

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

On the effectiveness of OpenMP teams for cluster-based many-core accelerators

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

On the effectiveness of OpenMP teams for cluster-based many-core accelerators / Capotondi, Alessandro; Marongiu, Andrea. - ELETTRONICO. - (2016), pp. 7568399.667-7568399.674. (Intervento presentato al convegno 14th International Conference on High Performance Computing and Simulation, HPCS 2016 tenutosi a Innsbruck, Austria nel 18-22 July 2016) [10.1109/HPCSim.2016.7568399].

Availability:

This version is available at: <https://hdl.handle.net/11585/575144> since: 2019-02-18

Published:

DOI: <http://doi.org/10.1109/HPCSim.2016.7568399>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the post peer-review accepted manuscript of:

Capotondi, Alessandro, and Andrea Marongiu. "On the effectiveness of OpenMP teams for cluster-based many-core accelerators." *2016 International Conference on High Performance Computing & Simulation (HPCS)*. IEEE, 2016.

The published version is available online at: <https://ieeexplore.ieee.org/document/7568399>

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works

On the Effectiveness of OpenMP Teams for Cluster-Based Many-Core Accelerators

Alessandro Capotondi
DEI - Università di Bologna
Email: alessandro.capotondi@unibo.it

Andrea Marongiu
IIS - ETH Zurich
Email: a.marongiu@iis.ee.ethz.ch

Abstract—With the introduction of more powerful and massively parallel embedded processors, embedded systems are becoming HPC-capable. Heterogeneous on-chip systems (SoC) that couple a general-purpose *host* processor to a many-core accelerator are becoming more and more widespread, and provide tremendous peak performance/watt, well suited to execute HPC-class programs. The increased computation potential is however traded off for ease programming. Application developers are indeed required to manually deal with outlining code parts suitable for acceleration, parallelize them efficiently over many available cores, and orchestrate data transfers to/from the accelerator. In addition, since most many-cores are organized as a collection of *clusters*, featuring fast local communication but slow remote communication (i.e., to another cluster’s local memory), the programmer should also take care of properly mapping the parallel computation so as to avoid poor data locality. OpenMP v4.0 introduces new constructs for computation offloading, as well as directives to deploy parallel computation in a cluster-aware manner. In this paper we assess the effectiveness of OpenMP v4.0 at exploiting the massive parallelism available in embedded heterogeneous SoCs, comparing to standard parallel loops over several computation-intensive applications from the linear algebra and image processing domains.

I. INTRODUCTION

Nowadays multi- and many-core designs are widely used in most computing domains, from high-performance (HPC) to mobile/embedded systems [1]. Energy efficiency is the key driver for platform evolution, be it for decreasing the energy bills of large data centers [2] or for improving battery life for high-end embedded devices [3]. Architectural heterogeneity has proven an effective design paradigm to successfully tackle many technology walls in the past decade [4] [5]. One of the most common heterogeneous system templates envisions single-chip coupling of a powerful, general-purpose *host* processor to one (or more) *programmable many-core accelerator(s)* (PMCA) featuring tens-to-hundreds of simple and energy efficient processing elements (PE). PMCAs deliver much higher performance/watt, compared to host processors, for a wide range of computation-intensive parallel workloads. To overcome the scalability bottlenecks encountered when interconnecting such a large amount of PEs, several recent embedded many-core accelerators leverage tightly-coupled *clusters* as building blocks. Representative examples include NVIDIA X1 [6], Kalray’s MPPA 256 [7], PEZY-SC [8], ST Microelectronics STHORM [9], or Toshiba 32-core accelerator for multimedia applications [10]. These products leverage a hierarchical design, which groups PUs into small-medium

sized subsystems (*clusters*) with shared L1 memory and high-performance local interconnection. Scalability to larger system sizes employs *cluster* replication and a scalable interconnection medium like a *network-on-chip* (NoC).

The tremendous GOps/Watt figures that such architectures can achieve are traded-off for an increased programming complexity: extensive and time-consuming rewrite of applications is required, using specialized programming paradigms. OpenCL [11], one of the most representative examples of such category of programming models, aims at providing a standardized way of programming accelerators, however it offers a very low-level coding style. Higher-level programming styles are offered by directive-based approaches such as OpenACC [12] or OpenMP [13], which has included in the latest specification document to manage accelerators.

Directive-based programming models have shown their effectiveness in filling the gap between *programmability* and *performance* on heterogeneous many-core architectures. Directives do not alter exiting code written for homogenous CPUs, which enables rapid and maintainable code development thanks to an incremental parallelization style. Several initiatives from academia and from industry follow this path achieving ease of programming at small or no performance loss compared to optimized code written with low level API [14] [15] [16] [17] [18] [19].

Offloaded regions of codes and may include sequential, parallel, and nested parallel regions. Nested parallelism is particularly important to effectively use many cores organized as a fabric of *clusters*. Indeed, while communication to local L1 memory leverages fast and high-bandwidth channels such as crossbars, inter-cluster communication is subject to non-uniform memory access (NUMA) effects, as it relies on multi-hop transactions over a network-on-chip, which offers lower bandwidth and higher latency.

