , defined as following:

$$EFF_{p,f} = \frac{Thr_{p,f}}{W_{p,f}}, p \in P, f \in f_p$$

The goal tackled in this work is to find $\hat{p} \in P$ and $\hat{f} \in f_{\hat{p}}$ such that:

$$EFF_{\hat{p},\hat{f}} \geq EFF_{p,f}, \forall p \in P \wedge \forall f \in f_p$$

In order to solve this optimization problem we need to address a set of technical issues related to the monitoring and controllability of the running application on the target system; specifically, we need i) the support for OpenCL for all the processors with the possibility to constrain and move the execution at runtime; ii) to measure the throughput of one iteration of the application and its power consumption; iii) a smart algorithm to explore the power efficiency curves and rapidly identify the best operating point.

4. CONTROLLER IMPLEMENTATION

The self-adaptive approach proposed in this paper has been implemented in a specific controller C++ class directly instantiated within the application source code. Figure 2 depicts the overall structure of the controller and its integration within the system. Moreover, in order to enable the

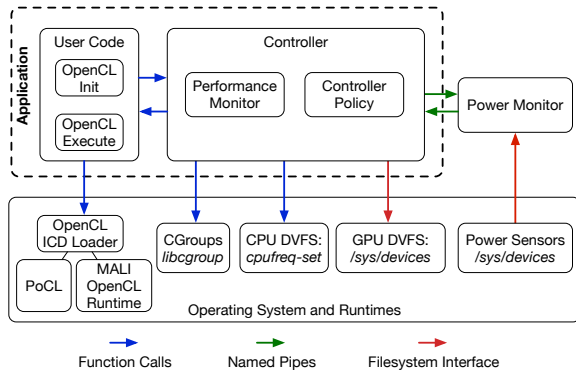


Figure 2: Overview of the implemented system.

actuation of the mapping on all OpenCL devices, the application has to be implemented according to a specific template. Listing 1 shows the defined application template and how the controller is instantiated and used. All the details of the controller in Figure 2 are discussed in the following paragraphs.

OpenCL runtime. To enable the support for all the devices available in the Exynos chip, we have installed both ARM OpenCL Mali SDK [1] and Portable OpenCL library [8], and we have enabled the concurrent discovery of both the platforms with the OpenCL ICD Loader provided by [16].

OpenCL template. The defined application structure is shown in Listing 1; it slightly enhances the standard OpenCL template. In particular, the latter one requires the programmer to select and configure the desired platform and device to be used to execute the kernel. We extend this template to allow the defined controller to dynamically select at each iteration of the application which device to use.

To enable this capability, in the OpenCL initialization step, all platforms and devices are discovered and set up, as shown in the piece of code reporting the `cl_init()` function (Lines 8-19). In particular, the function iterates on all platforms and sets up all devices² and related OpenCL objects, such as the context, the memory objects and the program objects; all these objects are stored in arrays.

Then, a specific variable, `curr_device`, is used to specify the index of the current device to be used in the execution of the kernel. Thus, the `run_kernel()` function uses the execution context specified in such a variable to run the kernel (Lines 21-38).

Cgroups actuation. In order to enable cluster-level mapping on the big.LITTLE CPU, we have exploited OS facilities for task mapping to force the usage of a subset of the cores. Linux OS provides two different mechanisms: `sched_set_affinity()` and `cgroups.sched_set_affinity()` cannot be used in OpenCL applications, since it needs to know thread IDs; indeed, threads are generated within the OpenCL runtime and their IDs are not visible externally. Instead, `cgroups` offers the possibility to assign a set of cores, and, more in general, further resources such as CPU quota and memory amount, by specifying only the application PID; all the threads spawned by that PID are then managed automatically. Therefore, as in [12] we have integrated `cgroups` in the proposed controller.

²For the sake of space, in the listing at Line 14 it is assumed to have a single device per platform.

