A Survey and Comparative Study of Hard and Soft Real-time Dynamic Resource Allocation Strategies for Multi/Many-core Systems

AMIT KUMAR SINGH, University of York and University of Southampton
PIOTR DZIURZANSKI, University of York and Staffordshire University
HASHAN ROSHANTHA MENDIS, University of York
LEANDRO SOARES INDRUSIAK, University of York

Multi/Many-core systems are envisioned to satisfy the ever increasing performance requirements of complex applications in various domains such as embedded and high performance computing (HPC). Such systems need to cater for increasingly dynamic workloads, requiring efficient dynamic resource allocation strategies in order to satisfy hard or soft real-time constraints. This article provides an extensive survey of hard and soft real-time dynamic resource allocation strategies proposed over the last two decades and highlights the emerging trends for multi/many-core systems. The survey covers a taxonomy of the resource allocation strategies and considers their various optimization objectives, which have been used to provide comprehensive comparison. The strategies employ various principles such as market and biological concepts to perform the optimizations. The trend followed by the resource allocation strategies, open research challenges, and likely emerging research directions have also been provided.

1 INTRODUCTION

A paradigm shift to the adoption of multi/many-core systems can be observed in various domains such as embedded and high performance computing (HPC). The reason behind such adoption lies in the fact that the performance requirements of applications cannot be satisfied by simply increasing the frequency of a single-core processor, which leads to high power and heat dissipation. In multi/many-core systems, chip manufactures are trying to overcome these bottlenecks by

39

integrating multiple cores operating at low frequencies, where the cores can cohesively communicate with each other [89] [22]. These systems provide increased parallelism that motivates us to partition applications into many small tasks and allocate them onto different cores in order to perform parallel executions towards satisfying the increased performance requirements [117].

The many-core processors have been designed by several chip manufactures. Some examples include Intel's Teraflop 80-core processor [183], Tilera's TILE-Gx family 100-core processor [181], AMD's Opteron 16-core processor [4], and Kalray's MPPA 256-core processor [50]. Recently, a joint effort between IBM and UCDavis has revealed KiloCore 1000-core chip [20]. The large number of cores are usually connected by an on-chip interconnection network [15, 18, 190]. These many-core processors are designed to be exploited in various application domains towards realizing different systems, referred to as many-core systems. Additionally, different types of cores have been integrated to exploit their distinct features towards meeting the functional and non-functional requirements [172]. The integration of different types of cores leads to the development of heterogeneous multi/many-core systems that become a formidable computing alternative where applications witness large improvement over their homogeneous (consisting of identical cores) counterpart. Further, the technological advancements will enable integration of higher number of cores in the same chip.

1.1 Resource Allocation for Multi/Many-core Systems

Resource allocation (mapping) process defines assignment and ordering of the tasks and their communications onto resources of multi/many-core system[1] in view of some optimization criteria such as compute performance and energy consumption. The many-core systems (contain relatively large number of cores) usually extend multi-core systems that contain small number of cores and thus resource allocation techniques for these systems can be interchangeably employed if they have architectural similarities [189]. However, to better exploit the many-core resources, the techniques tailored for multi-cores might need some modifications due to different interconnects. For these systems, usually, the applications need to be partitioned (parallelized) into multiple tasks that can be executed concurrently on different cores. Such partitioning is referred to as functional partitioning and can be furnished with the help of state-of-the-art application parallelization tools, e.g., MPSoC Application Programming Studio (MAPS) [34] and MNEMEE project tool-chain [115], and/or manual analysis. This procedure requires detailed application knowledge and involves finding the tasks, adding synchronization and inter-task communication in the tasks, management of the memory hierarchy communication and checking of the parallelized code (tasks) to ensure for correct functionality [122]. In case the multi/many-core system is heterogeneous, i.e. contains different types of cores, a task *binding* process that specifies the core types on them the task can be allocated along with the cost of allocation is required [172]. To compute the allocation cost, the binding process analyses the implementation costs (e.g., performance, power and resource utilization) of each task on different supported core types such as general purpose processor (GPP), digital signal processor (DSP) and coarse grain re-configurable hardware.

An example *resource allocation* along with the application *parallelization* is shown in Fig. 1. The *parallelization* procedure partitions a *sequential application* described in a high level programming language (e.g., C/C++) into various connected tasks. The connections between the tasks reflect the dependencies in the corresponding sequential application. The example partitioned application is shown as *Application Task Graph* that consists of ten tasks ($t_0,t_1,...,t_9$). For the *Application*

---

[1]By multi/many-core system, we mean a system on a multi/many-core chip or several multi/many-core chips connected with each other, where each chip contains a set of connected cores.
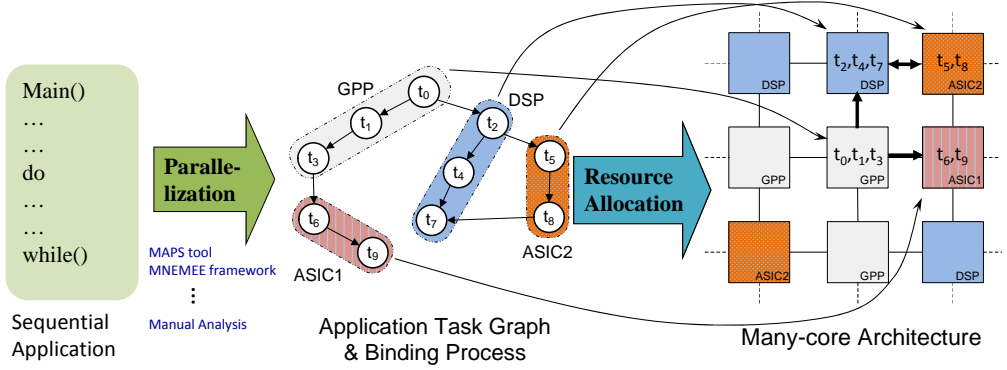
Fig. 1. Resource allocation on many-core system.

Task Graph, the *binding* process has specified different core types for various tasks based on the implementation costs, e.g. GPP for tasks t_0, t_1 and t_3. The *resource allocation* process assigns tasks and their communications on part of a heterogeneous many-core system. The communicating tasks are mapped (allocated) on the same core or close to each other in order to optimize for the communication delay and energy. Differently from functional partitioning, data partitioning can also be employed to perform parallel processing of the data by the same function, e.g., applying a filter function on partitioned segments of an image.

The resource allocation process is carried out either at design-time (statically) or run-time (dynamically). Most of the existing literature for resource allocation fall under static resource allocation (e.g., [33, 39, 75, 86, 118, 129, 136, 180]). However, they cannot handle dynamic workloads and changing environments, e.g. adding a new application into the system at run-time.

Dynamic resource allocation approaches can handle aforementioned issues as the assignment of tasks and their communications on the multi/many-core system resources is done at run-time. In addition, they offer several other advantages such as adaptability to the available resources over time (in case performance requirements of a running application is changed or current allocation is not sufficiently close to optimal) and ability to avoid defective parts of multi/many-core systems and enable foreseeable upgrades [169]. The allocation has been handled either by performing all the processing at the run-time, i.e. on-the-fly processing, or by using previously analyzed results [81, 169]. The results have been analyzed by employing efficient design-time design space exploration (DSE) strategies to encounter for different run-time scenarios [120, 145, 159, 168, 175, 192, 195, 196]. For *on-the-fly* processing, efficient heuristics have been devised to assign new arriving tasks on the system resources [31, 37, 96, 133, 141, 158, 184]. These heuristics do not use any prior analysis results and thus provide rather low quality of resource allocations. However, since they do not use precomputed platform specific analysis results, they cope well to allocate unknown applications on any platform. In contrast, a better quality of resource allocations are achieved by using previously analyzed results, but the applications to be supported on a platform should be known in advance to perform analysis.

Real-time dynamic resource allocation is desired in systems where performance (timing) constraints need to be satisfied to fulfil safe system operations (e.g., in automotive engine management, operating medical equipments and flight control software) and end user demands (frame rate in video processing). This necessitates the development of efficient resource allocation strategies that take an application model, multi/many-core platform model, constraints (e.g., timing and power),

performance model of inter-process communication (e.g., execution time and energy consumption) and estimate of the worst case execution time (WCET) of the process implementations on different cores (e.g., GPP, DSP, ASIC) as input and provide real-time performance guaranteeing resource allocations or optimized resource allocations. A significant amount of research for real-time dynamic resource allocation on single-core system was done in 1980s and 1990s [7]. For multi/many-core systems, it started at the same time, but huge attention was paid after the release of multi/many-core processors by the silicon vendors such as dual-core POWER4 processor by IBM in 2001. Despite the fact that several articles have been published and significant progress has been made for real-time allocation on multi/many-core systems, there still remains many open questions and research challenges.

1.2 Dynamic Resource Allocation Problem and Challenges

It has been well proven that resource allocation is one of the most complex problems in large many-core and distributed systems, and in general it is considered NP-hard [62]. It has also been identified as one of the most urgent problems to be solved for implementing multi/many-core based embedded systems [119] [123]. A well-tuned search algorithm may need to evaluate hundreds of thousands of distinct allocations before it finds one solution that meets the systemâĂŹs performance requirements [120, 145]. Since such evaluation is expected to take a long time, maybe hours to days, it cannot be applied to find the solution quickly, which is desired in the contexts of dynamic resource allocation. Further, it is difficult and challenging to identify the ways that can help to *achieve accurate status of resources* during run-time. This status may be utilization or memory usage of different cores into the system. An inaccurate status of resources may result in an allocation that might not be efficient at run-time.

It is also challenging to satisfy performance requirements of each application when various combinations of simultaneously active applications referred to as use-cases need to be supported into the system at run-time. For each use-case, since optimal solution cannot be explored at run-time due to limited computation power and evaluation time, it needs to be explored by advanced design-time DSE approaches and then to be used at run-time. However, an explosion in the number of use-cases is witnessed with increasing number of applications, e.g., 2^n use-cases for n applications. This makes analysis of all the possible use-cases infeasible for a large number of applications. In order to handle these situations, dynamic resource allocation employing on-the-fly processing needs to be applied even though optimal solutions are not guaranteed.

The resource allocation problem is being addressed by several research groups across the globe, which is evident from the massive available literature in this direction. Applications from various domains (e.g., automotive and video processing) are being targeted to allocate them on multi/many-core based systems in order to exploit their parallel processing capability. For a given domain, the knowledge of applications is used to devise an efficient resource allocation strategy. However, getting application domain knowledge might be time consuming and challenging. The research progress in the direction of dynamic resource allocation is being published in several top ranked conferences/journals to move beyond state-of-the-art.

1.3 Classification of Dynamic Resource Allocation Strategies

The dynamic resource allocation (DRA) strategies can be classified with a number of taxonomies, which could be based on optimization criteria (performance or energy), target architecture (homogeneous or heterogeneous), criticality (hard or soft real-time), etc. Broadly, the classification can be done based on criticality and other taxonomies can be included at some hierarchy in the criticality based classification. For example, hard real-time dynamic resource allocation can target
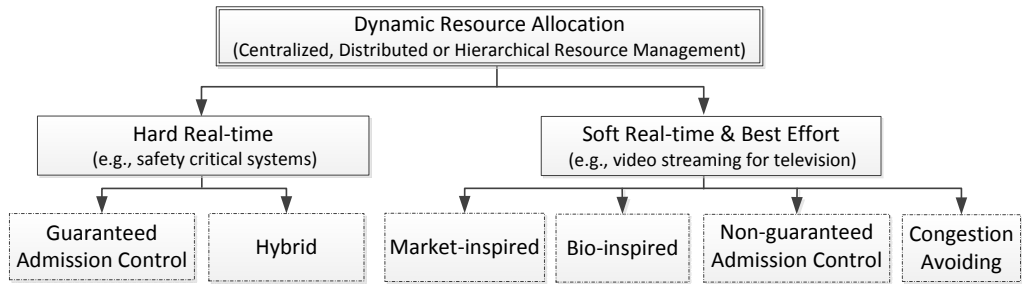
Fig. 2. A taxonomy of dynamic resource allocation strategies.

homogeneous or heterogeneous systems and carry out optimization for performance or energy. Fig. 2 shows a classification of the DRA strategies based on the criticality of the systems. The systems that need to guarantee the timing constraints (e.g., automotive engine management, operating medical equipments and flight control software, realized on a multi/many-core architecture) require *hard real-time* resource allocation approaches, whereas *soft real-time & best effort* resource allocation approaches are desired where deadline miss can be tolerated (e.g, video streaming for television and HPC systems).

For *hard real-time resource allocation*, the existing works reported in the literature can be classified into several categories. However, a careful observation of these works has led them to categorize broadly into *Guaranteed Admission Control* and *Hybrid* approaches, as shown in Fig. 2. In guaranteed admission control, concepts from scheduling theory are used to ensure that only requirement-satisfying applications are entered into the system during the course of resource allocation. The hybrid approaches utilize design-time computed allocations in order to identify a timing constraint satisfying allocation at run-time.

The existing *soft real-time and best effort resource allocation* approaches can be broadly categorized into *Market-inspired*, *Bio-inspired*, *Non-guaranteed Admission Control* and *Congestion Avoiding* strategies after observing the basic employed principles, as shown in Fig. 2. The strategies into different categories utilize some basic principles to optimize for one or several performance metrics in order to fulfil the end user demands.

The strategies under both the above categories perform computation either at both design-time and run-time or only at run-time depending upon the known system status at design-time. For example, if the applications to be executed in the platform are fixed and known at design-time, e.g., workloads of an application specific HPC data center, the allocations can be computed at design-time and used at run-time; otherwise, the allocations needs to be computed at run-time by applying best effort heuristics.

The resource dynamic management process is carried out by employing a *centralized* [31, 139, 141], *distributed* [32, 99, 144], or *hierarchical* [64, 149] approach. In centralized management, one core of the platform is used as the manager that handles the allocation process. For distributed management, the platform is divided into regions (clusters) and one core in each cluster manages the allocation process inside the cluster. The cluster managers communicate with each other through a global manager to find the best cluster for allocating an application. The hierarchical approach exploits the features of both the centralized and distributed approaches.

There are some resource allocation surveys reported in the literature, but they have several limitations, e.g., focus only on hard real-time resource allocation [48], cover strategies based on only one basic principle [194], focus on a specific domain [66, 79] or optimization criteria [69], and not

explicitly and extensively covering hard and soft real-time aspects [111, 146, 153, 169, 200]. Based on the aforementioned taxonomy, this article presents an in-depth *survey* and *comparative study* of dynamic resource allocation strategies, which have been reported in the literature over the last two decades. The strategies that try to specifically optimize *compute performance*, *energy consumption*, or *both of them* have been considered towards focusing on the most important metrics of real-time multi/many-core based systems. They also might optimize for some other performance metrics such as reliability and temperature, but they are not considered to limit the scope of comparison. Optimizing for compute performance is of paramount importance in order to meet the timing deadlines or to minimize the time taken to finish some applications. The compute performance may refer to total *execution time*, *latency*, *delay*, *period*, *throughput*, *exploration time*, worst-case response time (WCRT), etc., which are related to timing information. Optimizing for the energy consumption of modern computing systems, e.g, embedded and high performance centers is important as they are usually operated by stand-alone power supply like battery or a huge amount of energy is required to operate such systems. The energy optimization needs to be performed in order to increase the operational time of the systems and reduce the energy costs.

To include aspects of other taxonomies, e.g., target architecture (homogeneous or heterogeneous) and resource control mechanism (centralized, distributed, or hierarchical), the strategies classified based on criticality (hard or soft real-time, as shown in Figure 2) are analyzed to highlight the considered *optimization goal*, *target architecture type*, *control mechanism*, and their consideration for *only computation or both computation and communication optimization*. The strategies have also been compared and analyzed to highlight their strengths and weaknesses. The above investigations have also enabled to observe the trend followed by the strategies and identify significant open issues and promising future research directions.

Paper Organization: Section 2 and Section 3 cover analysis and elaboration of hard real-time and soft real-time & best effort resource allocation strategies, respectively. A comparative study of strategies falling into different categories has been performed into Section 4. Section 5 provides the upcoming trends that could be followed as the future research and open research challenges. Finally, Section 6 provides some concluding remarks.

2 HARD REAL-TIME RESOURCE ALLOCATION

The majority of works reported in the literature for hard real-time resource allocation assume the workloads to be known in advance, i.e., the allocation decisions by taking the timing constraints are computed at design-time [21, 51, 67, 85, 154]. However, the main focus in this article has been on dynamic resource allocation that needs to compute allocations at run-time for dynamically arriving workloads. Based on the earlier classification as in Figure 2, hard real-time dynamic resource allocation approaches are described next along with the advantages and drawbacks of approaches in each category. Further, a combined discussion of all the approaches is provided.