OpenMP 4.0 offers constructs to deploy the parallelism and to distribute the work among clusters in a cluster-aware manner. Specifically, the `teams` construct allows the creation of a team of worker threads each belonging to a different cluster. Each team master can then create nested parallel teams, whose threads are recruited from local resources. In a scratchpad-based architecture, the master thread is typically responsible for bringing data in and out via DMA transfers, thus it is extremely important that the thread-to-core mapping follows a cluster-aware policy such as the one enabled by

the teams construct. Distributing work among threads in a locality aware manner can be done at the loop-level using the `distribute` clause.

In this paper we explore the benefits of using the new OpenMP 4.0 directives for heterogeneous architectures in an embedded cluster-based many-core accelerator, considering several benchmarks from the linear algebra and image processing domain, and showing different programming patterns enabled by OpenMP 4.0. We highlight the benefits of such recent additions to the specifications, comparing the results to a flat parallelization scheme, i.e., one which uses all the processors available from a single logical thread team. In this case only a single master thread for the whole platform is available to orchestrate data transfers, which generates computation with poor locality. We also compare the distributed approach to the use of standard OpenMP nested parallel regions and show that in absence of cluster awareness these perform even worse.

The rest of the paper is organized as follows. Section II describes the target heterogeneous SoC and cluster-based manycores. Section III discusses the key background notions for OpenMP v4.0. Section IV introduces the considered benchmarks and discuss acceleration and parallelization schemes. Section V discusses the evaluation of the described schemes. Section VI describes related work and Section VII concludes the paper.

II. ARCHITECTURAL TEMPLATE

In this work we consider as a many-core based heterogeneous system the ST Microelectronics STHORM platform [9], but the results discussed later can be applied to a broader class of devices which share with STHORM a common architectural template. Figure 1 shows the block diagram of the target heterogeneous embedded system template. A powerful general-purpose processor (the *host*) is coupled to a programmable many-core accelerator composed of several tens of simple processors, where critical computation kernels of an application can be offloaded to improve overall performance/watt [9] [7] [8] [10] [20] [21].

Similar to GPGPUs, the many-core accelerator leverages a multi-cluster design to overcome scalability limitations [9] [7]. Processors within a cluster are tightly-coupled to local L1 scratchpad memory, which implies low-latency and high-bandwidth communication. Globally, the many-core accelerator leverages a partitioned global address space (PGAS). Every remote memory can be directly accessed by each processor, but inter-cluster communication travels through a NoC, and is subject to non-uniform memory access (NUMA) latency and bandwidth. Unlike the typical GPU data-parallel cores, that rely on a common fetch/decode phase, the processors considered here are simple independent RISC cores, perfectly suited to execute both single-instruction, multiple-data (SIMD) and multiple-instruction, multiple-data (MIMD) types of parallelism. This allows to efficiently support a programming model such as OpenMP, that leverages not only data-level parallelism, but also sophisticated forms of dynamic and irregular parallelism (e.g., tasking).

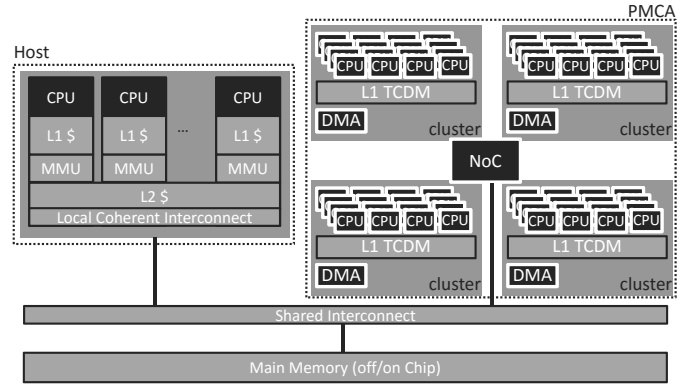


Fig. 1. Heterogeneous embedded SoC template

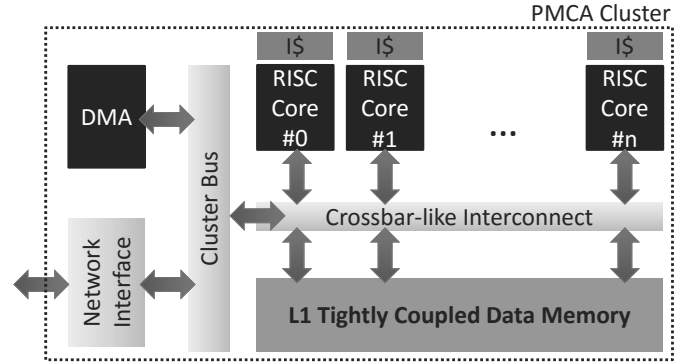


Fig. 2. On-chip shared memory cluster

The simplified block diagram of the target *cluster* is shown in Figure 2. It contains sixteen RISC32 processor cores, each featuring a private instruction cache. Processors communicate through a multi-banked, multi-ported Tightly-Coupled Data Memory (TCDM). This shared L1 TCDM is implemented as explicitly managed SRAM banks (i.e., scratchpad memory), to which processors are interconnected through a low-latency, high-bandwidth data interconnect which allows 2-cycle L1 accesses (one for request, one for response). This is compatible with pipeline depth for load/store for most processors, hence it can be executed in TCDM without stalls – in absence of conflicts. The interconnection supports up to 16 concurrent processor-to-memory transactions within a single clock cycle, given that the target addresses belong to different banks (one port per bank). Multiple concurrent reads at the same address happen in the same clock cycle (broadcast). A real conflict takes place only when multiple processors try to access different addresses within the same bank. In this case the requests are serialized on the single bank port. To minimize the probability of conflicts i) the interconnection implements address interleaving at the word-level; ii) the number of banks is M times the number of cores ($M=2$ by default).