Listing 1: Application template

```

1 //OpenCL objects
2 cl_platform_id platform_ids [MAX_PLATFORMS];
3 cl_device_id device_ids [MAX_DEVICES];
4 ...
5 cl_uint num_platforms;
6 cl_uint num_devices;
7
8 void cl_init () {
9     int i;
10    //setup OpenCL environment for all devices
11    clGetPlatformIDs(0, NULL, &num_platforms);
12    clGetPlatformIDs(num_platforms, platform_ids,
13                     NULL);
14    for (i=0; i<num_platforms; i++){
15        clGetDeviceIDs(platform_ids[i],
16                      CL_DEVICE_TYPE_ALL, 1, &device_ids[i],
17                      &num_devices);
18        //setup other OpenCL objects for device[i]
19        //i.e. context, queues, memory, programs
20        ...
21    }
22 }
23
24 void run_kernel(int currDevice){
25     //load application data
26     ...
27     //setup the workgroup sizes
28     localWorkSize[0] = ...
29     globalWorkSize[0] = ...
30     //write memory objects
31     clEnqueueWriteBuffer(cmdQueue[currDevice],
32                          mem_obj[currDevice], CL_TRUE, ...);
33     ...
34     // Set the arguments of the kernel
35     clSetKernelArg(clKernel[currDevice], 0,
36                   sizeof(cl_mem), (void *)&mem_obj[
37                     currDevice]);
38     ...
39     //execute kernel
40     clEnqueueNDRangeKernel(cmdQueue[currDevice],
41                            clKernel[currDevice], 2, NULL,
42                            globalWorkSize, localWorkSize, 0, NULL,
43                            NULL);
44     clFinish(cmdQueue[currDevice]);
45     //read memory objects
46     clEnqueueReadBuffer(cmdQueue[currDevice],
47                         mem_obj2[currDevice], CL_TRUE, ...);
48 }
49
50 int main(){
51     int curr_device, i;
52     Controller controller;
53     //applications variables and objects
54     ...
55     cl_init();
56     //setup CGroup, Heartbeat and policy
57     controller.init();
58     //application's main loop
59     for (int i=0; i<iterations; i++){
60         curr_device = controller.get_curr_config();
61         run_kernel(curr_device);
62         controller.send_heartbeat(DATA_SIZE);
63     }
64     //delete all objects
65     controller.destroy();
66     ...
67 }

```

Performance monitor. Instruction per Cycle (IPC) or other classical low-level metrics computed by the OS do not represent a useful information to the final user to perceive the actual progress of an application. As an example, IPC is not able to show if the video application in execution is providing a minimum Quality of Service (QoS) in terms of frame/s. Therefore, to enable run-time performance monitoring, we have integrated in the controller the Heartbeat mechanism [6], a state-of-the-art solution to acquire high-




Figure 3: Typical performance/power figure of the Polybench applications.




Figure 4: Power efficiency of three different Polybench applications.

level information from the application.

The basic idea of the Heartbeat mechanism is to measure the throughput of periodic application by i) measuring the duration of the execution of each single loop and ii) computing the ratio between the amount of processed data and such a duration. Therefore, the controller initialized all necessary data structures (timers and accumulators) during the initialization (Line 47). Then, at the end of the loop (Line 52), the invoked `send_heartbeat()` function collects the new timestamp and the size of the processed data (directly specified by the programmer) and based on such information computes the current throughput.

Power monitor. The considered Exynos chip integrated various sensors to monitor the status of the hardware platform, such as power consumption and temperature of the various clusters. Such sensors are exposed to the programmer through the virtual file systems of Linux OS.

In order to trace the power consumption of the big and LITTLE clusters and of the GPU, we have implemented an external monitor acting as a separate process daemon and periodically (i.e. every 50ms) collecting power values from the interface provided by the `sys` virtual file system of Linux OS. The external monitor can be triggered via a message over a *named pipe* (a Linux interprocess communication mechanism) and instructed to collect power information aggregating them over a period of time. Another message on the same pipe can stop the acquisition returning the average power consumption over the considered period.