2.1 Guaranteed Admission Control based Resource Allocation

A guaranteed admission control ensures that all the admitted applications or tasks in a system will meet their respective deadlines without forcing other running applications/tasks to miss theirs. To achieve such a control, an admission controller that fully guarantees the schedulability of the admitted applications or tasks is desired. To ensure the schedulability of the application, schedulability analysis has been extensively used in order to determine whether the application or taskset is schedulable or not when to be allocated on platform core(s) [109]. Various kinds of schedulability tests have been employed in the literature, e.g., response time tests [6, 48, 80] and utilization tests [9, 112]. An application is schedulable if its end-to-end worst-case response time is
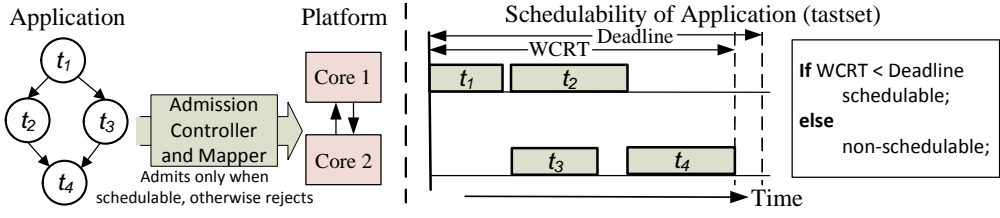
Fig. 3. Guaranteed admission control based resource allocation.

Table 1. Guaranteed admission control based resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimization.

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[105]	Execution time	Comp.	Homogeneous	Centralized
[84]	Execution time	Comp.	Homogeneous	Centralized
[132]	Execution time	Comp. & Comm.	Homogeneous	Centralized
[134]	Execution time	Comp. & Comm.	Heterogeneous	Centralized
[127]	Execution time	Comp. & Comm.	Homogeneous	Centralized
[126]	Execution time	Comp. & Comm.	Homogeneous	Centralized
[56]	Execution time	Comp. & Comm.	Homogeneous	Centralized

less than or equal to its deadline [53]. Similarly, a taskset is schedulable if all its tasks are schedulable. The response time for each task can be estimated by employing widely available standard techniques that takes the interference of the tasks with higher priority into account. Fig. 3 provides an overview of the admission controller based resource allocation and schedulability analysis that determines the end-to-end worst-case response time (WCRT) by taking the interference with running tasks into account in order to identify if the taskset meets the deadline. First tasks are mapped by following an allocation policy. Then, schedulability analysis is performed to determine WCRT. If the WCRT is less than the application deadline, then it is considered to be schedulable and only then the admission controller admits it into the system. Once admitted, it is mapped on the platform cores by following the allocation policy used during the schedulability analysis.

The schedulability analysis is performed by taking the allocations of tasks and scheduling algorithm into account, where the former determines tasks to platform cores assignment and latter defines execution order of tasks assigned to a core. Therefore, the schedulability of a task or application (taskset) depends upon the employed allocation and scheduling algorithm. At run-time, the schedulability of a task considers the entire knowledge on the working scenario at that time, e.g., slack and utilization of cores [56]. The slack of a core is usually defined as the difference between worst-case execution time and actual execution time of the task currently running on the core. In case of a taskset, the scedulability analysis for each task is done when it arrives into the system. One can also perform worst-case schedulability tests which may be conservative, but safe, as a form of online admission test. In these instances, the real resource usage is unknown, but only the worst-case timing properties are used. For hard real-time systems, the schedulability tests ensuring timing guarantees, e.g., worst-case response time test [6, 127] and exact schedulability test [13, 49, 57] can be employed. Feedback-based admission controller originated from control theory can also be used to facilitate hard real-time allocation as more accurate system status is known by feedbacks [56, 113, 199]. Feedback mechanisms monitor the capacity of computing resources and

quality-of-service levels in order to guarantee a bounded time response and stability even if the exact knowledge of a system workload and service capacity is not available a priori [173]. Thus, by careful fine-tuning of control parameters, they can be successfully applied even to systems with real-time constraints.

2.1.1   Existing Works. Table 1 lists works employing guaranteed (hard real-time) admission control based resource allocation of applications on multi/many-core systems. Such dynamic resource allocation works for multi/many-core systems are limited, unlike for single core systems, e.g. [199]. There has also been claims that some costly schedulability tests can be used for small task sets in on-line admission control [16]. These approaches try to find a schedulable allocation at run-time by taking limited available platform resources into account. They use assorted allocation and scheduling algorithms, e.g. highest priority task first and earliest deadline task first, and try to optimize only for the response time that has also been referred to as execution time. While optimizing for the execution time, they also try to optimize other real-time metrics, like schedulable applications, deadline misses, utilization, etc. In these approaches, all the computations are performed at run-time. A few works have been recently reported that utilize design-time computed results for dynamic resource allocation while applying schedulability analysis, but these fall in the category of hybrid resource allocations and are listed in the next subsection.

2.1.2   Advantages and Drawbacks. The schedulability analysis can be used to allocate time critical applications (e.g., avionics and medical) on required systems. Since such analysis can provide WCRT of an application in relatively short amount of time, it can be known whether the application is schedulable or not. In case the application is not schedulable (WCRT > deadline), a different allocation and scheduling algorithm can be tried. Further, if an application is rejected by an admission controller, the resource working time is not wasted with the application that will probably violate its deadline and a possibility of early signalling the lack of admittance can be employed by an outer system to perform an appropriate action (e.g., perform computation on a cloud) minimizing the negative impact of the task rejection. However, the drawback is that users need to have some knowledge of the application domain such that they can analyse the applications and well define the allocation and scheduling algorithms to be applied in order to meet the deadline. The scheduling analysis can also incur large computation (timing) overhead for a complex application/platform containing large number of tasks/cores, and complex dependencies to manage lots of data flows. Therefore, possibilities to reduce the analysis overhead for such complex scenarios can be explored [49, 102]. Another drawback is that the platform resources might be very little utilized as only deadline meeting applications are admitted and their number might be quite lower as hard real-time schedulability tests are pessimistic that take only the worst-case conditions into account.

2.2   Hybrid Resource Allocation

Depending upon the amount of computation involved at design-time and run-time, the existing literature for hybrid approaches can be classified into two categories: *1) Design-time Allocations Computation and Run-time Selection* and *2) Design-time Deadline Distribution and Run-time Allocation*. The details of these class of approaches are provided subsequently.

2.2.1   Design-time Allocations Computation and Run-time Selection. This kind of hybrid approaches utilize design-time computed allocations in order to identify a timing constraint satisfying allocation at run-time. In this approach, the applications to be supported on a platform should be known at design-time in order to perform advance design space exploration (DSE). Fig. 4 provides
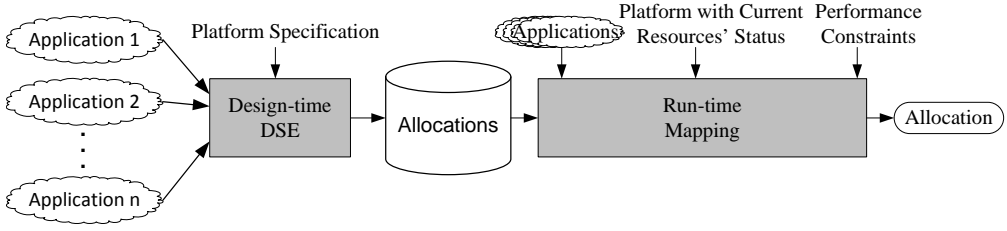
Fig. 4. Hybrid resource allocation: Design-time allocations computation and run-time selection.

Table 2. Comparison of hybrid resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimizing execution time (ET) and/or energy consumption (EC).

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[77]	ET	Comp. & Comm.	Heterogeneous	Centralized
[175]	ET	Comp. & Comm.	Homogeneous	Centralized
[187]	ET	Comp. & Comm.	Heterogeneous	Centralized
[57]	ET	Comp. & Comm.	Homogeneous	Centralized
[42]	EC	Comp. & Comm.	Homogeneous	Centralized
[43]	EC	Comp. & Comm.	Homogeneous	Centralized
[159]	EC	Comp.	Heterogeneous	Centralized
[78]	EC	Comp. & Comm.	Heterogeneous	Centralized
[163]	EC	Comp. & Comm.	Homogeneous	Centralized
[87]	EC	Comp. & Comm.	Homogeneous	Centralized
[55]	EC	Comp. & Comm.	Homogeneous	Centralized
[193]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[195]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[168]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[93]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[148]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[171]	ET & EC	Comp. & Comm.	Homogeneous	Centralized

an overview of this hybrid resource allocation approach, which takes the advantages of both design-time and run-time computations. The heavy computations pertaining to mappings exploration are performed at design-time so that only light ones are left for run-time. At design-time, for each application, DSE is performed by taking the application and architecture specifications as input in order to explore *allocations* with some design objectives. The explored allocations are stored and used as guidelines to efficiently allocate the applications at run-time [103, 168]. The same DSE strategy can be applied to all the applications (Application 1 to Application n) one after another. For the applications to be allocated at run-time, light-weight heuristics are required to select the most efficient allocation for each of them from the storage (precomputed set) of allocations. The selection of an allocation is based on the user demands representing real-time performance constraints of the application and current status (availability) of the platform resources. The selected allocation is used to configure the platform in order to execute the application.

Existing Works. Within the focus of this paper, the performance and energy optimizing existing resource allocation approaches have been considered. Table 2 lists the relevant approaches under this category that optimize performance, energy and both of them. For an application, all the approaches employ DSE to generate multiple allocations (operating points), which are used at run-time based on the platform status and user requirements. There has also been efforts just to develop DSE approaches optimizing for various metrics, such as execution time [5, 91, 145, 188], execution time and energy consumption [63, 196], and resource utilization [175, 192]. These efforts do not explore the ways to use the DSE results at run-time. However, they can be used to generate allocations to be used at run-time.

The DSE takes very long time if the application/platform size is large and hence the exploration may not finish within a limited time. Further, the exploration time increases with the heterogeneity in the platform (different types of cores) and number of optimization goals in order to find efficient allocations. To overcome the exploration time bottleneck, there have been efforts to accelerate the DSE process by incorporating estimations along with time consuming simulations to evaluate the allocations [71, 145, 164]. The accuracy of the results by these approaches depends upon the number of employed simulations. The exploration time could be further reduced by employing pure estimations [98, 131], but the evaluation results will not be accurate.

Along with efficient resource allocation, dynamic voltage and frequency scaling (DVFS) potential of cores has also been exploited to optimize energy consumption [42, 43, 163] . DVFS can be applied both at design-time and run-time based on the available time slack. At design-time, the slack is defined as the difference between application͆s timing constraint and achieved execution time. At run-time, the slack is dynamically created due to tasks finishing earlier than their worst-case execution times (WCETs).

Advantages and Drawbacks. At run-time, since only selection of the allocation from the storage is required, a light-weight run-time platform manager can be employed to configure the applications efficiently. The hybrid approach allocates applications more efficiently than on-the-fly heuristics that perform all the computations at run-time. However, flexibility in these approaches is limited, since all potential applications must be known in entirety at design-time and analysis results will be applicable only to the analyzed platform. Therefore, design-time analysis needs to be repeated when the application set or platform changes. Further, storing analysis results introduces additional memory overhead.

2.2.2   *Design-time Deadline Distribution and Run-time Allocation.* In order to meet a hard real-time (HRT) deadline, some researchers have tried to distribute deadlines to each task of the application at design-time and the resource management approach aims to meet these deadlines at run-time during the resource allocation and application execution process. There has also been efforts to distribute task's deadline further to its functional blocks [1]. Fig. 5 provides an overview of the design-time deadlines assignment to different tasks of an application and run-time resource allocation to meet the deadlines. Based on the application deadline (D_{App}), the computation complexity of its tasks are evaluated and then deadline is assigned to each task based on its computation complexity. At run-time, the tasks are allocated to platform cores to perform execution. If individual tasks meet their deadlines, i.e. WCRT of each task is less than their respective deadline (D), the application meets its deadline.

Existing Works. Table 3 lists deadline distribution based resource allocation approaches reported in the literature. We would like to highlight that even though some of the reviewed techniques were not originally developed targeting multi/many-core systems, they are based on similar assumptions
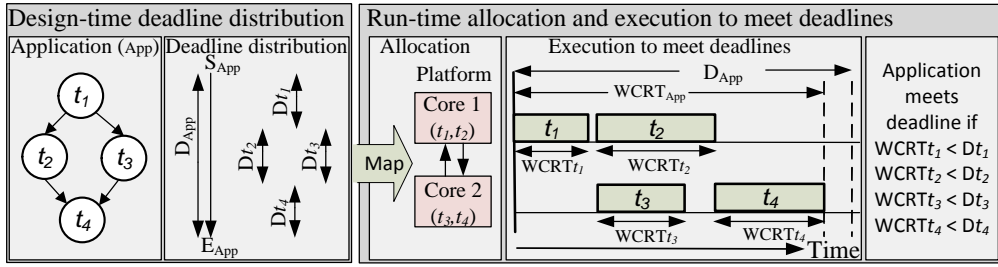
Fig. 5. Design-time deadline distribution and run-time resource management.

Table 3. Deadline distribution based resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimization.

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[155]	Execution time	Comp.	Homogeneous	Centralized
[92]	Execution time	Comp.	Homogeneous	Centralized
[88]	Execution time	Comp.	Homogeneous	Centralized
[160]	Execution time	Comp.	Homogeneous	Centralized
[27]	Execution time	Comp.	Homogeneous	Centralized
[108]	Execution time	Comp.	Homogeneous	Centralized
[73]	Execution time	Comp.	Homogeneous	Centralized

and can or have been applied to the multi/many-core allocation problem. Similar to guaranteed-admission control based resource allocation approaches, these approaches also target to optimize only the execution time (response time) as their main focus is to meet deadlines. The deadline distribution considers mainly two types of deadlines, global and local, which define deadline for a taskset and individual tasks, respectively. With deadlines available for each task, the listed approaches apply various kinds of allocation and scheduling algorithms, e.g., earliest deadline first.

Advantages and Drawbacks. Deadline distribution enables to consider the allocation of individual application tasks independently by considering their deadlines and scheduling order. This facilitates optimization for each individual task in order to meet its local deadline towards meeting the global deadline of the application or taskset. Thus, only local deadlines need to be taken into account. However, this might lead to a situation where all the local deadlines are not met, but the end-to-end deadline might be met. A few drawbacks might be observed when the exact task properties are known, e.g., subtask deadline calculation might be wrong if exact execution time of task is not known and the calculation might be too pessimistic if only WCET is known. Further, in case of a complex application containing a large number of tasks, the deadline distribution to individual tasks considering dependencies amongst the tasks might be quite cumbersome. Additionally, the complexity of the run-time resource management to meet the deadlines also increases.

2.3 Discussions and Summary

These hard real-time approaches have been extensively studied in the literature and applied to meet the deadlines for time-critical applications. The concepts from one category of approaches can be used in another one as well in order to accomplish the aim in a particular category. For example,

deadline distribution and schedulability analysis have been jointly exploited in [150], deadline distribution and admission control in [121], and hybrid and schedulability analysis is employed in [57]. For all the categories, it can also be deduced that complexity of the resource allocation increases with the complexity of the considered application. Out of the above categories, since hybrid resource allocation needs to perform computations at design-time, the taskset needs to be known in advance. However, most of the guaranteed admission control based approaches perform allocation and schedulability analysis directly at run-time as the analysis can be performed in a short amount of time [28].

3 SOFT REAL-TIME AND BEST EFFORT RESOURCE ALLOCATION

Based on the classification as shown in Figure 2, soft real-time and best effort resource allocation approaches are described subsequently. Advantages and drawbacks of approaches in each category are also described. Further, a combined discussion of all the approaches is also provided.

3.1 Market-inspired Resource Allocation

Market-inspired resource allocation mechanisms use market concepts to perform the allocation. These mechanisms use available platform capacity measured by low-level heuristics as bids within an auction-like allocation process in order to find the allocation that can provide guarantees to satisfy the required level of quality of service (QoS) and can maximize the overall system utility (profit). Fig. 6 demonstrates the process of market-inspired dynamic resource allocation where different applications need to be allocated into a many-core system representing a typical HPC data center. The shown system executes a set of *applications* submitted by various *users* at different moments of time. The applications are submitted to the *platform resource manager* that allocates resources to them. To incorporate market and value concepts in the allocation process, applications are assigned values and bids from resources are placed to the allocation engine (Manager Processor) in order to maximize the value/profit returned by a many-core system. The values of applications represent their importance level.