Processors can synchronize by means of standard read/write operations to an area of the TCDM which provides *test-and-set* semantics (a single atomic operation returns the content of the target memory location and updates it).

Mnemonic	Description
STRAS	Matrix multiplication using Strassen decomposition
GSID	Generalized squared interpoint distance
LRFR	Local reference frame radius (surface matching)
HIST	Histogram interpolation
NCC	Normalized cross-correlation algorithm
CT	Object tracking based on a specific color
FAST	Corner detector [22]

TABLE I
BENCHMARK SET.

Since the L1 TCDM is typically very small (256KB for STHORM) it is impossible to permanently host all data therein or to host large data chunks. The software must thus explicitly orchestrate data transfers from main memory to L1, to ensure that the most frequently referenced data at any time are kept close to the processors. To allow for performance- and energy-efficient transfers, the cluster is equipped with a DMA engine.

The OpenMP v4.0 implementation that we consider for our exploration is based on this work [14] and has been extended to include all the features for kernel offloading.

III. BACKGROUND

OpenMP v4.0 [13] introduces *offloading* directives to program accelerators. Similar to any previous OpenMP construct, these directives apply to the code block that they enclose. The key construct is the `target` directive, which highlights the structured code block that should be compiled and loaded for execution onto a device. The `map` clause can be additionally used to specify which data items have to be transferred to and from the device. In addition, the `target data` directive allows to allocate and transfer data before the actual offload takes place (i.e., a sort of data pre-fetching). The `device` clauses allows to specify the exact device to use if more than one is present in the system.

Within a `target` region most standard OpenMP constructs for parallelism can be used. Thus, upon offload a single thread is created that starts execution of the target region, until a `parallel` construct is encountered. Since many accelerators are organized into *clusters*, and since inter-cluster communication is typically costlier than internal transactions, OpenMP v4.0 also introduces directives to abstractly expose architecture organization at the program level. The `teams` directive groups the threads of a device into sets (*teams*) that are later mapped onto physical clusters, thus achieving uniform and high-locality inter-thread communication. The programmer can control the number of teams (`num_teams` clause) and the maximum number of threads in each team (`thread_limit` clause) along with the `teams` directive, respectively. One of the threads in each team is designated team *master* and the structured block following the directive is executed by all team masters across the different teams. Upon team start only team masters execute, sequentially, one per cluster. When a `parallel` directive is encountered, all the threads in each team start execution, to collaborate in the execution of the enclosed structured block.

As most of the parallel work in offloaded kernels is typically found within loops, the `distribute` directive is provided to distribute loop iterations across teams, and then across threads therein. Note that the same thing could not be simply achieved by nesting two `parallel for` constructs, as this would require manually rewriting the loop as a nested loop (with outer and inner loops).

These new constructs allows to achieve a cluster-aware mapping of threads but also loops, without requiring that the programmer explicitly handles these aspects. In the next section we illustrate how these constructs can be used to efficiently offload computation to a many-core accelerator.

IV. BENCHMARKS AND ACCELERATION SCHEMES

In this section we briefly describe the six benchmarks used for our exploration, and the acceleration schemes enabled by the OpenMP v4.0 offload directives. The benchmarks were selected from the linear algebra, image processing and computer vision domains, and are representative of the computational kernels typically offloaded to many-core accelerators. A brief description can be found in Table I.

FAST is particularly sensitive to input data, in terms of the available degree of parallelism. The two parameters that impact the performance the most are input image size and corner density. The former influences the overall number of iteration. Being the core computation kernel of FAST particularly fine-grained, a very small number of iterations per threads results in visible parallelization overheads. The latter influences the actual parallel work, which is protected by an `if` statement that quickly filters out image block that clearly don't contain a corner. For this reason, we consider here six variants of the benchmark execution, with as many different input images. Table II describes the input images used as data set for FAST.

Mnemonic	Size	Description
1.5_S	QVGA	1.5% corner density image
6_S	QVGA	6% corner density image
15_S	QVGA	15% corner density image
1.5_L	VGA	1.5% corner density image
6_L	VGA	6% corner density image
15_L	VGA	15% corner density image

TABLE II
FAST IMAGE INPUT DATASET.

Since in STHORM the host and the accelerator physically share the main L3 memory, the offload infrastructure by default simply passes pointers to data structures therein, rather than copying them to the accelerator space. However, for improved performance and energy efficiency, data has to be moved in the TCDM. In absence of a data cache this has to be explicitly done in the program via DMA transfers.