Controller. The controller has been implemented in a single class encapsulating all the discussed mechanism and the decision policy. It exposes the following methods:

- `init()`, invoked at Line 47, sets up the environment by initializing the data structures of the Heartbeat, *cgroups* and of the decision policy; moreover the function connects to the external power monitor by means of the named pipe.
- `get_curr_config()`, invoked at Line 50, analyzes all

collected metrics (power and throughput) and executes the decision policy to identify on which device to run the application kernel during the current loop iteration; moreover, it actuates on the *cgroup* library.

- `send_heartbeat()`, invoked at Line 52, processes current Heartbeat at the end of the main loop to compute the throughput.
- `destroy()`, invoked at Line 55, deallocates all data structures.

5. CONTROLLER POLICY

This section describes the controller policy that allows to solve the problem tackled in this paper, defined in Section 3. We will first discuss a preliminary profiling phase carried out on the target applications in order to understand which control strategy should to be adopted, and, then, we will describe the policy itself.

5.1 Preliminary Analysis

In a preliminary phase, we carried out an experimental evaluation of the behavior of such applications on the considered heterogeneous platform. We measured the execution time, power consumption and power efficiency of each considered application on each processing unit (big, LITTLE and GPU) at each available frequency level.

During this analysis, we noticed that the power profile is almost the same for all the considered applications presenting a quadratic relation with respect to the frequency, as shown for GEMM application in Figure 3(a). In the same way the execution time of the applications follows the same trend for all the benchmarks where an increase in the frequency level turns into an improvement of the execution time up until a certain value, as shown in Figure 3(b) for the GEMM application.

Combining these two curves we found how the power efficiency of the benchmarks varies (Figure 4). Also in this case,

the trend of the curves is similar for all the benchmarks. An interesting aspect is that, depending on the actual values of the execution time and power consumption, three different situations can be found where either the big, the LITTLE, or the GPU outperforms the other units (as shown in Figure 4). Finally, all these curves present a maximum which is located around the middle of the frequency range for all architectures. The goal of the policy is then to find at runtime this maximum without any previous profiling information.

5.2 Policy Definition

The controller policy is triggered each time the user code requests a configuration to use (i.e. at the beginning of each iteration of the application loop).

Considering the structure of the power efficiency curves, if we fix a device, we can use a ternary search algorithm to solve our optimization problem. This algorithm can be used to find a maximum of a mathematical continuous function F (with a single maximum point), and, actually, the curves characterizing the computational efficiency of the considered applications (refer to Figure 4) present such a shape.

The algorithm needs to find three initial points $a, b, c \mid a < b < c \wedge F(b) \geq F(a) \wedge F(b) \geq F(c)$; to do this we execute the first three iterations of the application loop on a given device at its minimum, maximum and middle frequencies.

Then, at each iteration we acquire a new estimation of the power efficiency of the application at the last operating frequency for a given device. The new a, b, c points will be chosen by sorting the estimation and picking the two frequencies next to the current maximum efficiency found. At this point one operating frequency is chosen in the ranges (a, b) or (b, c) , priority is given to the first interval; if no frequency exists in those two ranges, it means that the three points a, b, c are contiguous to each others, and b is the actual maximum we are looking for.

The ternary search algorithm has been implemented in a function called `getNextConfiguration()` that is used to pick for a given device the frequency to use. This function is then used in the `getConfiguration()` routine (Listing 2) that decides which device to use to execute the application kernel alongside its frequency. In the `getConfiguration()` function, the controller checks for all the available devices (`devicesList`) whether there exists an estimation for the best efficiency. If such estimation does not exist the controller uses the ternary search on the device to find the maximum as explained before (Line 2). When all the estimations are available, the algorithm selects the device with the best power efficiency (Lines 9-14) and asks that device for the next configuration (Line 16).

The whole configuration decision process is invoked when the application code calls the `get_curr_config()` (Line 50 in Listing 1). The code of this function, which bridges the application code with the controller policy is represented in Listing 3. The API calls the `get_curr_config()` function and then invokes the function to set the desired configuration: only the frequency for the GPU, both frequency and *cggroups* configuration for the CPU. Finally, it returns to the user the device id to use when invoking the OpenCL kernel, as explained in the application template section.