3.1.1 Existing Works. Table 4 lists some market-inspired resource allocation approaches that optimize for execution time, energy consumption, or both of them. The main goal of the execution time optimizing approaches is to maximize value (profit) by early completion of applications. Some researchers assume a fixed value of an application [179], whereas others consider values that can change with time, described with so-called value curve of the application [26, 36, 83, 97]. The changing value over time reflects the impact of the computation over the business processes and adds complexity to the allocation process. With changing value over time, an early completion leads to high value, whereas late completion results in a low value. An enormous literature exists for optimizing value as the main focus of cloud data centers have been to maximize the profit. However, since energy consumption of data centers is quite huge (around 1.5% of the worldwide electricity consumption [100]) and rapidly increasing, optimizing for energy consumption along with the value is of paramount importance and recently has been a focus point of various researchers. Along with allocation, DVFS potential has been exploited to achieve energy savings [29, 167].

3.1.2 Advantages and Drawbacks. These resource allocation approaches are proven to provide promising results in the overload situation where demand for available resources is higher than the supply. Such situation is normally encountered in HPC data centers [194]. The notion of values of applications and auction process facilitate in deciding to hold the low value applications for late allocation and allocating limited resources to the high value applications. However, with value curve associated with each application, the resource allocation becomes complex while considering
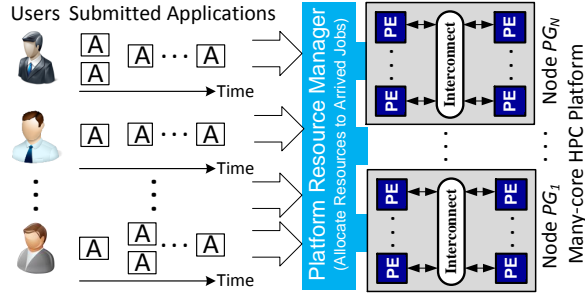
Fig. 6. Market-inspired resource allocation in a cloud data center that contains different nodes (servers) with dedicated cores (or Processing Elements, PEs) to execute applications submitted by multiple users.

Table 4. Market-inspired resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimizing execution time (ET) and/or energy consumption (EC).

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[179]	ET (Value)	Comp.	Heterogeneous	Centralized
[11]	ET (Value)	Comp.	Homogeneous	Centralized
[25]	ET (Value)	Comp. & Comm.	Homogeneous	Centralized
[165]	ET (Value)	Comp. & Comm.	Homogeneous	Centralized
[2]	EC	Comp.	Homogeneous	Centralized
[29]	EC	Comp.	Homogeneous	Centralized
[186]	EC	Comp.	Homogeneous	Centralized
[97]	Value & EC	Comp.	Homogeneous	Centralized
[166]	Value & EC	Comp. & Comm.	Homogeneous	Centralized
[167]	Value & EC	Comp. & Comm.	Homogeneous	Centralized

enormous number of applications arrived at the same time due to overload situation. Depending upon the arrival rate and value curve pattern of applications, appropriate resource allocation approach needs to be identified in order to perform optimization for desired metrics.

3.2 Bio-inspired Resource Allocation

Biologically inspired resource allocation approaches draw inspiration from nature and apply the observed characteristics when solving specific computational problems [68]. They are often based on the characteristics of self-organizing biological systems where global patterns emerge from interactions at a lower-level in the system [8, 30]. This indicates that the resource allocations are changed, i.e. reallocations are performed based on observed biological phenomenon. They have been well explored to balance communication loads in networks and distributed systems [46, 140], communication loads in many-core systems [151], and both computation and communication loads in embedded systems [128]. These approaches usually employ distributed resource management to overcome the limitations of centralized and clustered (hierarchical distributed) managements for dynamic applications and large scale many-core systems. Fig. 7 illustrates an example to reallocate some late tasks in a 3×3 many-core system by employing biological inspiration that is based on pheromone signalling mechanism as seen in social insects (e.g honey bees) in order to improve the overall performance [128]. At each reallocation interval, *late tasks identification* is done on each node
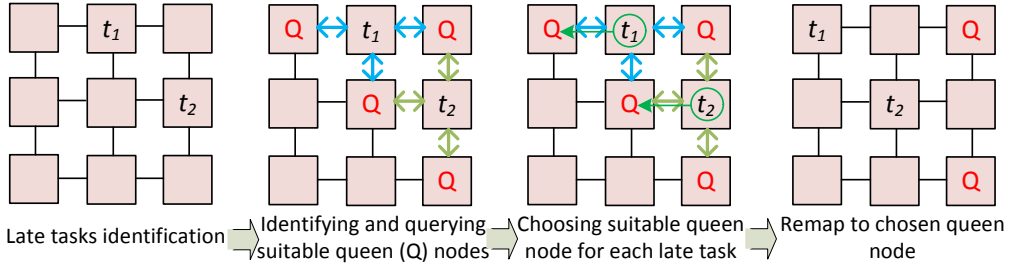
Fig. 7. Bio-inspired resource allocation. An example showing remapping of late tasks t_1 and t_2 on suitable queen nodes (cores) having higher resource availability.

Table 5. Bio-inspired resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimizing execution time (ET) and/or energy consumption (EC).

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[68]	ET	Comp.	Homogeneous	Distributed
[24]	ET	Comp.	Heterogeneous	Distributed
[135]	ET	Comp.	Homogeneous	Distributed
[138]	ET	Comp.	Heterogeneous	Distributed
[17]	ET	Comp.	Homogeneous	Distributed
[12]	EC	Comp.	Homogeneous	Distributed
[90]	ET & EC	Comp. & Comm.	Homogeneous	Distributed
[128]	ET & EC	Comp. & Comm.	Homogeneous	Distributed

to find late tasks in their task queues. Then, for late tasks, e.g. t_1 and t_2, the process of *identifying and querying suitable queen (Q) nodes* is carried out, where each Q node represents its pheromone level in terms of processing capability that is obtained periodically by executing light-weight set of rules on each node. Only the nodes having pheromone level greater than a threshold contribute to the Q nodes and suitable Q nodes are the ones in close proximity. Thereafter, suitability of each Q node is determined to reallocate the late tasks in order to *choose the most suitable queen node for each late task*. The neighbouring Q nodes are chosen to reallocate (remap) the late tasks such that they are still allocated close to each other in order to maintain low communication overhead in case the tasks communicate with each other.

3.2.1 Existing Works. Table 5 lists bio-inspired resource allocation approaches for many-core systems. It has been observed that they have been lesser explored for many-core systems than networks and distributed systems [46, 140]. These approaches have performed optimizations by considering several biological phenomenon, e.g., particle swarm optimization (PSO) [90], pheromone signalling (PS) [128], etc. In [128], energy consumption is not directly optimized, but communication overhead is reduced that leads to optimized communication energy consumption. These approaches have been studied to implement a fully distributed resource allocation in order to overcome the limitations of centralized and clustered (hierarchical) resource managements and have shown promising results in several scenarios.

3.2.2 Advantages and Drawbacks. The bio-inspired approaches perform allocation in a completely distributed and self-organizing way. Thus, they alleviate the limitations of centralized and

hierarchical (mixture of centralized and distributed) resource allocation approaches that suffer from the issues of scalability, huge monitoring traffic, hot-spots, etc. The main drawbacks of these approaches lie in the fact that they are difficult to implement, study and calibrate in real systems due to several assumptions on the running system. Additionally, in static environments with a small number of cores, they might lead to bad results as compared to other design-time (static) optimization approaches. Further, for a system of small size, they might perform worse than the centralized resource management approaches that have a better view of the system resources. For various system sizes, in [99], it has been shown that applications achieve better performance by centralized approach than the distributed approach.

3.3 Non-guaranteed Admission Control based Resource Allocation

A non-guaranteed admission controller does not guarantee the schedulability of the admitted applications or tasks. Therefore, some of the admitted applications might not meet their deadlines. For guaranteed admission control, computation costly schedulability analysis (e.g., exact schedulability test) is employed, which takes several measures into account, e.g., computation/communication requirements of tasks and interference amongst tasks, such that timing guarantees are always fulfilled for the admitted and running tasks. The costly analysis might incur long delay between the application release (arrival) and its allocation process. Further, there might be very low utilization of system resources as the tasks are admitted only when the system can allocate a necessary resource budget to meet timing requirements. To overcome the issues of costly schedulability analysis and low utilization, low cost tests can be employed, which would let a higher number of tasks to enter into the system to increase the resources utilization, but some of them might miss their deadlines. Fig. 8 shows such a process, representing non-guaranteed admission control based resource allocation, where some of the admitted applications might miss their deadlines as light-weight schedulability tests are employed that are usually based on status of the platform resources and thus does not ensure schedulability. Feedback-based admission controller can also be used to provide monitored system information in order to facilitate for better admission control decisions.

3.3.1 Existing Works. Table 6 lists non-guaranteed admission control based resource allocation approaches and works applying light-weight schedulability tests for soft real-time systems. Some of the approaches also employ feedback concepts originated from control-theory. For overall execution time optimization, some approaches admit and allocate the tasks such that a load balance is achieved in the multi/many-core system in order to achieve high performance. The energy optimization is mainly based on the feedback utilized to perform dynamic voltage and frequency scaling (DVFS). Monitoring (feedback) information has also been used to perform adaptive resource allocation in order to improve performance [76, 107]. It has been observed that several works regarding feedback admission control exist for single-core systems, but we focus on multi/many-core systems. Additionally, admission control based approaches to jointly optimize both execution time and energy consumption are lacking.

3.3.2 Advantages and Drawbacks. The non-guaranteed admission control enables high system utilization by admitting a higher number of applications. Additionally, light-weight schedulability analysis reduces the delay between the start of schedulability test and time of admission into the system. This might lead to lower response time. However, in case of using feedback that is some monitored activity in the system, fast and accurate run-time monitoring is desired, which might not always be easily achievable. Further, monitoring activities can flood the network and increase energy consumption.
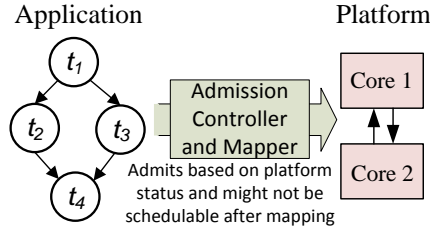
Fig. 8. Non-guaranteed admission control based resource allocation.

Table 6. Non-guaranteed admission control based resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimization.

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[104]	Execution time	Comp.	Homogeneous	Centralized
[101]	Execution time	Comp. & Comm.	Homogeneous	Centralized
[127]	Execution time	Comp. & Comm.	Homogeneous	Centralized
[110]	Energy consumption	Comp.	Homogeneous	Centralized

3.4 Congestion Avoiding Resource Allocation

The congestion avoiding heuristics assign new arriving tasks or application tasks on many-core system resources such that the congestion in the cores and links is minimized. This might lead to optimized overall execution time and energy consumption. These heuristics does not guarantee for a QoS requirement, but can be used to meet soft real-time requirements as they optimize congestion leading to optimized execution time. Fig. 9 shows an example mapping of the computations (tasks) and communications (edges) of an arrived application on a many-core system's cores and links, respectively, while trying to reduce congestion in the cores/links. Tasks t_1 and t_2 are mapped on one core assuming that they are highly communicating tasks (huge data needs to be transferred between them) and can fit on the core, i.e. the core has sufficient memory to allocate them. This reduces the communication overhead between the tasks as data transfer happens via shared local memory on the core. The other tasks are mapped in the close proximity assuming that higher amount of data needs to be transferred between t_3 and t_4 as compared to t_2 and t_4, i.e., t_3 and t_4 are mapped close to each other while trying to reduce the communication distance between t_2 and t_4 as well. The congestion on the cores is also reduced by homogeneously distributing loads of tasks on the cores.

3.4.1 Existing Works. Table 7 lists recent works applying congestion avoiding heuristics to allocate tasks/edges on the resources at run-time. The heuristics employ various fundamental principles based on the characteristics of the application and current status of system resources, e.g., map highly communicating tasks on the same core or neighbouring cores, edges on least utilized (highly available) links, and tasks on least utilized cores. These works optimize for one or several performance metrics. Further, these heuristics perform all the computations at run-time and thus can handle highly dynamic workloads.

3.4.2 Advantages and Drawbacks. The congestion avoiding heuristics cope well to map unknown applications (not available at design-time) on any platform as they do not use any platform specific analysis results computed in advance. However, these heuristics may not be able to guarantee
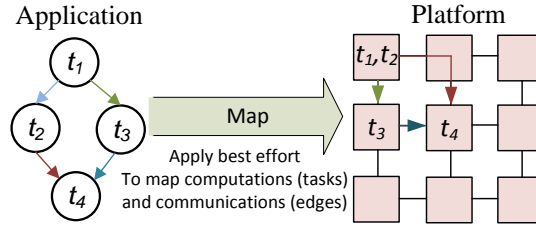
Fig. 9. Congestion avoiding resource allocation in a many-core system for an arrived application.

Table 7. Congestion avoiding resource allocation approaches that consider computation (Comp.) and/or communication (Comm.) for optimizing execution time (ET) and/or energy consumption (EC).

References	Optimization Goal	Comp. & Comm. Consideration	Architecture	Control
[133]	ET	Comp. & Comm.	Homogeneous	Centralized
[141]	ET	Comp. & Comm.	Heterogeneous	Centralized
[3]	ET	Comp. & Comm.	Heterogeneous	Distributed
[74]	ET	Comp.	Homogeneous	Centralized
[162]	ET	Comp.	Homogeneous	Centralized
[144]	ET	Comp. & Comm.	Homogeneous	Distributed
[178]	ET	Comp.	Heterogeneous	Centralized
[184]	ET	Comp. & Comm.	Heterogeneous	Centralized
[19]	ET	Comp.	Heterogeneous	Centralized
[76]	ET	Comp. & Comm.	Heterogeneous	Centralized
[99]	ET	Comp. & Comm.	Heterogeneous	Distributed
[37]	ET	Comp.	Heterogeneous	Centralized
[40]	EC	Comp. & Comm.	Homogeneous	Centralized
[116]	EC	Comp. & Comm.	Homogeneous	Centralized
[176]	EC	Comp. & Comm.	Homogeneous	Centralized
[143]	EC	Comp. & Comm.	Homogeneous	Centralized
[172]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[125]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[23]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[41]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[158]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[31]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[177]	ET & EC	Comp. & Comm.	Heterogeneous	Centralized
[32]	ET & EC	Comp. & Comm.	Homogeneous	Distributed
[130]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[59]	ET & EC	Comp. & Comm.	Homogeneous	Centralized
[139]	ET & EC	Comp. & Comm.	Homogeneous	Centralized

for schedulability, i.e., for strict timing deadlines due to lack of any prior analysis and limited computational power at run-time.

3.5   Discussions and Summary

The soft real-time and best effort resource allocation approaches can be applied to perform optimization for one or several performance metrics depending upon the need of the application domain. For example, in battery operated embedded systems, the timing constraints need to be fulfilled while optimizing for the energy consumption. These approaches employ different kinds of resource management, e.g., centralized, distributed or mix of both (hierarchical) depending upon the size of the many-core system and workload to be executed on it. It has been observed that distributed management leads to better results than centralized one when many-core systems of large sizes are considered [99]. However, in a relatively smaller system, the centralized management might perform better as the distributed approach incurs additional communication overhead amongst the several agents without offering significant advantages [94, 99]. In case the performance needs to be improved at run-time by employing adaptive resource allocation, task migrations are performed, where the tasks are migrated without completely stopping and restarting on the destination core [23, 144]. These approaches can also be modified to perform optimizations for other performance metrics, e.g., reliability, fault-tolerance, temperature, security, etc. by taking appropriate measures into account.

4   COMPARATIVE STUDY AND SUMMARY

This section shows comparative results of various resource allocation approaches falling under hard real-time and soft real-time & best effort categories. We show some example comparisons, where mainly considered approaches are those that employ the same application and multi/many-core system model, or the application models that can be easily converted to a unified model such that the same evaluation tool can be used. It should be noted that a single tool chain is not used to evaluate approaches from various categories, but the tool chain to evaluate a set of approaches within a category is the same and taken from in-house/open-source tool chains. This enables a fair comparison of approaches within a category. Further, approaches within each category are compared separately, i.e., approaches across different categories are not compared as they try to achieve objectives by following different principles. This enables comparison of approaches following similar principles. It should also be noted that some approaches are not considered for comparison due to the reasons such as their well proven inferiority, lack of availability of application models, architecture models and tool chains. The used application model, architecture model and the tool chain are listed while presenting respective results in next subsections.