Each of the considered benchmarks operates on input and output data sets that are too large to fit in the TCDM. Thus, such data structures are divided in *stripes*, which are transferred in and out of the TCDM following a traditional double buffering scheme.

```
#pragma omp target \
    map(in, out)
for( stripe = 0;
    stripe < N_STRIPES;
    stripe++)
{
    dma_in(in[stripe]);
    KER(in[stripe], out[stripe]);
    dma_out(out[stripe]);
}
```

```
void KER(in, out){
    #pragma omp parallel for
    for(i = 0; i < ... ; i++){
        [ALGORITHM]
    }
}
```

flat parallel thread team

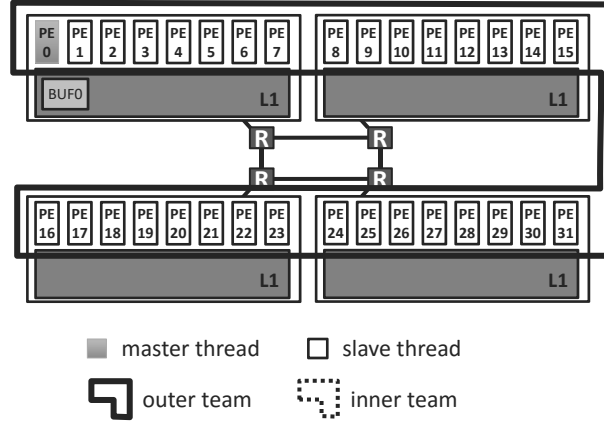


Fig. 3. Flat parallelization pattern.

```
#pragma omp distribute teams \
    map(in, out)\
    num_teams(4)
#pragma omp distribute
for( stripe = 0;
    stripe < N_STRIPES;
    stripe++)
{
    dma_in(in[stripe]);
    KER(in[stripe], out[stripe]);
    dma_out(out[stripe]);
}
```

```
void KER(in, out){
    #pragma omp parallel for
    for(i = 0; i < ... ; i++){
        [ALGORITHM]
    }
}
```

distributed nested team

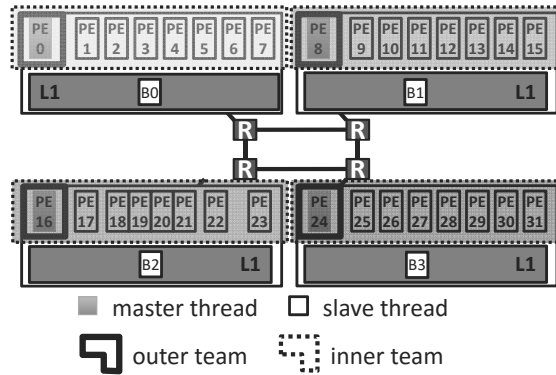


Fig. 4. Distribute parallelization pattern.

DMA transfers of data stripes are typically taken care of by a single thread (the *master*), from within an outer loop. Additional threads are involved in parallel computation when the transfer is complete. To parallelize the target benchmarks we have used three different approaches.

The simplest approach to use all available cores is that of creating a large parallel region which recruits them all. We call this approach *flat* parallelism, as it does not take into account the hierarchical structure of the cluster organization (interconnect, memory). Figure 3 shows how this parallelization scheme deploys threads onto available cores. The *target* directive starts execution of the enclosed region onto a single core,

which orchestrates DMA transfers then jumps into the *KER* function. Here the main loop is found, and it is parallelized with a *parallel for* construct. By default, if no number of threads is specified all the available threads are involved. Note that since the master thread manages the DMA transfers with no awareness of the clusters, the data used by all threads in the parallel region is held in a single buffer (BUF0) that physically reside in the TCDM of the cluster that hosts the master thread. As a consequence, the threads that live in the same cluster as the master will enjoy fast data access, whereas threads belonging to other clusters will experience longer access times, leading to unbalanced computation.

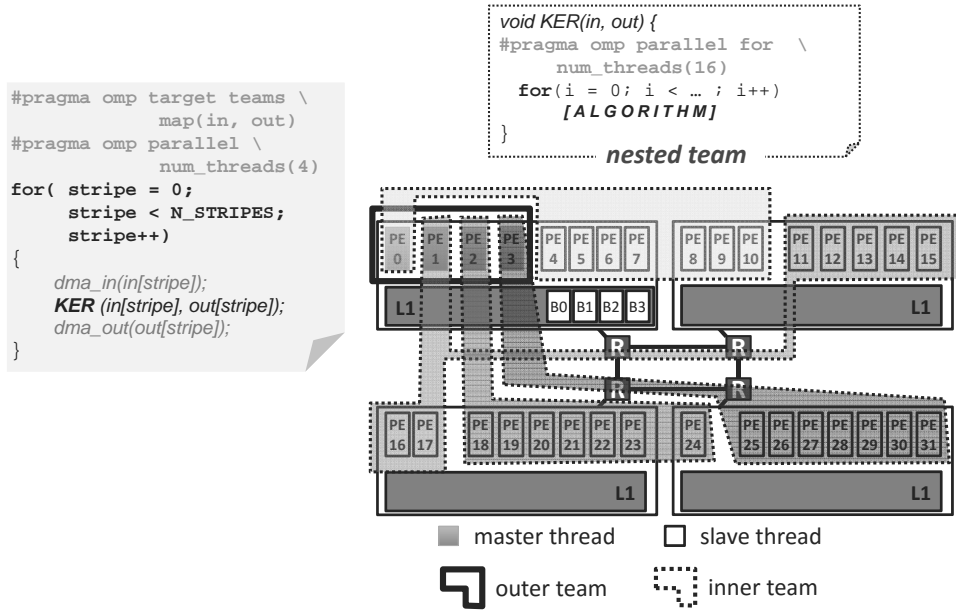


Fig. 5. Nested parallelization pattern.