6. EXPERIMENTAL RESULTS

This section illustrates the results of our controller. Experiments have been conducted on an Odroid XU3 board [5]

Listing 2: Selection of next configuration

```

1 Configuration Controller::getConfiguration() {
2     for(auto d : devicesList){
3         if(!d->hasConverged())
4             return d->getNextConfiguration();
5     }
6
7     std::vector<std::pair<float, int>>
8         bestEfficiency;
9
10    for(int i=0; i<devicesList.size(); i++){
11        float xMax, yMax;
12        xMax = devicesList[i]->maxEstimation;
13        yMax = devicesList[i]->
14            getEstimationAtFrequency(xMax);
15        bestEfficiency.push_back( std::pair<float,
16            int>(yMax, i) );
17    }
18
19    std::sort(bestEfficiency.begin(),
20        bestEfficiency.end());
21    std::pair<float, int> bestDevice =
22        bestEfficiency[bestEfficiency.size()-1];
23
24    return devicesList[bestDevice.second]->
25        getNextConfiguration();
26 }

```

Listing 3: Code of the `get_curr_config()` API

```

1 unsigned int OpenClController::get_curr_config
2     () {
3     this->currentConfiguration = getConfiguration
4     ();
5
6     // Set frequency using either cpufreq-set for
7     // CPU or using filesystem for GPU
8     // For CPU devices also calls the proper
9     // cggroups functions
10    this->currentConfiguration.d->
11        setConfiguration(this->
12            currentConfiguration.frequency);
13
14    // Return the device id to use in the
15    // application
16    return this->currentConfiguration.device->
17        openCl_deviceId;
18 }

```

hosting an Exynos 5420 chip. We used a set of 7 applications from the Polyhedral benchmark suite [4], namely 3DCONV, 3MM, ATAX, BICG, GEMM, MVT, SYR2K. All these applications have been extended with the inclusion of the controller proposed in this work.

In order to test the performance of our controller, we modified each application to execute the computational kernel for a fixed number of 50 iterations to see if the controller was able to converge to the most power efficient solution possible. The best configuration for power efficiency has been identified with an offline profiling phase which performed an exhaustive exploration of all the possible configurations.

Figure 5 illustrates the comparison among the profiled power efficiency and the one found at runtime by our controller. Note that the power measures are subject to a little variability due to different working condition (i.e. background system processes, out of our control). As the figure shows, the runtime power efficiency adheres for all the benchmarks to the best power efficiency found during the profiling phase. In particular, for the first five benchmarks (from 3DCONV to GEMM) we have that the GPU has the best power efficiency and the controller is able to converge and use the same device at runtime. In the last two benchmarks we have that MVT has the best efficiency on the BIG

Table 1: Time to converge for the different benchmarks.

Benchmarks	Iterations
ATAX	20
BICG	19
3DCONV	17
GEMM	19
3MM	19
MVT	18
SYR2K	20



Figure 5: Comparison of the solution found by our controller with profiled information.

processors while SYR2K³ is optimal on the LITTLE ones; nonetheless the controller found also in these situations a solution near to the optimum. More in details the controller reaches a power efficiency which is at least 90% of the best value found in profiling. The reason behind the 10% error is due to runtime measurements variability. This causes the policy to converge to a frequency that is near to the optimal one; e.g. GEMM converges to 350 instead of 480 MHz.

In all these experiments the time needed to execute our policy is included in the execution time of the kernel iteration and concurs in defining the power efficiency of the application. Since all the controller functionalities are invoked also when the policy has converged, we can state that the overhead introduced by our solution is negligible.

Concerning the convergence time we have that the controller converges by testing less than 20 different configurations; this represents about 50% of the overall design space (that is composed of 38 configurations) as shows in Table 1. At the opposite, the offline profiling strategy has always to explore all the 38 configurations. Furthermore, we have to keep in mind that the offline profiling requires that the working conditions are exactly the same at the moment the application is in execution. At the opposite, runtime adaptation allows to converge also when the working conditions change.