4.1   Hard Real-time Approaches

4.1.1   Guaranteed Admission Control based Approaches. The recently reported admission control based approaches that can be applied to applications represented as task graphs are considered for comparison and listed in Table 8. For the task graph, these approaches can use outputs from controllers to choose the core for critical path jobs or cores for the remaining jobs. The decision whether to use outputs from controllers or not for critical path and other jobs leads to four possible alternatives listed in Table 8 under column Approaches. We abbreviate them with four letter acronyms, where the two first letters denote whether the core selection for critical path jobs is done without (open loop - OL) or with (closed loop - CL) controllers and similarly the two remaining letters inform if the core selection for jobs outside the critical path is performed without (OL) or with (CL) controllers. Specifically, proportional–integral–derivative (PID) controllers are used to perform the comparative study. A Transaction-Level Modelling (TLM) simulation model developed in SystemC language has been used to evaluate the efficiency of the approaches. The controller

Table 8. Admission control based approaches considered for comparison with various open loop (OL) and closed loop (CL) combinations.

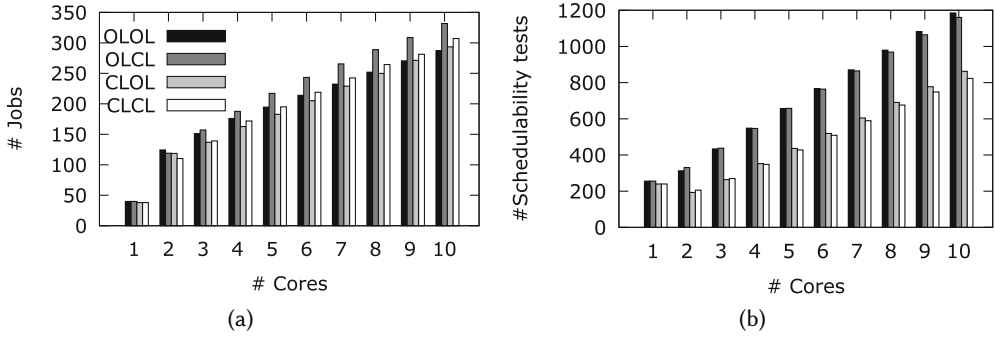| References	Approaches	Abbreviation
[127]	Open Loop for critical path jobs & Open Loop for other jobs	OLOL
[56]	Open Loop for critical path jobs & Closed Loop for other jobs	OLCL
[56]	Closed Loop for critical path jobs & Open Loop for other jobs	CLOL
[56]	Closed Loop for critical path jobs & Closed Loop for other jobs	CLCL

Fig. 10. Number of executed jobs and number of schedulability test executions.

components (proportional, integral, and derivative) are tuned by analysing the corresponding open-loop system response to a bursty workload.

Number of Executed Jobs and Number of Schedulability Test Executions. Figure 10 (a) and (b) show the number of jobs executed before their deadlines and number of schedulability tests when multi-core systems with different number of cores are considered. The shown results are for grid workload of an engineering design department of a large aircraft manufacturer that contains 100 tasks of 827 to 962 jobs in total, where job execution time varies from 1ms to 99ms, which has been scaled down for faster execution and was originally in the orders of hours. In the open-loop (OL) configuration, cores are scanned in a lexicographical order as long as the first one capable of executing the job satisfying its timing constraints is not found, whereas in the closed-loop (CL) configurations the tasks are checked with regards to the decreasing value of the corresponding controller outputs. A couple of observations can be made from Figure 10. *1)* For the number of executed jobs, OLOL configuration approach seems to be particularly beneficial in the systems with lower number of cores (heavier loaded with jobs). However, in the systems with more than two cores, the OLCL configuration leads to the best results. Its superiority in comparison with CLCL stems from the fact that an over-pessimistic rejection of critical path jobs leads to fast rejection of the whole task. Thus, the cost of a false negative estimation is rather high. The OLCL configuration admits 11% more jobs than OLOL, whereas CLCL is only slightly (about 1.5%) better than the baseline OLOL. *2)* For the number of schedulability tests, the difference between OLOL and OLCL is almost unnoticeable, but the configurations with control-theory-aided selection of a core for the critical path jobs, i.e. CLCL, leads to significant, over 30% reduction. This indicates the benefits of using the controller outputs. From the results, it follows that two configurations OLCL and CLCL dominate each other, the former in terms of number of executed jobs, the latter in terms of number

Table 9. Hybrid approaches (design-time allocations computation and run-time selection) considered for comparison.

| References	Approaches	Abbreviation
[175]	Hybrid Execution Time Optimization	HETOpt
[42]	Hybrid Energy Consumption Optimization 1	HECOpt1
[163]	Hybrid Energy Consumption Optimization 2	HECOpt2
[193]	Hybrid Execution Time & Energy Consumption Optimization 1	HETECOpt1
[168]	Hybrid Execution Time & Energy Consumption Optimization 2	HETECOpt2
[171]	Hybrid Execution Time & Energy Consumption Optimization 3	HETECOpt3
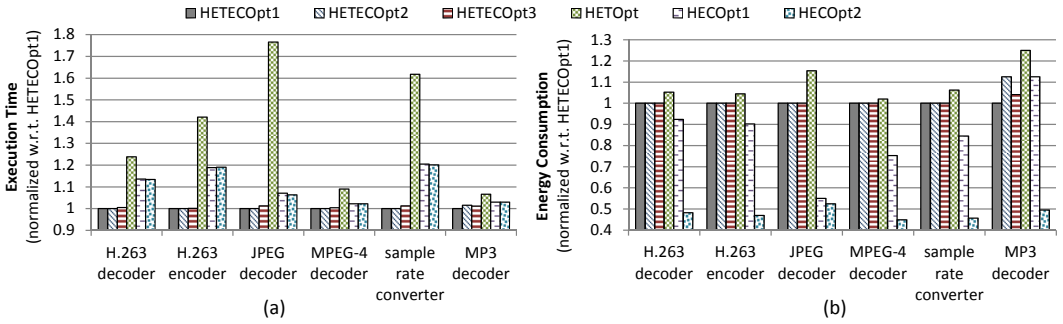
Fig. 11. Execution time and energy consumption comparison.

of schedulability tests. Depending upon which goal is more important, one of them is advised to be selected.

4.1.2 Hybrid Approaches: Design-time Allocations Computation and Run-time Selection. The compared hybrid approaches that perform design-time allocations computation and run-time selection are listed in Table 9. Some of these approaches optimize only execution time or energy consumption, and some of them optimize both the metrics. HETOpt tries to find a load balanced allocation to optimize execution time. The energy consumption optimization in HECOpt1 and HECOpt2 is performed by exploiting expected future slack and current slack, respectively. Appropriate voltage/frequency levels of used cores are identified to exploit the slack. In HETECOpt1, at design-time, exhaustive exploration is performed to identify the best allocation in terms of execution time and energy consumption, whereas HETECOpt2 prunes the design space to perform the exploration within a limited time. HETECOpt3 identifies the best allocation at run-time by exploiting the design-time extracted execution traces. To evaluate these approaches, they are implemented using publicly available SDF³ tool set [174] along with the required application and architecture models.

Execution Time and Energy Consumption Comparison. Figure 11 (a) and (b) show comparison of worst-case execution time and energy consumption for various streaming multimedia applications when different approaches tabulated in Table 9 are employed to map them on a 4×4 2D mesh multi-core chip. The multimedia applications are characterized by throughput constraints [10]. The results obtained by different approaches are normalized with respect to (w.r.t.) the result obtained by HETECOpt1. A couple of observations can be made from Figure 11 (a) and (b). *1)* HETOpt leads to the worst result in terms of both execution time and energy consumption as it tries to perform load balanced resource allocation of tasks on cores without taking parallelism of tasks

Table 10. Hybrid approaches (design-time deadline distribution and run-time allocation) considered for comparison.

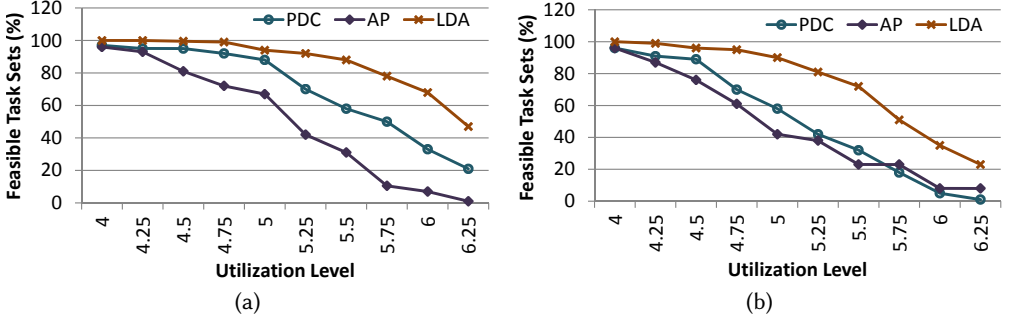| References	Approaches	Abbreviation
[88]	Pipeline Delay Composition	PDC
[27]	Automatic Partitioning	AP
[73]	Local-Deadline Assignment	LDA

Fig. 12. Percentage of feasible task sets for (a) balanced and (b) imbalanced workloads.

and their communication overhead into account. *2)* Both execution time and energy consumption optimization approaches HETECOpt1, HETECOpt2, and HETECOpt2 lead to similar results for most of the applications as all of them compute efficient allocations at design-time. *3)* Energy consumption by HECOpt1 and HECOpt2 is lower than other approaches. The reason lies in the fact that HECOpt1 and HECOpt2 employ DVFS on cores to reduce the energy consumption while respecting the application deadlines, whereas, other approaches just try to find the best allocation and do not employ DVFS. *4)* HECOpt2 leads to minimum energy consumption as DVFS is applied both at design-time and run-time while taking DVFS overhead and deadline into account.

4.1.3 Hybrid Approaches: Design-time Deadline Distribution and Run-time Allocation. The compared hybrid approaches that perform design-time deadline distribution and run-time allocation are listed in Table 10. These approaches distribute the end-to-end deadline of an application/job to its sub-functions and then perform resource allocation in order to meet the deadline. In PDC, end-to-end delay of a job in a multi-stage pipeline is bounded as a function of job execution times on different stages. The AP approach partitions a parallel real-time application into a set of sequential flows. In LDA, a locally optimal algorithm is employed to assign local deadlines to the jobs. The evaluation results for these approaches are employed from [73], where they are implemented in C++ and consider stream-type and general-type workloads to emulate different kinds of application scenarios.

Number of Executed Task Sets. Figure 12 (a) and (b) show comparison of feasible task sets at various utilization levels for balanced and imbalanced workloads when assorted approaches are employed. Each workload contains a total of 100 task sets of 50 tasks each, where individual instances of a task and sub-task are referred to as jobs and sub-jobs, respectively. A total of 8 cores are considered. In a balanced workload, the execution time of a job is randomly distributed along its execution path to balance the core loads, whereas cores' loads are imbalanced in an imbalanced

workload. A couple of observations can be made from the figure. *1)* LDA finds far more feasible sets than the other two methods for both the balanced and imbalanced workloads. *2)* LDA performs much better than PDC and AP at high utilization levels where there are more jobs in the system. It should also be noted that sometimes LDA may not be able to find a feasible solution even though such solutions indeed exist, since LDA finds local sub-job deadlines for each core independently instead of using a global approach.

4.2 Soft Real-time Approaches

4.2.1 Market-inspired Approaches. We have considered mainly market-inspired and value based approaches for comparison. Table 11 lists the compared approaches that have been applied to applications containing dependent tasks. The value optimization approaches optimize only for value while applying various principles, e.g., ValOpt1 chooses the highest value application first, ValOpt2 first chooses the application having maximum value density computed as value over the amount of required computational resources, and ValOpt3 chooses the application with the minimum remaining value first. The approach in [163] optimizes only for energy and has been extended to optimize both value and energy consumption for a fair comparison, where first a value optimizing allocation is found and then DVFS is applied to optimize energy consumption. In ValEnJoinOpt, value and energy consumption are jointly optimized by employing a genetic algorithm. In ValEnAdaptOpt, in addition to joint optimization, adaptation (reallocation) is also performed based on the execution status of running applications and available system resources in order to explore the scope for further optimization. These approaches are implemented in a C++ prototype and integrated with a SystemC functional simulator. They are evaluated by considering job models from historical data of an industrial HPC data center at High Performance Computing Center Stuttgart (HLRS).

Value and Energy Consumption Comparison. Figure 13 (a) and (b) show the influence of the number of available nodes (servers) on the overall value and energy consumption when various approaches are employed. Each node contains a total of 10 cores. To sufficiently stress the platform, we consider all the applications arriving over a month. Further, to remain close to the reality, it is considered that higher number of applications arrives in peak times, i.e. weekdays and daytimes as compared to off-peak times, i.e. weekends and night times. The value and energy consumption results are normalized w.r.t. the value and energy by ValOpt1 approach at 2 nodes. A couple of observations can be made from Figure 13. *1)* Overall value by all the approaches increases with the number of nodes due to increased processing capability leading to completion of a higher number of applications before their value becomes zero. *2)* ValEnAdaptOpt approach achieves a higher overall value than other approaches. This is due to the fact that adaptation leads to early completion of executing applications and thus higher values for them. Further, earlier completion leaves resources for the queued applications to be allocated and completed sooner, leading to higher values. *3)* The energy consumption is not the lowest by ValEnAdaptOpt as it completes execution of higher number of applications and energy is consumed for executing them. Further, if both the value and energy consumption metrics are to be jointly optimized as value achieved per unit of energy consumption, i.e. value divided by energy, ValEnAdaptOpt leads to the best results.

4.2.2 Bio-inspired Approaches. Table 12 lists the bio-inspired approaches considered for comparison. These approaches are applied to perform reallocation on top of an initial allocation done based on least utilized heuristic in order to improve average execution time and energy consumption. In RR, at every remapping interval, each core selects the most late task from its task queue and randomly selects another core to remap it. MCR partitions the whole many-core chip into virtual

Table 11. Market-inspired and value based approaches considered for comparison.

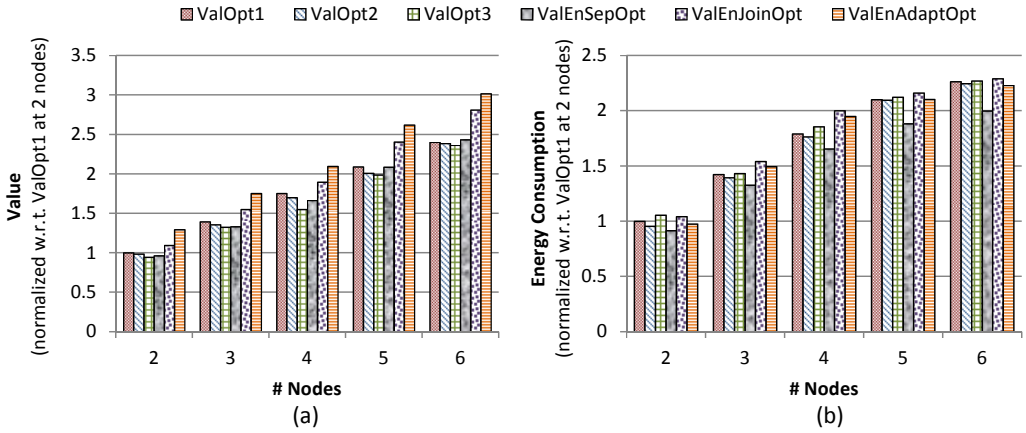| References	Approaches	Abbreviation
[179]	Value Optimization 1	ValOpt1
[11]	Value Optimization 2	ValOpt2
[25]	Value Optimization 3	ValOpt3
[163]	Value and Energy Separate Optimization	ValEnSepOpt
[166]	Value and Energy Joint Optimization	ValEnJoinOpt
[167]	Value and Energy Adaptive Optimization	ValEnAdaptOpt

Fig. 13. Value and energy consumption comparison.

clusters and each cluster is managed by a local manager that receives states of its cores every time a task completes execution and performs mapping of tasks. In case the cluster has no available cores, it sends request to neighboring clusters and loans the closest available core to map the task. To have a fair comparison, MCR approach of [32] has been modified to take the following aspects into account: relocation to maximum two hop distance, remapping a late task is to a PE with positive slack, and placing cluster manager in the center of the cluster [128]. SCR is essentially MCR with only one cluster. In PSR, the reallocation is performed based on the pheromone level of the cores. These approaches are evaluated by a discrete-event abstract simulator to map a video stream workload on a 10×10 mesh of cores [128], where cluster size of 2×5 (i.e. 10 clusters) is considered for MCR and the video stream contains a set of jobs arriving at different moments of time.