Figure 4 shows the second parallelization approach, which adds awareness of the clustered nature of the platform to the code. Here the `teams` directive is used to create an outer parallel team that recruits threads from different clusters. These threads will become local masters of these clusters, and will orchestrate DMA transfers to/from the local TCDM. The `distribute` directive is used to partition the outermost loop among local masters, and this will make each master have its own data buffer in the local L1 memory. When a new `parallel` construct is encountered, an inner thread team is created, that shares high locality computation with the local master.

A third parallelization approach is considered for the sake of comparison: standard nested parallel regions. In principle is possible to specify the creation of an outer parallel region with as many threads as clusters, which will act as local master to those regions. Additional parallelism can then be recreated when required by nesting a `parallel` construct within the first. Note that however this scheme lacks a notion of the cluster organization, and threads for the outer and inner regions will be recruited in an unspecified order. In the STHORM implementation this order is sequential, considering the list of all the processors available. Thus, creating an outermost region of four threads recruits the local masters from the same cluster. As a consequence, the code for DMA management will create four data buffers that reside in the same TCDM. Innermost teams will be composed of threads that physically belong to more than one cluster, which will create significantly higher cost for their runtime management (in addition to poor data locality). Figure 5 shows how this approach deploys threads and computation to the platform.

V. EXPERIMENTS

In this section we describe the results collected by running the various benchmarks on STHORM when the three deployment approaches are considered. The experiments rely on an extended version of the multi-ISA toolchain for STHORM proposed by Marongiu et al. [14]. The toolchain supports both OpenMP `teams` and `distribute` directives and nested parallelism. As a main metric of performance we consider speedup of the parallel application versus the sequential.

Results for this experiment are shown in Figure 6.

A. Effectiveness of the `teams distribute` construct

The most notable finding is that the cluster-aware workload deployment enabled by the `distribute` directive allows to achieve very high speedups and thus to make an effective use of many cores. Four out of seven benchmarks achieve nearly ideal speedup (above $60\times$), considering the best result for FAST speedup. As already explained, FAST leverages very fine-grained parallelization, for which the overhead introduced by runtime support for nested parallelism has a higher impact. Thus, when the image size is very small (QVGA) the speedups are limited (up to $60\times$). The corner density is also confirmed to have a big impact on performance, as shown by the variance among the three configurations (1.5%, 6%, 15%). Note that already for moderately large images (VGA) the speedups get as high as close to ideal.

The only application that achieves poor speedup in this configuration is CT, thus it is worth a bit more of investigation. Color-based tracking consists of a cascade of four functional kernels. Color space conversion (CSC), threshold-based color filter (cvTHR), motion vector calculation (cvMOM) and motion vector to reference frame addition (cvADD). Each of these kernels contains little computation, thus to improve the