7. CONCLUSIONS AND FUTURE WORK

This paper has presented a novel runtime controller integrated within OpenCL applications able to enable its automated adaptation. The controller features a novel policy allowing the application to adapt by acting on the mapping and the DVFS of the processing units to optimize the per-

³All the SYR2K power efficiency values have been multiplied by 100 for sake of clarity.

formance/power consumption trade-off. Experimental results have demonstrated the efficiency of the controller to quickly converge to the optimal solution with less than 10% of error. Future work deals with the adoption of further actuation knobs for resource usage, such as quota assignment and finer-grained mapping, the improvement of the proposed policy and controller to support the concurrent execution of several applications.

Acknowledgments

This work has been partially funded by the EU FP7 SAVE project (#610996-SAVE).

8. REFERENCES

- [1] ARM. Exynos 5 Octa. <https://developer.arm.com/products/software/mali-sdks/mali-opencl-sdk>.
- [2] C. Bolchini, G. C. Durelli, A. Miele, G. Pallotta, and M. D. Santambrogio. An orchestrated approach to efficiently manage resources in heterogeneous system architectures. In *Intl. Conf. on Computer Design*, pages 200–207, 2015.
- [3] R. Brochard and N. Nikolaev. FreeOCL. <https://forge.imag.fr/projects/ocl-icd>.
- [4] S. Grauer-Gray, L. Xu, R. Searles, S. Ayalasonmayajula, and J. Cavazos. Auto-tuning a high-level language targeted to GPU codes. In *Proc. of Innovative Parallel Computing*, pages 1–10, 2012.
- [5] Hardkernel co. Odroid XU3. http://www.hardkernel.com/main/products/prdt_info.php?g_code=G140448267127.
- [6] H. Hoffmann, J. Eastep, M. D. Santambrogio, J. E. Miller, and A. Agarwal. Application heartbeats for software performance and health. *ACM Sigplan Notices*, 45(5):347–348, 2010.
- [7] HSA Foundation. <http://www.hsafoundation.com/>, 2015.
- [8] P. Jäskeläinen, C. S. de La Lama, E. Schnetter, K. Raiskila, J. Takala, and H. Berg. pocl: A Performance-Portable OpenCL Implementation. *International Journal of Parallel Programming*, 43(5):752–785, 2015.
- [9] G. Jo, W. J. Jeon, W. Jung, G. Taft, and J. Lee. OpenCL Framework for ARM Processors with NEON Support. In *Proc. of Workshop on Programming Models for SIMD/Vector Processing*, 2014.
- [10] Khronos Group. OpenCL. <https://www.khronos.org/opencl/>, 2016.
- [11] T. Lepley, P. Paulin, and E. Flamand. A novel compilation approach for image processing graphs on a many-core platform with explicitly managed memory. In *Proc. Conf. on Compilers, Architecture and Synthesis for Embedded Systems (CASES)*, pages 1–10, 2013.
- [12] S. Libutti, G. Massari, and W. Fornaciari. Co-scheduling tasks on multi-core heterogeneous systems: An energy-aware perspective. *IET Computers Digital Techniques*, 10(2):77–84, 2016.
- [13] E. Paone, F. Robino, G. Palermo, V. Zaccaria, I. Sander, and C. Silvano. Customization of OpenCL Applications for Efficient Task Mapping Under Heterogeneous Platform Constraints. In *Proc. Conf. on Design, Automation & Test in Europe (DATE)*, pages 736–741, 2015.
- [14] A. Prakash, S. Wang, A. E. Irimiea, and T. Mitra. Energy-efficient execution of data-parallel applications on heterogeneous mobile platforms. In *Proc. Int. Conf. on Computer Design (ICCD)*, pages 208–215, 2015.
- [15] Samsung. Exynos 5 Octa. <http://www.samsung.com/global/business/semiconductor/product/application/detail?productId=7978&iaId=2341>.
- [16] B. Videau and V. Danjean. OpenCL ICD Loader. <https://forge.imag.fr/projects/ocl-icd>.
- [17] Xilinx. Zynq-700 All Programmable SoC. <http://www.xilinx.com/products/silicon-devices/soc/zynq-7000.html>.