Execution Time and Energy Consumption Comparison. To evaluate average execution time and energy consumption, the relevant metrics job lateness and communication overhead are considered. By optimizing the job lateness and communication overhead, average execution time and communication energy can be optimized, where job lateness is computed as the difference between the job deadline and its response time. Figure 14 (a) and (b) compare distribution of cumulative job lateness and communication overhead, respectively, when various approaches are employed. For job lateness, all the approaches show positive and negative results. Therefore, the remapping techniques have failed to improve lateness of jobs under certain workload situations. However, majority of the distribution for PSR and MCR is in the positive region. In over 60% of the workload scenarios, PSR produces positive improvement to the job lateness of the video streams. Further, MCR shows a

Table 12. Bio-inspired approaches considered for comparison.

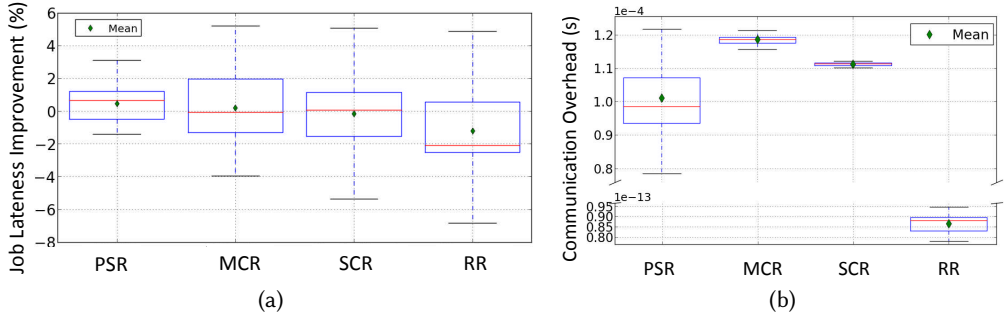| References	Approaches	Abbreviation
Reference implementation	Random Reallocation	RR
[32]	Multiple Cluster based Reallocation	MCR
[128]	Single Cluster based Reallocation	SCR
[128]	Pheromone Signalling based Reallocation	PSR

Fig. 14. Job lateness and communication overhead comparison.

better job lateness improvement over SCR as the monitoring traffic is shorter in route-length and hence is less disruptive to the data communication. The random remapper shows the worst results. It is interesting to note that there are a few scenarios where random remapping produced significant job lateness improvements. For communication overhead, PSR shows a significant reduction when compared to some other approaches. The maximum overhead of PSR is comparable to that of MCR. Both the MCR and SCR show a higher and narrower distribution of communication overhead than PSR. A higher upper whisker in PSR shows that under certain workload scenarios the overhead can be costly and similar to the MCR. The lower communication overhead distribution of the SCR when compared to MCR is due to the lack of inter-cluster communication. In SCR, communicating tasks mapped at the middle of the system will suffer due to the network congestion caused by the incoming monitoring traffic and thus communication overhead issues will become severe for larger system sizes. The RR has lowest communication overhead as it only incurs overhead when notifying the task dispatcher regarding remapping decisions.

4.2.3 Non-guaranteed Admission Control Approaches. Table 13 lists the relevant approaches considered for comparison. NAC is considered to analyse the effects of admission controls over no admission control. In GAC, it is ensured that all the admitted tasks/jobs will meet their deadlines. The EQF approach divides the total remaining slack among the subtasks in proportion to their estimated execution times in the hope to reduce individual task deadline miss rate. In HAC, a heuristic based admission control is employed in order to achieve trade-off between predictability and utilization. These approaches are evaluated by a discrete-event abstract simulator to map a number of video streams on a 3×3 mesh of cores [127], where each video stream contains a set of jobs arriving at different moments of time.

Admitted Video Streams (relates to Predictability) and Energy Consumption (relates to System Utilization) Comparison. Figure 15 (a) shows the number of video streams that were admitted

Table 13. Non-guaranteed admission control approaches considered for comparison.

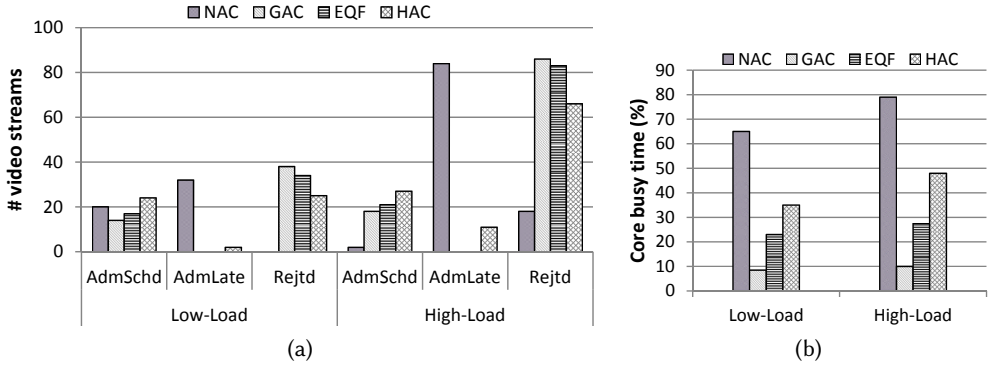| References	Approaches	Abbreviation
Reference implementation	No Admission Control	NAC
[109]	Guaranteed Admission Control	GAC
[95]	Equal Flexibility	EQF
[127]	Heuristic Admission Control	HAC

Fig. 15. Video stream admission and utilization comparison.

and successfully schedulable (AmSchd), admitted but late (AdmLate) and rejected (Rejtd) when difference approaches are employed under light (Low) and heavy (High) load conditions. The HAC is employed by setting fixed values of constants in the adopted heuristic [127]. A couple of observations can be made from Figure 15 (a). *1)* For both low and high load conditions, when GAC is employed, none of the admitted video streams incur any lateness, i.e., AdmLate is zero. *2)* In NAC, all the incoming video streams are admitted unless the global input buffers do not have available free space to hold the new tasks. Hence, a large number of admitted video streams miss their deadlines, specially under the high-load situation, and only few video streams are rejected. *3)* The rejected jobs are higher in high-load situation as the task-queues and input buffers become saturated earlier. *4)* EQF and GAC have similar service guarantees, where no admitted streams incur lateness (AdmLate = 0), however, on average, the rejection rate of EQF is lower than GAC, giving a more tighter guaranteed decision. Figure 15 (b) shows the percentage busy time of all the cores for different approaches under low and high load conditions. The percentages is computed by the ratio between the total core busy time and total simulation. A couple of observations can be made by looking Figure 15 (a) and (b). *1)* Admitting more streams into the system (whether they are late or schedulable) improves system utilization. In the high load condition, NAC has peak utilization of about 80% after which the buffers begin to overflow and streams are rejected. Further, admitting a few high resolution video streams may cause the system to be more busy than admitting relatively low resolution ones. *2)* GAC has the lowest system busy time which corresponds well with the high number of stream rejections shown in Figure 15 (a). *3)* The system busy time is better by EQF over GAC because of extra admitted streams. Further, HAC shows even higher busy time due to more admitted video streams.

Table 14. Congestion avoiding approaches considered for comparison.

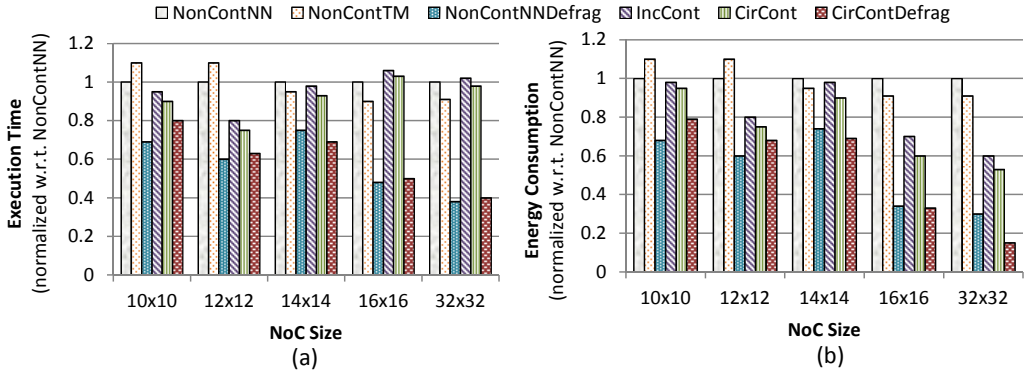| References	Approaches	Abbreviation
[31]	Non-contiguous Nearest Neighbour	NonContNN
[130]	Non-contiguous with Task Migration	NonContTM
[139]	NonContNN with Defragmentation	NonContNNDefrag
[41]	Incremental Contiguous	IncCont
[176]	Circular Contiguous	CirCont
[139]	CirCont with Defragmentation	CirContDefrag

Fig. 16. Execution time and energy consumption comparison.

4.2.4 Congestion Avoiding Approaches. Table 14 lists the promising congestion avoiding resource allocation approaches considered for comparison. NonContNN tries to find the allocation for a task on the nearest core close to the core containing its communicating task. NonContTM employs task migration to obtain a better allocation. NonContNNDefrag employs the defragmentation approach proposed in [139] on top of NonContNN. In IncCont, for the tasks of an incoming application, an incremental allocation is performed in a selected region that may be long and thin in shape, which may incur a higher communication energy. In contrast, CirCont selects a region that is nearly circular in shape, which leads to small average distance between any pair of cores inside the selected region and thus save energy consumption. CirContDefrag employs the defragmentation approach of [139] on top of CirCont. These approaches are implemented by extending Noxim [61], a SystemC-based Network-on-Chip (NoC) simulator. Their evaluation is done by considering applications represented as task graphs that have varying arrival times [139].

Execution Time and Energy Consumption Comparison. Figure 16 (a) and (b) show overall execution time and energy consumption when 100 random applications are allocated on systems of various sizes by employing the congestion avoiding approaches reported in Table 14. The inter-application arrival period is set randomly and average number of tasks in each application is 8. In the case of non-contiguous mapping, on average, NonContNNDefrag reduces execution time by 35.6% and 34% when compared to NonContNN and NonContTM, respectively. Further, NonContNNDefrag reduces energy consumption by 40.6% and 39% over NonContNN and NonContTM, respectively. In NonContNN and NonContTM, free cores might be scattered and thus tasks of the incoming application are allocated to non-contiguous regions. This increases the inter-task communication overhead and thus degrades the performance. On the other hand, NonContNNDefrag leads to better

performance as inter-task communication distance is reduced by grouping the free cores into one contiguous region before allocating of the incoming application. In the case of contiguous mapping, on average, CirContDefrag shows a 28% reduction in execution time when compared to CirCont. Moreover, on an average, CirContDefrag reduced energy consumption by 28% when compared to CirCont. CirContDefrag improves the CirCont approach by reshaping the irregular free core region into a regular and contiguous one with low communication distance; otherwise, CirCont has to wait until a regular free core region is not available even if sufficient scattered cores are available after multiple allocations and deallocations.

4.3 Discussions and Summary

The compared approaches under various categories indicate their superiority for performing different kinds of resource allocations. Therefore, based on the kind of application domain and run-time scenario, one category of approaches are recommended to be employed. The references to recommended previous studies for some common scenarios/applications are as follows. In case the platform is overloaded (i.e., demand for available resources is higher than the supply) in a centralized control environment, market-inspired approaches are suitable as more profitable jobs can be allocated to the limited resources while holding the low value jobs for late allocation [2, 29, 97, 167, 179]. In a distributed control environment where reallocations need to be performed based on observed biological phenomenon and related platform/application-level metrics (e.g., throughput and lateness), bio-inspired approaches can be employed to perform efficient resource allocation [12, 24, 90, 128, 138] . For typical embedded systems where the applications are known in advance, efficient allocations for each application or use-case (combination of active applications) can be computed at design-time with well-defined search heuristics. Such measures can be taken to achieve better results as compared to on-the-fly processing approaches [77, 148, 169, 171, 187, 195]. However, on-the-fly approaches are recommended to be applied when the applications are not known in advance as they lend well in such scenarios [31, 32, 59, 133, 139, 162].

Further, depending upon the size of the system, a decision to choose between centralized and distributed control can be made by performing experiments with various system sizes. Usually, distributed control leads to better results for large scale systems as it alleviates the monitoring traffic problem around the centralized manager [32, 94, 99].

5 UPCOMING TRENDS FOR FUTURE RESEARCH AND OPEN CHALLENGES

This section addresses some of the upcoming trends and challenges to be faced to take the mapping methodologies into the next era.

5.1 Hybrid Resource Allocation

The analysis of dynamic resource allocation approaches indicate that hybrid strategies, that combine design space exploration of design-time techniques with the run-time management in order to select mapping configurations that are best suited to newly arriving applications, lead to better results than on-the-fly strategies. The already known mapping configurations also improve predictability. Further, since they involve minimum computation at run-time, they facilitate for light-weight run-time platform manager. This speeds up the mapping process, i.e., mapping time, significantly over on-the-fly strategies.

The trend for the hybrid resource allocation was introduced earlier [169]. Since then, due to its potential, significant advances have taken place for embedded systems [57, 87, 93, 148, 171, 187]. It has also shown its potential for managing data center resources [166, 167]. With advancement, these approaches have tried to address the involved challenges such as reducing exploration time

for large scale applications/platforms and storage overheads for explored designs. However, since the problem size is continuously increasing because of the additional complexity of applications and the way they share increasingly sophisticated platforms, further studies are required.

Although the advantages of hybrid strategy seem promising, it comes with its own trade-offs due to inherent pseudo-dynamic nature and inability to handle new applications without available design-time exploration. With no doubt, hybrid strategies seem to be followed in the field of mapping methodologies but due to their nascent development and lack of in-depth examination, further development of design-time and on-the-fly mapping methodologies will continue hand-in-hand with hybrid strategies.

5.2 Large Scale Many-core Architectures

It is evident that technological enhancement will enable integration of hundreds and even thousands of cores in a single chip [22]. Recently, some large scale architectures have been introduced, e.g., Angstrom [72], Kalray's MPPA [50], and KiloCore chip [20]. These architectures impose a big challenge to manage their resources at run-time in a scalable manner. Some researchers have tried to address the scalability concern by developing distributed resource management strategies [3, 58, 94, 99]. These distributed strategies need to be further investigated such that the systems can be made more predictable and real-time challenges of foreseeable future can be addressed.

The resource allocation for large scale many-core architectures has already been a trend [169]. However, due to the integration of higher number of cores in recent years [20, 50] and its expected increase in upcoming years, resource allocation approaches need to advance in order to address the challenge of efficient exploitation of the abundant amount of cores.

In order to overcome the issues of large scale 2-dimensional (2D) many-core chips such as large area, power and signal transmission delay, integration of multiple layers of cores into a single device is taking place in order to realize 3-dimensional (3D) many-core architectures [44, 45, 198]. Despite having several advantages of 3D integration, the 3D high integration density brings major concern in the temperature increase that causes thermal hot spots and high temperature gradients. This might lead to an unreliable system and degraded performance. Efficient thermal management of 3D architectures is challenging and requires investigation of efficient resource allocation strategies to cater the reliability and performance concerns. The performance concern is specially important for real-time applications where predictable execution is desired.

The resource allocation studies for 3D architectures have also advanced. These advanced studies have tried to exploit the 3D architecture through various heuristics. For example, in [38], interconnect energy optimization is achieved by allocating heavily communicating edges to fast vertical links, and in [170], potential of 3D-neighborhood correlation available in spatial and temporal domains for 3D videos is exploited.

Further, these large scale architectures are expected to contain several types of cores, e.g., general purpose processor (GPP) cores, digital signal processing (DSP) cores, graphic processing unit (GPU) cores and field programmable gate array (FPGA), in order to meet functional and non-functional demands by exploiting their distinct features. Such heterogeneous integration will further increase the resource management challenges. Recently, small scale heterogeneous architectures have been released by industries, e.g., Samsung Exynos 5422 System-on-Chip [156] that powers the popular Samsung Galaxy S5 smartphone and the chip contains 4 ARM Cortex-A15 cores, 4 ARM Cortex-A7 cores and a six-core ARM Mali T628 MP6 GPU. Additionally, some development chips have been released, e.g., MediaTek's Helio X20 chip [124] containing Cortex-A72 cores, Cortex-A53 cores and Mali-T880 MP4 GPU cores, Xilinx's Zynq UltraScale+ EG MPSoC devices [191] containing ARM Cortex-A53 cores, Cortex-R5 cores and a Mali-400 MP2 GPU, and Intel's Xeon-FPGA hybrid

chip [82] containing 12-core Intel microprocessor with an Altera Arria10 FPGA. The number of heterogeneous cores in future architectures are going to increase towards realization of large scale heterogeneous architectures. Further, academia is also investigating large scale heterogeneous architectures, e.g. Invasive architecture containing thousands of cores [70] and Loki many-core architecture [14]. In order to exploit these architectures, new framework such as OpenCL [142] has been developed for writing programs that can execute across heterogeneous processing cores CPU, GPU, DSP, and FPGAs. Existing works have explored the possibility of executing OpenCL programs across CPUs and GPUs [65, 114, 147]. In future, efficient resource allocation strategies will need to be developed by taking abundant amount of cores and heterogeneity into account [54]. This will address the need to satisfy the ever increasing performance requirements of modern and future applications that can efficiently utilize large number of processor cores.