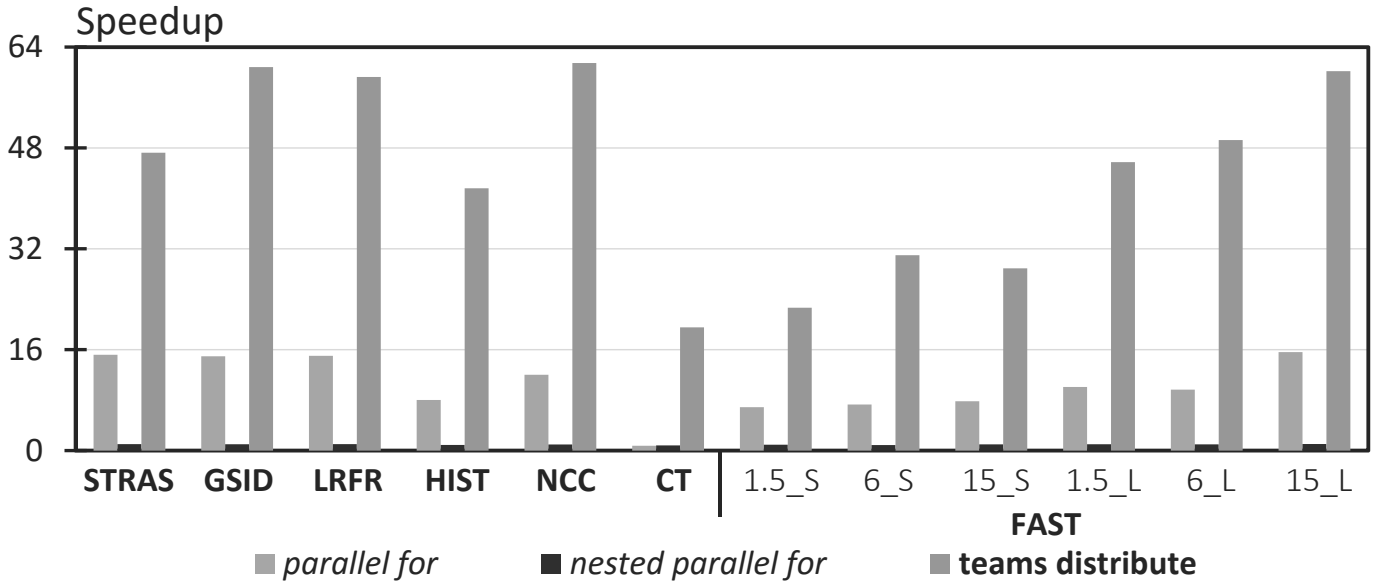


Fig. 6. Comparison of various approaches to nested parallelism support.

computation to communication ratio (CCR) we merge the CSC, cvThresh and cvMOM kernels into a single kernel (i.e., a single data stripe transfer is required to execute all the kernels in sequence). The last kernel, cvADD can not be merged with the previous kernels because it requires as an input the motion vectors for the whole image. Figure 7 illustrates the described parallelization scheme, with the first three kernels merged in a single `teams` region, plus a second `teams` regions composed of the sole last kernel. The figure also shows the breakdown of the speedup for these two `teams` regions. The CCR for cvADD is very small (only an addition is performed per pixel), and this justifies the small speedup achieved for this kernel, which overall impacts the total speedup for the application.

B. Comparison with flat parallel for construct

The comparison with the flat `parallel for` construct shows a much lower efficiency (speedup is always below $16\times$). As explained in Section IV this is due to the poor locality of computation generated by a deployment scheme which only envisions a global master for the entire many-core platform. This master will manage data stripes transfers into the local TCDM, but several threads from the same logical team reside on remote clusters. Such threads will have to traverse the NoC and compete with several other transactions, both for data requests coming from other threads and for instruction cache refills. It has to be pointed out that it is not only the actual parallel computation that encounters such remote communication issues. The implementation of the OpenMP runtime support also relies on data structures that are hosted in the TCDM of the cluster that hosts the master thread. Thus, every time that the parallel code requires explicit or implicit thread synchronization (e.g., barriers, end of parallelization constructs, dynamic loop scheduling, locks, etc.), additional remote transactions are generated. These results are even

more important in the light of the fact that the non-expert programmer will always tend to use the flat `parallel for` approach as a default.

C. Comparison with nested parallel for construct

Probably the most surprising result is that achieved with the nested `parallel for` construct. Due to the above mentioned reasons regarding poor data locality and remote team management it was expected that the speedups would be limited. The extent to which this would impact performance could not entirely be expected. Nested parallel regions have traditionally been used in large HPC systems to improve the performance, however this was achieved: i) on top of multi-level cache hierarchies, that in part mitigate NUMA effects compared to scratchpad-based systems (where every access to a remote data structure can be seen as a miss in a cache-based system); ii) in combination with language or runtime constructs to control thread-to-core binding. Thus, while logically nested parallel regions and distributed teams are equivalent – in terms of how the work is split at the outermost level among *local masters*, and how innermost teams work in strict collaboration with these masters – physically the lack of control of where such masters and their slaves are mapped in the platform leads to extremely poor results. Note that, compared to the flat `parallel for` construct, in this case the impact of runtime library overhead is much more pronounced, as managing and synchronizing nested parallel teams generates much higher communication volumes [23].

VI. RELATED WORK

The latest OpenMP 4.0 specifications introduce relevant features for accelerator exploitation, but not many devices are currently 4.0-enabled. Among commercial devices Texas Instrument Keystone II [24] and Intel Xeon Phi [25] are

```

#pragma omp target teams
num_teams(4)
#pragma omp distribute
for(stripe = 0;
    stripe < N_STRIPES;
    ++stripe)
{
    dma_in(in[stripe]);

    CSC(in[stripe], tmp1[stripe]);
    cvTHR(tmp1[stripe], tmp2[stripe]);
    cvMOM(tmp2[stripe], xy[stripe]);
}

#pragma omp barrier

#pragma omp distribute
for(stripe = 0;
    stripe < N_STRIPES;
    ++stripe)
{
    dma_in(in[stripe]);
    dma_in(track[stripe]);

    cvADD(in[stripe], track[stripe], out[stripe]);

    dma_out(out[stripe]);
}

```

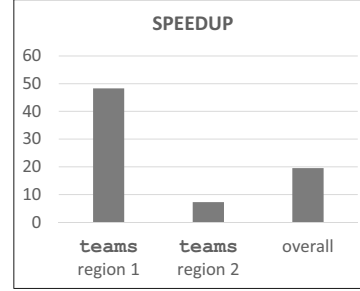
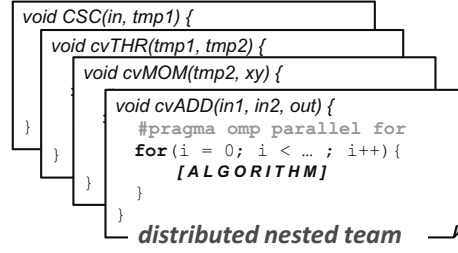


Fig. 7. Breakdown of Color-Tracking kernels speedup.

probably the most representative examples. Stotzer [26] and Schmidl [27] present a performance assessment of flat parallelism for these architecture. These architecture, different from the embedded manycores considered in our work, rely on a coherent shared memory system and on multi-level data-cache hierarchy.