5.3 Joint Consideration of Computation and Communication Loads

Several earlier researches on multi/many-core systems indicate that only computation loads were considered and optimizations used to be performed only for such loads. Some of these works can be seen in Tables 1 to 14, where computation (comp.) is entered in the column computation and communication (comp. and comm.) consideration. These works consider loads as a set of independent tasks, which used to be true for single core systems where each task is executed sequentially. However, to efficiently exploit multi/many-core systems, the applications are partitioned into multiple tasks that usually have execution order and data dependencies, and such partitioned application are represented as task graphs, dataflow graphs, etc. Due to the dependency, once computation of a task is finished, it sends data (representing communication load) to other task that starts execution after receiving the data. This necessitates joint consideration of computation and communication loads during the process of resource allocation.

5.4 Multi-objective Resource Allocation

The resource allocation problem for multi/many-core systems has focused a lot to optimize execution time in order to meet the timing constraints, which is one of the important objective to achieve safe operation in time-critical systems or better user experience in embedded systems. However, modern multi/many-core systems need to be optimized for several other metrics along with the execution time due to various reasons such as increasing demand of energy, shrinking transistor sizes leading to unreliable systems due to leakage power and thermal issues, and security threats due to interaction with untrusted devices.

5.4.1 Jointly Optimizing Execution Time and Energy Consumption. The joint optimization of execution time and energy consumption is of paramount importance due to battery operated systems [59, 90, 148, 171], and limited availability of power for remote devices and HPC data centers [97, 166]. This helps to enhance the operational time while still meeting the timing requirements. As mentioned in Tables 1 to 14, recently, there has been extensive focus for such joint optimization. However, with the increase of application/platform complexity and power requirement of many-core systems, the developments of efficient resource allocation approaches to jointly optimize these two important metrics is expected to continue in foreseeable future.

5.4.2 Consideration for Multiple Objectives. As mentioned earlier, the requirements to jointly optimize for several performance metrics will grow in future. Along with the execution time and energy consumption, other important metrics are supposed to be temperature, reliability, fault-tolerance, and security. For example, three metrics execution time, energy consumption and temperature are optimized in [161]. The temperature is to be optimized to surmount its effect

on other metrics such as execution time [137] and reliability [185], and to reduce the cooling cost of many-core based HPC data centers [35]. Reliability aware resource allocation is desired to increase the mean time to failure of a system [169]. However, in case a fault has happened, resource reallocation is desired during the system operation in order to make the system fault-tolerant [47, 52, 60, 78, 106, 152, 157, 169, 182, 197]. The demand to support increasing number of applications within a device and then their communication with the applications in other devices has created security threats due to possible attack in the communication channels and interaction with untrusted devices. This indicates that resource allocation approaches need to optimize for security as well, which can be achieved by developing security-aware resource allocation approaches that should take into account other metrics as well.

The optimization for multiple objectives also increases the design space and thus the exploration time. Therefore, in case the design space exploration can be performed at design-time, e.g. for known application set, the design space need to be efficiently pruned so that Pareto-fronts for several conflicting objectives can be derived within a limited time. The design space increases even further with the heterogeneity of systems and thus exploration strategies might need to establish an upper limit on the heterogeneity in order to maintain low complexity.

5.5 Additional Challenges

Some additional challenges also need to be addressed to take the mapping methodologies into the next era. For example, development of efficient programming models for large scale and 3D architectures, efficient synchronization and control of concurrently executing tasks on such architectures and debugging of several concurrent executions if results are not as expected.

6 CONCLUSION

This paper provides a survey and comparative study of dynamic resource allocation strategies for multi/many-core systems. Specially, hard and soft real-time strategies optimizing for execution time and energy consumption are surveyed and compared. The kind of target architecture, e.g., homogeneous or heterogeneous, and resource control mechanism, e.g., centralized or distributed, have also been identified. The comparative study of these strategies has shown the superiority of some of them. Based on the analysis of the surveyed and compared resource allocation strategies, upcoming trends and open challenges are identified.

The research directions highlighted in this survey are expected to advance in future in order to tackle the identified open challenges. These advances will require to develop efficient resource allocation mechanisms to cater the need of future hard and soft real-time multi/many-core systems.

REFERENCES

[1] Waheed Ahmed, Muhammad Shafique, Lars Bauer, and Jörg Henkel. 2011. Adaptive resource management for simultaneous multitasking in mixed-grained reconfigurable multi-core processors. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. ACM, 365–374.

[2] B. Aksanli and T. Rosing. 2014. Providing regulation services and managing data center peak power budgets. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1–4.

[3] Mohammad Abdullah Al Faruque, Rudolf Krist, and Jörg Henkel. 2008. ADAM: run-time agent-based distributed application mapping for on-chip communication. In *Proceedings of ACM Design Automation Conference (DAC)*. 760–765.

[4] AMD. 2011. AMD Opteron 6000 series processors. (2011). http://www.amd.com/en-us/products/server/opteron/6000 (Last visited: 12 February, 2016).

[5] F. Angiolini, Jianjiang Ceng, R. Leupers, F. Ferrari, C. Ferri, and L. Benini. 2006. An Integrated Open Framework for Heterogeneous MPSoC Design Space Exploration. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1 –6.

[6] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A.J. Wellings. 1993. Applying new scheduling theory to static priority pre-emptive scheduling. *Software Engineering Journal* 8, 5 (Sep 1993), 284–292.

[7] Neil C. Audsley, Alan Burns, Robert I. Davis, Ken W. Tindell, and Andy J. Wellings. 1995. Fixed Priority Pre-emptive Scheduling: An Historical Perspective. *Real-Time Syst.* 8, 2-3 (1995), 173–198.

[8] Ozalp Babaoglu, Geoffrey Canright, Andreas Deutsch, Gianni A Di Caro, Frederick Ducatelle, Luca M Gambardella, Niloy Ganguly, Márk Jelasity, Roberto Montemanni, Alberto Montresor, and others. 2006. Design patterns from biology for distributed computing. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 1, 1 (2006), 26–66.

[9] Theodore P Baker and Sanjoy K Baruah. 2007. Schedulability analysis of multiprocessor sporadic task systems. *Handbook of Real-Time and Embedded Systems* (2007).

[10] Mohamed A. Bamakhrama and Todor Stefanov. 2012. Managing latency in embedded streaming applications under hard-real-time scheduling. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 83–92.

[11] Nikhil Bansal and Kirk R. Pruhs. 2010. Server Scheduling to Balance Priorities, Fairness, and Average Quality of Service. *SIAM J. Comput.* 39, 7 (2010), 3311–3335.

[12] Donato Barbagallo, Elisabetta Di Nitto, Daniel J Dubois, and Raffaela Mirandola. 2010. A bio-inspired algorithm for energy optimization in a self-organizing data center. In *Self-Organizing Architectures*. Springer, 127–151.

[13] Sanjoy Baruah, Vincenzo Bonifaci, Alberto Marchetti-Spaccamela, and Sebastian Stiller. 2010. Improved multiprocessor global schedulability analysis. *Real-Time Systems* 46, 1 (2010), 3–24.

[14] Daniel Bates, Alex Bradbury, Andreas Koltes, and Robert Mullins. 2015. Exploiting tightly-coupled cores. *Journal of Signal Processing Systems* 80, 1 (2015), 103–120.

[15] L. Benini and G. De Micheli. 2002. Networks on chips: a new SoC paradigm. *Computer* 1 (2002), 70–78.

[16] M. Bertogna, M. Cirinei, and G. Lipari. 2005. Improved schedulability analysis of EDF on multiprocessor platforms. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECTRS)*. 209–218.

[17] Benjamin Betting and Uwe Brinkschulte. 2014. Analyzing the Overhead of Self-Optimization through Task Migration within a Decentralized Task Control Mechanism for Dependable System-on-Chip Architectures. In *IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC)*. IEEE, 84–91.

[18] Tobias Bjerregaard and Shankar Mahadevan. 2006. A survey of research and practices of Network-on-chip. *ACM Comput. Surv.* 1 (2006).

[19] Carolina Blanch, Rogier Baert, Paul Coene, Maja D'Hondt, Zhe Ma, and Roel Wuyts. 2011. Runtime Scheduling for Video Decoding on Heterogeneous Architectures. In *Proceedings of the International Conference on Real Time and Networks Systems (RTNS)*. Citeseer, 195–204.

[20] B. Bohnenstiehl, A. Stillmaker, J. Pimentel, T. Andreas, B. Liu, A. Tran, E. Adeagbo, and B. Baa. 2016. A 5.8 pJ/Op 115 Billion Ops/sec, to 1.78 Trillion Ops/sec 32nm 1000 Processor Array. In *IEEE Symposia on VLSI Technology and Circuits*.

[21] Vincenzo Bonifaci, Alberto Marchetti-Spaccamela, Sebastian Stiller, and Andreas Wiese. 2013. Feasibility analysis in the sporadic DAG task model. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECRTS)*. IEEE, 225–233.

[22] Shekhar Borkar. 2007. Thousand core chips: a technology perspective. In *Proceedings of ACM Design Automation Conference (DAC)*. 746–749.

[23] Eduardo Wenzel Briáo, Daniel Barcelos, and Flávio Rech Wagner. 2008. Dynamic task allocation strategies in MPSoC for soft real-time applications. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1386–1389.

[24] Uwe Brinkschulte, Mathias Pacher, and Alexander Von Renteln. 2007. Towards an artificial hormone system for self-organizing real-time task allocation. In *Software Technologies for Embedded and Ubiquitous Systems*. Springer, 339–347.

[25] Andrew Marc Burkimsher. 2014. *Fair, responsive scheduling of engineering workflows on computing grids*. Ph.D. Dissertation. UK.

[26] A. Burns, D. Prasad, A. Bondavalli, F. Di Giandomenico, K. Ramamritham, J. Stankovic, and L. Strigini. 2000. The Meaning and Role of Value in Scheduling Flexible Real-time Systems. *J. Syst. Archit.* 46, 4 (2000), 305–325.

[27] Giorgio Buttazzo, Enrico Bini, and Yifan Wu. 2011. Partitioning real-time applications over multicore reservations. *IEEE Transactions on Industrial Informatics* 7, 2 (2011), 302–315.

[28] Giorgio C. Buttazzo. 2011. *Hard Real-Time Computing Systems*. Real-Time Systems Series, Vol. 24. Springer US.

[29] Rodrigo N. Calheiros and Rajkumar Buyya. 2014. Energy-Efficient Scheduling of Urgent Bag-of-Tasks Applications in Clouds Through DVFS. In *Proceedings of IEEE International Conference on Cloud Computing Technology and Science (CLOUDCOM)*. 342–349.

[30] Scott Camazine, Nigel R. Franks, James Sneyd, Eric Bonabeau, Jean-Louis Deneubourg, and Guy Theraula. 2001. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, USA.

[31] Ewerson Luiz de Souza Carvalho, Ney Laert Vilar Calazans, and Fernando Gehm Moraes. 2010. Dynamic Task Mapping for MPSoCs. *IEEE Des. Test* (2010), 26–35.

[32] Guilherme Castilhos, Marcelo Mandelli, Guilherme Madalozzo, and Filipe Moraes. 2013. Distributed resource manage-ment in NoC-based MPSoCs with dynamic cluster sizes. In *Proceedings of IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 153–158.

[33] Jeronimo Castrillon, Andreas Tretter, Rainer Leupers, and Gerd Ascheid. 2012. Communication-aware mapping of KPN applications onto heterogeneous MPSoCs. In *Proceedings of ACM Design Automation Conference (DAC)*. 1266–1271.

[34] J. Ceng, J. Castrillon, W. Sheng, H. Scharwächter, R. Leupers, G. Ascheid, H. Meyr, T. Isshiki, and H. Kunieda. 2008. MAPS: an integrated framework for MPSoC application parallelization. In *Proceedings of ACM Design Automation Conference (DAC)*. 754–759.

[35] Muhammad Tayyab Chaudhry, Teck Chaw Ling, Atif Manzoor, Syed Asad Hussain, and Jongwon Kim. 2015. Thermal-aware scheduling in green data centers. *ACM Computing Surveys (CSUR)* 47, 3 (2015), 39.

[36] Ken Chen and Paul Muhlethaler. 1996. A Scheduling Algorithm for Tasks Described by Time Value Function. *Real-Time Syst.* 10, 3 (1996), 293–312.

[37] Liang Chen, T. Marconi, and T. Mitra. 2012. Online scheduling for multi-core shared reconfigurable fabric. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 582 –585.

[38] Yuanqing Cheng, Lei Zhang, Yinhe Han, and Xiaowei Li. 2013. Thermal-constrained Task Allocation for Interconnect Energy Reduction in 3-D Homogeneous MPSoCs. *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)* 21, 2 (2013), 239–249.

[39] Junchul Choi, Hyunok Oh, Sungchan Kim, and Soonhoi Ha. 2012. Executing synchronous dataflow graphs on a SPM-based multicore architecture. In *Proceedings of ACM Design Automation Conference (DAC)*. 664–671.

[40] Chen-Ling Chou and Radu Marculescu. 2008. User-aware dynamic task allocation in networks-on-chip. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1232–1237.

[41] Chen-Ling Chou, U. Y. Ogras, and R. Marculescu. 2008. Energy- and Performance-Aware Incremental Mapping for Networks on Chip With Multiple Voltage Levels. *Trans. Comp.-Aided Des. Integ. Cir. Sys.* (Oct. 2008), 1866–1879.

[42] Pravanjan Choudhury, P. P. Chakrabarti, and Rajeev Kumar. 2007. Online Dynamic Voltage Scaling Using Task Graph Mapping Analysis for Multiprocessors. In *Proceedings of the International Conference on VLSI Design (VLSID)*. 89–94.

[43] Jason Cong and Karthik Gururaj. 2009. Energy Efficient Multiprocessor Task Scheduling Under Input-dependent Variation. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 411–416.

[44] Ayse K. Coskun, Jose L. Ayala, David Atienza, Tajana Simunic Rosing, and Yusuf Leblebici. 2009. Dynamic thermal management in 3D multicore architectures. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1410–1415.

[45] Marco Cox, Amit Kumar Singh, Akash Kumar, and Henk Corporaal. 2013. Thermal-aware mapping of streaming applications on 3D Multi-Processor Systems. In *Proceedings of IEEE/ACM/IFIP Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia)*. 11–20.

[46] Antoniel da Silva Rego, Joaquim Celestino, Andre Dos Santos, Eduardo Coelho Cerqueira, Anup Patel, and Mehdi Taghavi. 2012. BEE-C: a bio-inspired energy efficient cluster-based algorithm for data continuous dissemination in Wireless Sensor Networks. In *IEEE International Conference on Networks (ICON)*. IEEE, 405–410.

[47] Anup Das, Amit Kumar Singh, and Akash Kumar. 2013. Energy-aware dynamic reconfiguration of communication-centric applications for reliable MPSoCs. In *IEEE International Workshop on Reconfigurable and Communication-Centric Systems-on-Chip (ReCoSoC)*. 1–7.

[48] Robert I. Davis and Alan Burns. 2011. A Survey of Hard Real-time Scheduling for Multiprocessor Systems. *ACM Comput. Surv.* 43, 4, Article 35 (Oct. 2011), 44 pages.

[49] Robert I Davis, Attila Zabos, and Alan Burns. 2008. Efficient exact schedulability tests for fixed priority real-time systems. *IEEE Trans. Comput.* 57, 9 (2008), 1261–1276.

[50] Benoît Dupont De Dinechin, Duco Van Amstel, Marc Poulhiès, and Guillaume Lager. 2014. Time-critical computing on a single-chip massively parallel processor. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1–6.

[51] André de Matos Pedro, David Pereira, Luís Miguel Pinho, and Jorge Sousa Pinto. 2015. Logic-based schedulability analysis for compositional hard real-time embedded systems. *ACM SIGBED Review* 12, 1 (2015), 56–64.

[52] O. Derin, D. Kabakci, and L. Fiorin. 2011. Online task remapping strategies for fault-tolerant Network-on-Chip multiprocessors. In *IEEE/ACM Symposium on Networks on Chip (NoCS)*. 129 –136.

[53] Michael L Dertouzos and Aloysius Ka-Lau Mok. 1989. Multiprocessor online scheduling of hard-real-time tasks. *IEEE Transactions on Software Engineering* 15, 12 (1989), 1497–1506.

[54] Jiun-Hung Ding, Ya-Ting Chang, Zhou-dong Guo, Kuan-Ching Li, and Yeh-Ching Chung. 2014. An efficient and comprehensive scheduler on Asymmetric Multicore Architecture systems. *Journal of Systems Architecture* 60, 3 (2014), 305–314.

[55] Piotr Dziurzanski, Amit Kumar Singh, and Leandro Soares Indrusiak. 2016. Energy-Aware Resource Allocation in Multi-mode Automotive Applications with Hard Real-Time Constraints. In *IEEE International Symposium on Real-Time Distributed Computing (ISORC)*. 100–107.

[56] Piotr Dziurzanski, Amit Kumar Singh, and Leandro Soares Indrusiak. 2016. Feedback-Based Admission Control for Hard Real-Time Task Allocation under Dynamic Workload on Many-core Systems. In *Proceedings of the International Conference on Architecture of Computing Systems (ARCS)*. 157–169.

[57] Piotr Dziurzanski, Amit Kumar Singh, Leandro Soares Indrusiak, and Björn Saballus. 2015. Hard Real-time Guarantee of Automotive Applications During Mode Changes. In *Proceedings of the International Conference on Real Time and Networks Systems (RTNS)*. 161–170.

[58] Thomas Ebi, David Kramer, Wolfgang Karl, and Jörg Henkel. 2011. Economic learning for thermal-aware power budgeting in many-core architectures. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 189–196.

[59] M. Fattah, P. Liljeberg, J. Plosila, and H. Tenhunen. 2014. Adjustable contiguity of run-time task allocation in networked many-core systems. In *Proceedings of IEEE Asia and South Pacific Design Automation Conference (ASP-DAC)*. 349–354.

[60] Mohammad Fattah, Maurizio Palesi, Pasi Liljeberg, Juha Plosila, and Hannu Tenhunen. 2014. SHiFA: System-Level Hierarchy in Run-Time Fault-Aware Management of Many-Core Systems. In *Proceedings of ACM Design Automation Conference (DAC)*. 101:1–101:6.

[61] Fabrizio Fazzino, Maurizio Palesi, and David Patti. 2008. Noxim: Network-on-chip simulator. *URL: http://sourceforge. net/projects/noxim* (2008).

[62] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability; A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co.

[63] Beltra Giovanni, Luca Fossati, and Donatella Sciuto. 2010. Decision-theoretic design space exploration of multiprocessor platforms. *Trans. Comp.-Aided Des. Integ. Cir. Sys.* (2010), 1083–1095. Issue 7.

[64] Maximilian Götzinger, Amir M Rahmani, Martin Pongratz, Pasi Liljeberg, Axel Jantsch, and Hannu Tenhunen. 2016. The role of self-awareness and hierarchical agents in resource management for many-core systems. In *IEEE International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoC)*. IEEE, 53–60.

[65] Dominik Grewe, Zheng Wang, and Michael FP OŠBoyle. 2013. OpenCL task partitioning in the presence of GPU contention. In *International Workshop on Languages and Compilers for Parallel Computing*. 87–101.

[66] Abdul Hameed, Alireza Khoshkbarforoushha, Rajiv Ranjan, Prem Prakash Jayaraman, Joanna Kolodziej, Pavan Balaji, Sherali Zeadally, Qutaibah Marwan Malluhi, Nikos Tziritas, Abhinav Vishnu, and others. 2014. A survey and taxonomy on energy efficient resource allocation techniques for cloud computing systems. *Computing* (2014), 1–24.

[67] Ching-Chih Han and Kwei-Jay Lin. 1989. Scheduling parallelizable jobs on multiprocessors. In *Proceedings Real Time Systems Symposium (RTSS)*. IEEE, 59–67.

[68] Hans-Ulrich Heiss and Michael Schmitz. 1995. Decentralized dynamic load balancing: The particles approach. *Information Sciences* 84, 1 (1995), 115–128.

[69] Jörg Henkel, Lars Bauer, Nikil Dutt, Puneet Gupta, Sani Nassif, Muhammad Shafique, Mehdi Tahoori, and Norbert Wehn. 2013. Reliable On-chip Systems in the Nano-era: Lessons Learnt and Future Trends. In *Proceedings of ACM Design Automation Conference (DAC)*. 99:1–99:10.

[70] J. Henkel, A. Herkersdorf, L. Bauer, T. Wild, M. Hubner, R.K. Pujari, A. Grudnitsky, J. Heisswolf, A. Zaib, B. Vogel, V. Lari, and S. Kobbe. 2012. Invasive manycore architectures. In *Proceedings of IEEE Asia and South Pacific Design Automation Conference (ASP-DAC)*. 193 –200.

[71] Fernando Herrera and Ingo Sander. 2013. Combining analytical and simulation-based design space exploration for time-critical systems. In *Forum on Specification Design Languages (FDL)*. 1–8.

[72] Henry Hoffmann, Jim Holt, George Kurian, Eric Lau, Martina Maggio, Jason E. Miller, Sabrina M. Neuman, Mahmut Sinangil, Yildiz Sinangil, Anant Agarwal, Anantha P. Chandrakasan, and Srinivas Devadas. 2012. Self-aware computing in the Angstrom processor. In *Proceedings of ACM Design Automation Conference (DAC)*. 259–264.

[73] Shengyan Hong, T. Chantem, and Xiaobo Sharon Hu. 2015. Local-Deadline Assignment for Distributed Real-Time Systems. *IEEE Trans. Comput.* 64, 7 (July 2015), 1983–1997.

[74] S. Hong, S. H. K. Narayanan, M. Kandemir, and Ö. Özturk. 2009. Process variation aware thread mapping for chip multiprocessors. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 821–826.

[75] Jingcao Hu and Radu Marculescu. 2003. Energy-aware mapping for tile-based NoC architectures under performance constraints. In *Proceedings of IEEE Asia and South Pacific Design Automation Conference (ASP-DAC)*. 233–239.

[76] Jia Huang, Andreas Raabe, Christian Buckl, and Alois Knoll. 2011. A workflow for runtime adaptive task allocation on heterogeneous mpsocs. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1–6.

[77] Lin Huang and Qiang Xu. 2010. Performance yield-driven task allocation and scheduling for MPSoCs under process variation. In *Proceedings of ACM Design Automation Conference (DAC)*. 326 –331.

[78] Lin Huang, Rong Ye, and Qiang Xu. 2011. Customer-aware task allocation and scheduling for multi-mode MPSoCs. In *Proceedings of ACM Design Automation Conference (DAC).* 387 −392.

[79] Hameed Hussain, Saif Ur Rehman Malik, Abdul Hameed, Samee Ullah Khan, Gage Bickler, Nasro Min-Allah, Muhammad Bilal Qureshi, Limin Zhang, Wang Yongji, Nasir Ghani, Joanna Kolodziej, Albert Y. Zomaya, Cheng-Zhong Xu, Pavan Balaji, Abhinav Vishnu, Fredric Pinel, Johnatan E. Pecero, Dzmitry Kliazovich, Pascal Bouvry, Hongxiang Li, Lizhe Wang, Dan Chen, and Ammar Rayes. 2013. Review: A Survey on Resource Allocation in High Performance Distributed Computing Systems. *Parallel Comput.* 39, 11 (Nov. 2013), 709−736.

[80] Leandro Soares Indrusiak. 2014. End-to-end schedulability tests for multiprocessor embedded systems based on networks-on-chip with priority-preemptive arbitration. *Journal of Systems Architecture* 60, 7 (2014), 553 − 561.

[81] Leandro Soares Indrusiak, Piotr Dziurzanski, and Amit Kumar Singh. 2016. *Dynamic Resource Allocation in Embedded, High-Performance and Cloud Computing.* River Publishers.

[82] Intel. 2016. Intel Hardware Accelerator Research Program. (2016). https://www.nextplatform.com/2016/03/14/intel-marrying-fpga-beefy-broadwell-open-compute-future/.

[83] David E. Irwin, Laura E. Grit, and Jeffrey S. Chase. 2004. Balancing Risk and Reward in a Market-Based Task Service. In *IEEE International Symposium on High Performance Distributed Computing (HPDC).* 160−169.

[84] Damir Isovic and Gerhard Fohler. 2004. Quality aware MPEG-2 stream adaptation in resource constrained systems. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECRTS).* 23−32.

[85] Ralf Jahr, Martin Frieb, Mike Gerdes, Theo Ungerer, Andreas Hugl, and Hans Regler. 2014. Paving the way for multi-cores in industrial hard real-time control applications. In *IEEE International Symposium on Industrial Embedded Systems (SIES).* 1−4.

[86] Haris Javaid and Sri Parameswaran. 2009. A design flow for application specific heterogeneous pipelined multiprocessor systems. In *Proceedings of ACM Design Automation Conference (DAC).* 250−253.

[87] Haris Javaid, Muhammad Shafique, Jorg Henkel, and Sri Parameswaran. 2014. Energy-Efficient Adaptive Pipelined MPSoCs for Multimedia Applications. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 33, 5 (2014), 663−676.

[88] Praveen Jayachandran and Tarek Abdelzaher. 2008. Delay composition in preemptive and non-preemptive real-time pipelines. *Real-Time Systems* 40, 3 (2008), 290−320.

[89] Ahmed Jerraya, Hannu Tenhunen, and Wayne Wolf. 2005. Guest Editors' Introduction: Multiprocessor Systems-on-Chips. *Computer* 7 (2005), 36−40.

[90] Vaibhav Jha, Mohit Jha, and GK Sharma. 2014. Estimation of Optimized Energy and Latency Constraints for Task Allocation in 3d Network on Chip. *arXiv preprint arXiv:1405.0109* (2014).

[91] Zai Jian Jia, A.D. Pimentel, M. Thompson, T. Bautista, and A. Nunez. 2010. NASA: A generic infrastructure for system-level MP-SoC design space exploration. In *Proceedings of IEEE/ACM/IFIP Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia).* 41 −50.

[92] Jan Jonsson and Kang G Shin. 2002. Robust adaptive metrics for deadline assignment in distributed hard real-time systems. *Real-Time Systems* 23, 3 (2002), 239−271.

[93] Hanwoong Jung, Chanhee Lee, Shin-Haeng Kang, Sungchan Kim, Hyunok Oh, and Soonhoi Ha. 2014. Dynamic Behavior Specification and Dynamic Mapping for Real-Time Embedded Systems: HOPES Approach. *ACM Transactions on Embedded Computing Systems (TECS)* 13, 4s, Article 135 (2014), 26 pages.

[94] Michael Kadin, Sherief Reda, and Augustus Uht. 2009. Central vs. distributed dynamic thermal management for multi-core processors: which one is better?. In *Proceedings of ACM Great Lakes symposium on VLSI (GLSVLSI).* 137−140.

[95] B. Kao and H. Garcia-Molina. 1997. Deadline assignment in a distributed soft real-time system. *IEEE Transactions on Parallel and Distributed Systems* 8, 12 (Dec 1997), 1268−1274.

[96] Samarth Kaushik, Amit Kumar Singh, Wu Jigang, and Thambipillai Srikanthan. 2011. Run-time computation and communication aware mapping heuristic for NoC-based heterogeneous MPSoC platforms. In *IEEE International Symposium on Parallel Architectures, Algorithms and Programming (PAAP).* 203−207.

[97] Bhavesh Khemka, Ryan Friese, Sudeep Pasricha, Anthony A Maciejewski, Howard Jay Siegel, Gregory A Koenig, Sarah Powers, Marcia Hilton, Rajendra Rambharos, and Steve Poole. 2015. Utility maximizing dynamic resource management in an oversubscribed energy-constrained heterogeneous computing system. *Sustainable Computing: Informatics and Systems* 5 (2015), 14−30.

[98] Joonsoo Kim and Michael Orshansky. 2006. Towards formal probabilistic power-performance design space exploration. In *Proceedings of ACM Great Lakes symposium on VLSI (GLSVLSI).* 229−234.

[99] Sebastian Kobbe, Lars Bauer, Daniel Lohmann, Wolfgang Schröder-Preikschat, and Jörg Henkel. 2011. DistRM: distributed resource management for on-chip many-core systems. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES).* 119−128.

[100] Jonathan Koomey. 2011. Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times* (2011).

[101] Akash Kumar, Henk Corporaal, Bart Mesman, and Yajun Ha. 2010. *Multimedia Multiprocessor Systems: Analysis, Design and Management.* Springer Science & Business Media.

[102] Tei-Wei Kuo, Li-Pin Chang, Yu-Hua Liu, and Kwei-Jay Lin. 2003. Efficient online schedulability tests for real-time systems. *IEEE Transactions on Software Engineering* 29, 8 (2003), 734–751.

[103] Yu-Kwong Kwok and others. 2006. A semi-static approach to mapping dynamic iterative tasks onto heterogeneous computing systems. *J. Parallel Distrib. Comput.* 66, 1 (2006), 77–98.

[104] Karthik Lakshmanan, Dionisio de Niz, and Ragunathan Rajkumar. 2009. Coordinated task scheduling, allocation and synchronization on multiprocessors. In *IEEE Real-Time Systems Symposium (RTSS)*. 469–478.

[105] Sylvain Lauzac, Rami Melhem, and Daniel Mosse. 1998. An efficient RMS admission control and its application to multiprocessor scheduling. In *IEEE Symposium on Parallel and Distributed Processing (SPDP)*. 511–518.

[106] Chanhee Lee, Hokeun Kim, Hae-woo Park, Sungchan Kim, Hyunok Oh, and Soonhoi Ha. 2010. A task remapping technique for reliable multi-core embedded systems. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 307–316.

[107] Chanhee Lee, Sungchan Kim, and Soonhoi Ha. 2013. Efficient run-time resource management of a manycore accelerator for stream-based applications. In *Proceedings of IEEE/ACM/IFIP Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia)*. 51–60.

[108] Jinkyu Lee, Insik Shin, and Arvind Easwaran. 2012. Convex optimization framework for intermediate deadline assignment in soft and hard real-time distributed systems. *Journal of Systems and Software* 85, 10 (2012), 2331–2339.

[109] Hennadiy Leontyev and James H Anderson. 2008. A unified hard/soft real-time schedulability test for global EDF multiprocessor scheduling. In *IEEE Real-Time Systems Symposium (RTSS)*. 375–384.

[110] Yu-Chia Lin, Chuan-Yue Yang, Che-Wei Chang, Yuan-Hao Chang, Tei-Wei Kuo, and Chi-Sheng Shih. 2010. Energy-Efficient Mapping Technique for Virtual Cores. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECRTS)*. 66–75.

[111] Michele Lombardi and Michela Milano. 2012. Optimal methods for resource allocation and scheduling: a cross-disciplinary survey. *Constraints* 17, 1 (2012), 51–85.

[112] Jose M Lopez, Jose L Diaz, and Daniel F Garcia. 2004. Minimum and maximum utilization bounds for multiprocessor rate monotonic scheduling. *IEEE Transactions on Parallel and Distributed Systems (TPDS)* 15, 7 (2004), 642–653.

[113] Chenyang Lu, John A Stankovic, Sang H Son, and Gang Tao. 2002. Feedback control real-time scheduling: Framework, modeling, and algorithms*. *Real-Time Systems* 23, 1-2 (2002), 85–126.

[114] Chi-Keung Luk, Sunpyo Hong, and Hyesoon Kim. 2009. Qilin: exploiting parallelism on heterogeneous multiprocessors with adaptive mapping. In *Proceedings of IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 45–55.

[115] A. Mallik and others. 2011. MNEMEE - An Automated Toolflow for Parallelization and Memory Management in MPSoC Platforms. In *Proceedings of ACM Design Automation Conference (DAC)*.

[116] Marcelo Mandelli, Luciano Ost, Everton Carara, Guilherme Guindani, Thiago Gouvea, Guilherme Medeiros, and Fernando G Moraes. 2011. Energy-aware dynamic task mapping for NoC-based MPSoCs. In *IEEE International Symposium on Circuits and Systems (ISCAS)*. 1676–1679.

[117] G Manimaran, C Siva Ram Murthy, and Krithi Ramamritham. 1998. A new approach for scheduling of parallelizable tasks in real-time multiprocessor systems. *Real-Time Systems* 15, 1 (1998), 39–60.

[118] C.A.M. Marcon, E.I. Moreno, N.L.V. Calazans, and F.G. Moraes. 2008. Comparison of network-on-chip mapping algorithms targeting low energy consumption. *IET Computers Digital Techniques* (2008), 471 –482.

[119] R. Marculescu, U.Y. Ogras, Li-Shiuan Peh, N.E. Jerger, and Y. Hoskote. 2009. Outstanding Research Problems in NoC Design: System, Microarchitecture, and Circuit Perspectives. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 1 (2009), 3–21.