Bertolli et al. [28] propose a method to coordinate GPGPU threads mimicing the OpenMP 4.0 specification for Nvidia CUDA GPGPUs. [28] explores the utilization of the new `team` and `distribute` pragmas to implement efficiently dynamic parallelism on GPGPU accelerators. The focus of this work is however more on presenting a compiler implementation rather than assessing the effectiveness of the language constructs. Also Liao et al. [29] present an OpenMP 4.0 source to source compiler for Nvidia GPU. The compiler is based on the ROSE Compiler Infrastructure [30] and supports the OpenMP 4.0 `team` and `distribute` directives to deploy threads among CUDA cores. A more recent work from Yang et al. [31] presents a directive-based APIs *la* OpenMP that extends the CUDA language to enable dynamic nested parallelism and task level parallelism within a kernel. Ozen et al. [32] evaluate how different parallel programming interfaces, like OpenMP and other patterns for heterogeneous system can influence the deployment and the efficiency of kernel execution on GPGPUs in OmpSs. Unlike what is presented here, the focus for all these works is on GPGPU-like accelerator.

VII. CONCLUSION

Many-core embedded heterogeneous SoCs are getting closer and closer to their HPC counterparts, in terms of computation

capabilities, but efficiently programming them is a cumbersome task. OpenMP has always provided a user-friendly interface to application development, based on compiler directives that abstractly highlight parallelism in a sequential program. The latest OpenMP specifications introduce new constructs for computation offloading, as well as directives to deploy parallel computation with high data locality. This paper explored the capabilities of OpenMP v4.0 at exploiting the massive parallelism available in embedded heterogeneous SoCs. In particular, our experiments demonstrate that the new `teams` and `distribute` constructs allow to abstractly expose the clustered organization of most many-cores, thus achieving very efficient resource usage. Compared to standard parallel loops (the most widely used by inexperienced programmers) with no awareness of the hierarchical interconnect and memory organization, these new construct enable major improvements in terms of speedup. Nested parallel loops, that logically provide a similar abstraction to the `teams` and `distribute` constructs, in absence of architectural awareness surprisingly perform very poorly, in virtually every considered case.

ACKNOWLEDGMENT

This work was supported by EU project FP7 P-SOCRATES (g.a. 611016) and EU H2020 project HERCULES (g.a. 688860).

REFERENCES

- [1] Pete Decher. Embedding HPC: A rocket in your pocket. [Online]. Available: <http://www.embedded.com/design/prototyping-and-development/4230994/A-rocket-in-your-pocket>

- [2] A. Borghesi, C. Conficoni, M. Lombardi, and A. Bartolini, "Ms3: A mediterranean-style job scheduler for supercomputers-do less when it's too hot!" in *High Performance Computing & Simulation (HPCS), 2015 International Conference on*. IEEE, 2015, pp. 88–95.
- [3] A. Bartolini, M. Ruggiero, and L. Benini, "Hvs-dbs: human visual system-aware dynamic luminance backlight scaling for video streaming applications," in *Proceedings of the seventh ACM international conference on Embedded software*. ACM, 2009, pp. 21–28.
- [4] A. Munir, S. Ranka, and A. Gordon-Ross, "High-performance energy-efficient multicore embedded computing," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 23, no. 4, pp. 684–700, 2012.
- [5] J. Diaz, C. Munoz-Caro, and A. Nino, "A survey of parallel programming models and tools in the multi and many-core era," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 23, no. 8, pp. 1369–1386, 2012.
- [6] Nvidia Inc. (2014) Nvidia Tegra X1 - NVIDIA'S New Mobile Superchip. [Online]. Available: <http://international.download.nvidia.com/pdf/tegra/Tegra-X1-whitepaper-v1.0.pdf>
- [7] Kalray S.A., "Kalray MPPA Manycore 256." [Online]. Available: <http://www.kalrayinc.com/kalray/products/#processors>
- [8] PEZY Computing. (2014) PEZY-SC Many Core Processor. [Online]. Available: <http://www.pezy.co.jp/en/products/pezy-sc.html>
- [9] D. Melpignano, L. Benini, E. Flamand, B. Jego, T. Lepley, G. Haugou, F. Clermidy, and D. Dutoit, "Platform 2012, a many-core computing accelerator for embedded SoCs: performance evaluation of visual analytics applications," in *Proceedings of the 49th Annual Design Automation Conference*. ACM, 2012, pp. 1137–1142.
- [10] H. Xu, J. Tanabe, H. Usui, S. Hosoda, T. Sano, K. Yamamoto, T. Kodaka, N. Nonogaki, N. Ozaki, and T. Miyamori, "A low power many-core SoC with two 32-core clusters connected by tree based NoC for multimedia applications," in *VLSI Circuits (VLSIC), 2012 Symposium on*, 2012, pp. 150–151.
- [11] Khronos Group. (2014) The OpenCL Specification. [Online]. Available: <http://www.khronos.org/registry/cl/specs/opencl-2.0.pdf>
- [12] OpenACC, "The OpenACC Application Programming Interface." [Online]. Available: http://www.openacc.org/sites/default/files/OpenACC.2.0a_1.pdf
- [13] OpenMP ARB. (2013) OpenMP 4.0 Application Program Interface. [Online]. Available: <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>
- [14] A. Marongiu, A. Capotondi, G. Tagliavini, and L. Benini, "Simplifying Many-Core-Based Heterogeneous SoC Programming With Offload Directives," *Industrial Informatics, IEEE Transactions on*, vol. 11, no. 4, pp. 957–967, Aug 2015.
- [15] E. Ayguadé, R. M. Badia, P. Bellens, D. Cabrera, A. Duran, R. Ferrer, M. González, F. Igual, D. Jiménez-González, J. Labarta *et al.*, "Extending OpenMP to survive the heterogeneous multi-core era," *International Journal of Parallel Programming*, vol. 38, no. 5-6, pp. 440–459, 2010.
- [16] R. Dolbeau, S. Bihan, and F. Bodin, "HMPP: A hybrid multi-core parallel programming environment," in *Workshop on General Purpose Processing on Graphics Processing Units (GPGPU 2007)*, 2007.
- [17] S. Lee and R. Eigenmann, "OpenMPC: extended OpenMP for efficient programming and tuning on GPUs," *International Journal of Computational Science and Engineering*, vol. 8, no. 1, pp. 4–20, 2013.
- [18] R. Reyes, I. López-Rodríguez, J. J. Fumero, and F. de Sande, "accULL: an OpenACC implementation with CUDA and OpenCL support," in *Euro-Par 2012 Parallel Processing*. Springer, 2012, pp. 871–882.
- [19] B. Chapman, T. Curtis, S. Pophale, S. Poole, J. Kuehn, C. Koelbel, and L. Smith, "Introducing OpenSHMEM: SHMEM for the PGAS community," in *Proceedings of the Fourth Conference on Partitioned Global Address Space Programming Model*. ACM, 2010, p. 2.
- [20] Adapteva, "Epiphany III 16-core Chip Product." [Online]. Available: http://adapteva.com/docs/e16g301_datasheet.pdf
- [21] A. Heinecke, M. Klemm, and H.-J. Bungartz, "From GPGPU to many-core: Nvidia Fermi and Intel Many Integrated Core architecture," *Computing in Science & Engineering*, vol. 14, no. 2, pp. 78–83, 2012.
- [22] E. Rosten, R. Porter, and T. Drummond, "Faster and better: a machine learning approach to corner detection." *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, pp. 105–19, 2010.
- [23] A. Marongiu, P. Burgio, and L. Benini, "Fast and lightweight support for nested parallelism on cluster-based embedded many-cores," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, 2012, pp. 105–110.
- [24] Texas Instruments Inc. KeyStone II System-on-Chip 66AK2Hx. [Online]. Available: <http://www.ti.com/lit/ds/symlink/66ak2h12.pdf>
- [25] C. George, "Knights Corner, Intel's first many integrated core (MIC) architecture product," in *Hot Chips*, 2012.
- [26] E. Stotzer, A. Jayaraj, M. Ali, A. Friedmann, G. Mitra, A. P. Rendell, and I. Lintault, "Openmp on the low-power ti keystone ii arm/dsp system-on-chip," in *OpenMP in the Era of Low Power Devices and Accelerators*. Springer, 2013, pp. 114–127.
- [27] D. Schmidl, T. Cramer, S. Wienke, C. Terboven, and M. S. Müller, "Assessing the performance of OpenMP programs on the Intel Xeon Phi," in *Euro-Par 2013 Parallel Processing*. Springer, 2013, pp. 547–558.
- [28] C. Bertolli, S. F. Antao, A. E. Eichenberger, K. O'Brien, Z. Sura, A. C. Jacob, T. Chen, and O. Sallenave, "Coordinating GPU Threads for OpenMP 4.0 in LLVM," in *Proceedings of the 2014 LLVM Compiler Infrastructure in HPC*, ser. LLVM-HPC '14, 2014, pp. 12–21.
- [29] C. Liao, Y. Yan, B. R. de Supinski, D. J. Quinlan, and B. Chapman, "Early experiences with the openMP accelerator model," in *OpenMP in the Era of Low Power Devices and Accelerators*. Springer, 2013, pp. 84–98.
- [30] Lawrence Livermore National Laboratory, "ROSE Compiler Infrastructure." [Online]. Available: <http://rosecompiler.org/>
- [31] Y. Yang and H. Zhou, "CUDA-NP: Realizing nested thread-level parallelism in GPGPU applications," in *ACM SIGPLAN Notices*, vol. 49, no. 8. ACM, 2014, pp. 93–106.
- [32] G. Ozen, E. Ayguadé, and J. Labarta, "On the roles of the programmer, the compiler and the runtime system when programming accelerators in OpenMP," in *Using and Improving OpenMP for Devices, Tasks, and More*. Springer, 2014, pp. 215–229.