[120] G. Mariani, P. Avasare, G. Vanmeerbeeck, C. Ykman-Couvreur, G. Palermo, C. Silvano, and V. Zaccaria. 2010. An industrial design space exploration framework for supporting run-time resource management on multi-core systems. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 196–201.

[121] Dana Marinca, Pascale Minet, and Laurent George. 2004. Analysis of deadline assignment methods in distributed real-time systems. *Computer Communications* 27, 15 (2004), 1412–1423.

[122] G. Martin. 2006. Overview of the MPSoC design challenge. In *Proceedings of ACM Design Automation Conference (DAC)*. 274 –279.

[123] Peter Marwedel, Jürgen Teich, Georgia Kouveli, Iuliana Bacivarov, Lothar Thiele, Soonhoi Ha, Chanhee Lee, Qiang Xu, and Lin Huang. 2011. Mapping of applications to MPSoCs. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 109–118.

[124] MediaTek. 2016. MediaTek's Helio X20 Chip. (2016). http://www.mediatek.com/products/mediatek-helio.

[125]   A. Mehran, A. Khademzadeh, and S. Saeidi. 2008. DSM: A Heuristic Dynamic Spiral Mapping algorithm for network on chip. *IEICE Electronics Express* 13 (2008), 464–471.

[126]   A. Melani, M. Bertogna, V. Bonifaci, A. Marchetti-Spaccamela, and G.C. Buttazzo. 2015. Response-Time Analysis of Conditional DAG Tasks in Multiprocessor Systems. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECRTS)*. 211–221.

[127]   Hashan Roshantha Mendis, Leandro Soares Indrusiak, and Neil C Audsley. 2014. Predictability and utilisation trade-off in the dynamic management of multiple video stream decoding on network-on-chip based homogeneous embedded multi-cores. In *Proceedings of the International Conference on Real-Time Networks and Systems (RTNS)*. 161–170.

[128]   Hashan R Mendis, Leandro Soares Indrusiak, and Neil C Audsley. 2015. Bio-inspired distributed task remapping for multiple video stream decoding on homogeneous NoCs. In *Proceedings of IEEE/ACM/IFIP Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia)*. 1–10.

[129]   Brett H. Meyer, Adam S. Hartman, and Donald E. Thomas. 2010. Cost-effective slack allocation for lifetime improvement in NoC-based MPSoCs. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 1596–1601.

[130]   Mehdi Modarressi, Marjan Asadinia, and Hamid Sarbazi-Azad. 2013. Using task migration to improve non-contiguous processor allocation in NoC-based CMPs. *Journal of Systems Architecture* 59, 7 (2013), 468–481.

[131]   S. Mohanty, V. K. Prasanna, S. Neema, and J. Davis. 2002. Rapid design space exploration of heterogeneous embedded systems using symbolic search and multi-granular simulation. *SIGPLAN Not.* 37, 7 (2002), 18–27.

[132]   Orlando Moreira, Jan-David Mol, Marco Bekooij, and Jef Van Meerbergen. 2005. Multiprocessor resource allocation for hard-real-time streaming with a dynamic job-mix. In *IEEE Real Time and Embedded Technology and Applications Symposium (RTAS)*. 332–341.

[133]   Orlando Moreira, Jacob Jan-David Mol, and Marco Bekooij. 2007. Online resource management in a multiprocessor with a network-on-chip. In *Proceedings of ACM Symposium on Applied Computing (SAC)*. 1557–1564.

[134]   Orlando Moreira, Frederico Valente, and Marco Bekooij. 2007. Scheduling multiple independent hard-real-time jobs on a heterogeneous multiprocessor. In *Proceedings of ACM International Conference on Embedded Software (EMSOFT)*. 57–66.

[135]   Pierre-André Mudry and Gianluca Tempesti. 2009. Self-scaling stream processing: A bio-inspired approach to resource allocation through dynamic task replication. In *IEEE NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. 353–360.

[136]   Srinivasan Murali, Martijn Coenen, Andrei Radulescu, Kees Goossens, and Giovanni De Micheli. 2006. A methodology for mapping multiple use-cases onto networks on chips. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 118–123.

[137]   Srinivasan Murali, Almir Mutapcic, David Atienza, Rajesh Gupta, Stephen Boyd, Luca Benini, and Giovanni De Micheli. 2008. Temperature control of high-performance multi-core platforms using convex optimization. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 110–115.

[138]   Sasmita Kumari Nayak, Sasmita Kumari Padhy, and Siba Prasada Panigrahi. 2012. A novel algorithm for dynamic task scheduling. *Future Generation Computer Systems* 28, 5 (2012), 709–717.

[139]   J. Ng, X. Wang, A. K. Singh, and T. Mak. 2015. DeFrag: Defragmentation for Efficient Runtime Resource Allocation in NoC-Based Many-core Systems. In *Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)*. 345–352.

[140]   T Nishitha and P Chenna Reddy. 2012. Performance Evaluation Of AntHocNet Routing Algorithm in Ad Hoc Networks. In *IEEE International Conference on Computing Sciences (ICCS)*. 207–211.

[141]   Vincent Nollet, Prabhat Avasare, Hendrik Eeckhaut, Diederik Verkest, and Henk Corporaal. 2008. Run-time management of a MPSoC containing FPGA fabric tiles. *IEEE Trans. Very Large Scale Integr. Syst.* (2008), 24–33.

[142]   OpenCL. 2016. Open Computing Language (OpenCL) - The open standard for parallel programming of heterogeneous systems. (2016). https://goo.gl/A9wXRJ.

[143]   Luciano Ost, Marcelo Mandelli, Gabriel Marchesan Almeida, Leandro Moller, Leandro Soares Indrusiak, Gilles Sassatelli, Pascal Benoit, Manfred Glesner, Michel Robert, and Fernando Moraes. 2013. Power-aware dynamic mapping heuristics for NoC-based MPSoCs using a unified model-based approach. *ACM Transactions on Embedded Computing Systems (TECS)* 12, 3 (2013), 75:1–75:22.

[144]   Z. Peter, S. Gilles, U. Nurten, S.J. Nicolas, B. Pascal, and G. Manfred. 2009. A Decentralised Task Mapping Approach for Homogeneous Multiprocessor Network-On-Chips. *International Journal of Reconfigurable Computing* (2009), 3:1–3:14.

[145]   R. Piscitelli and A.D. Pimentel. 2012. Design space pruning through hybrid analysis in system-level design space exploration. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 781–786.

[146]   Ruxandra Pop and Shashi Kumar. 2004. A survey of techniques for mapping and scheduling applications to Network on Chip systems. *School of Engineering, Jönköping University, Research Report* 4 (2004), 4.

[147] Alok Prakash, Siqi Wang, Alexandru Eugen Irimiea, and Tulika Mitra. 2015. Energy-efficient execution of data-parallel applications on heterogeneous mobile platforms. In *Proceedings of IEEE International Conference on Computer Design (ICCD)*. 208–215.

[148] Wei Quan and Andy D. Pimentel. 2015. A Hybrid Task Mapping Algorithm for Heterogeneous MPSoCs. *ACM Trans. Embed. Comput. Syst.* 14, 1 (2015), 14:1–14:25.

[149] Wei Quan and Andy D Pimentel. 2016. A hierarchical run-time adaptive resource allocation framework for large-scale MPSoC systems. *Design Automation for Embedded Systems (DAES)* 20, 4 (2016), 311–339.

[150] Juan Maria Rivas, José Javier Gutiérrez García, José C Palencia Gutiérrez, and Michael González Harbour. 2010. Optimized Deadline Assignment and Schedulability Analysis for Distributed Real-Time Systems with Local EDF Scheduling. In *ESA*. 150–156.

[151] Matthew Rowlings, Andy Tyrrell, and Martin Trefzer. 2015. Social-Insect-Inspired Networking for Autonomous Load Optimisation. *Procedia CIRP* 38 (2015), 259–264.

[152] Siva Satyendra Sahoo, Akash Kumar, and Bharadwaj Veeravalli. 2016. Design and evaluation of reliability-oriented task re-mapping in MPSoCs using time-series analysis of intermittent faults. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 798–803.

[153] Pradip Kumar Sahu and Santanu Chattopadhyay. 2013. A survey on application mapping strategies for network-on-chip design. *Journal of Systems Architecture* 59, 1 (2013), 60–76.

[154] A. Saifullah, K. Agrawal, Chenyang Lu, and C. Gill. 2011. Multi-core Real-Time Scheduling for Generalized Parallel Task Models. In *IEEE Real-Time Systems Symposium (RTSS)*. 217–226.

[155] Manas Saksena and Seongsoo Hong. 1996. An engineering approach to decomposing end-to-end delays on a distributed real-time system. In *IEEE International Workshop on Parallel and Distributed Real-Time Systems*. 244–251.

[156] Samsung. 2014. Samsung Exynos 5422. (2014). www.samsung.com/exynos/.

[157] Lars Schor, Iuliana Bacivarov, Devendra Rai, Hoeseok Yang, Shin-Haeng Kang, and Lothar Thiele. 2012. Scenario-based design flow for mapping streaming applications onto on-chip many-core systems. In *Proceedings of ACM Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES)*. 71–80.

[158] Andreas Schranzhofer, Jian-Jia Chen, and Lothar Thiele. 2009. Power-Aware Mapping of Probabilistic Applications onto Heterogeneous MPSoC Platforms. In *IEEE Real Time and Embedded Technology and Applications Symposium (RTAS)*. 151–160.

[159] A. Schranzhofer, Jian-Jian Chen, and L. Thiele. 2010. Dynamic Power-Aware Mapping of Applications onto Heterogeneous MPSoC Platforms. *IEEE Transactions on Industrial Informatics* 4 (2010), 692 –707.

[160] Nicola Serreli, Giuseppe Lipari, and Enrico Bini. 2009. Deadline assignment for component-based analysis of real-time transactions. In *Workshop on Compositional Real-Time Systems*. Citeseer.

[161] Hafiz Fahad Sheikh and Ishfaq Ahmad. 2014. Efficient heuristics for joint optimization of performance, energy, and temperature in allocating tasks to multi-core processors. In *IEEE International Green Computing Conference (IGCC)*. 1–8.

[162] Hamid Shojaei, AmirHossein Ghamarian, Twan Basten, Marc Geilen, Sander Stuijk, and Rob Hoes. 2009. A parameterized compositional multi-dimensional multiple-choice knapsack heuristic for CMP run-time management. In *Proceedings of ACM Design Automation Conference (DAC)*. 917–922.

[163] Amit Kumar Singh, Anup Das, and Akash Kumar. 2013. Energy Optimization by Exploiting Execution Slacks in Streaming Applications on Multiprocessor Systems. In *Proceedings of ACM Design Automation Conference (DAC)*. Article 115, 7 pages.

[164] Amit Kumar Singh, Anup Das, and Akash Kumar. 2013. RAPIDITAS: RAPId Design-Space-Exploration Incorporating Trace-Based Analysis and Simulation. In *Proceedings of IEEE Euromicro Conference on Digital System Design (DSD)*. 836–843.

[165] Amit Kumar Singh, Piotr Dziurzanski, and Leandro Soares Indrusiak. 2015. Market-inspired Dynamic Resource Allocation in Many-core High Performance Computing Systems. In *IEEE International Conference on High Performance Computing & Simulation (HPCS)*. 413–420.

[166] Amit Kumar Singh, Piotr Dziurzanski, and Leandro Soares Indrusiak. 2015. Value and Energy Optimizing Dynamic Resource Allocation in Many-core HPC Systems. In *IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. 180–185.

[167] Amit Kumar Singh, Piotr Dziurzanski, and Leandro Soares Indrusiak. 2016. Value and Energy Aware Adaptive Resource Allocation of Soft Real-time Jobs on Many-core HPC Data Centers. In *IEEE International Symposium on Real-Time Computing (ISORC)*.

[168] Amit Kumar Singh, Akash Kumar, and Thambipillai Srikanthan. 2013. Accelerating throughput-aware runtime mapping for heterogeneous MPSoCs. *ACM Trans. Des. Autom. Electron. Syst.*, Article 9 (2013), 29 pages.

[169] Amit Kumar Singh, Muhammad Shafique, Akash Kumar, and Jörg Henkel. 2013. Mapping on Multi/Many-core Systems: Survey of Current and Emerging Trends. In *Proceedings of ACM Design Automation Conference (DAC)*. Article 1, 10 pages.

[170] A. K. Singh, M. Shafique, A. Kumar, and J. Henkel. 2016. Analysis and Mapping for Thermal and Energy Efficiency of 3-D Video Processing on 3-D Multicore Processors. *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)* 24, 8 (2016), 2745–2758.

[171] Amit Kumar Singh, Muhammad Shafique, Akash Kumar, and Jorg Henkel. 2016. Resource and Throughput Aware Execution Trace Analysis for Efficient Run-Time Mapping on MPSoCs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 35, 1 (2016), 72–85.

[172] L.T. Smit, G.J.M. Smit, J.L. Hurink, H. Broersma, D. Paulusma, and P.T. Wolkotte. 2004. Run-time mapping of applications to a heterogeneous reconfigurable tiled system on chip architecture. In *Proceedings of IEEE International Conference on Field-Programmable Technology (FPT)*. 421–424.

[173] John A Stankovic, Chenyang Lu, Sang H Son, and Gang Tao. 1999. The case for feedback control real-time scheduling. In *Proceedings of IEEE Euromicro Conference on Real-Time Systems (ECRTS)*. 11–20.

[174] S. Stuijk, M.C.W. Geilen, and T. Basten. 2006. SDF^3: SDF For Free. In *Proceedings of IEEE Conference on Application of Concurrency to System Design (ACSD)*. 276–278.

[175] Sander Stuijk, Marc Geilen, and Twan Basten. 2010. A Predictable Multiprocessor Design Flow for Streaming Applications with Dynamic Behaviour. In *Proceedings of IEEE Euromicro Conference on Digital System Design (DSD)*. 548–555.

[176] Guang Sun, Yong Li, Yuanyuan Zhang, Li Su, Depeng Jin, and Lieguang Zeng. 2010. Energy-aware run-time mapping for homogeneous NoC. In *IEEE International Symposium on System on Chip (SoC)*. 8–11.

[177] Timon D. ter Braak, Philip K. F. Hölzenspies, Jan Kuper, Johann L. Hurink, and Gerard J. M. Smit. 2010. Run-time spatial resource management for real-time applications on heterogeneous MPSoCs. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 357–362.

[178] Theocharis Theocharides, Maria K. Michael, Marios Polycarpou, and Ajit Dingankar. 2009. Towards embedded runtime system level optimization for MPSoCs: on-chip task allocation. In *Proceedings of ACM Great Lakes symposium on VLSI (GLSVLSI)*. 121–124.

[179] Theocharis Theocharides, Maria K. Michael, Marios Polycarpou, and Ajit Dingankar. 2010. Hardware-enabled Dynamic Resource Allocation for Manycore Systems Using Bidding-based System Feedback. *EURASIP J. Embedded Syst.* 2010, Article 3 (2010), 21 pages.

[180] Lothar Thiele, Lars Schor, Hoeseok Yang, and Iuliana Bacivarov. 2011. Thermal-aware system analysis and software synthesis for embedded multi-processors. In *Proceedings of ACM Design Automation Conference (DAC)*. 268–273.

[181] TILE-Gx. 2009. First 100-core Processor with the New TILE-Gx Family. (2009). http://www.tilera.com/ (Last visited: 12 February, 2016).

[182] Stavros Tzilis, Ioannis Sourdis, Vasileios Vasilikos, Dimitrios Rodopoulos, and Dimitrios Soudris. 2016. Runtime Management of Adaptive MPSoCs for Graceful Degradation. In *Proceedings of ACM Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES)*. 1–10.

[183] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Borkar. 2007. An 80-Tile 1.28TFLOPS Network-on-Chip in 65nm CMOS. In *Proceedings of IEEE International Solid-State Circuits Conference (ISSCC)*. 98–589.

[184] Feng Wang, Yibo Chen, C. Nicopoulos, X. Wu, Y. Xie, and N. Vijaykrishnan. 2011. Variation-Aware Task and Communication Mapping for MPSoC Architecture. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 2 (2011), 295 –307.

[185] Shengquan Wang and Jian-Jia Chen. 2010. Thermal-aware lifetime reliability in multicore systems. In *Proceedings of International Symposium on Quality Electronic Design (ISQED)*. 399–405.

[186] Xiaohang Wang, Baoxin Zhao, Terrence Mak, Mei Yang, Yingtao Jiang, and Masoud Daneshtalab. 2015. An efficient runtime power allocation scheme for many-core systems inspired from auction theory. *Integration, the VLSI Journal* 50 (2015), 147–157.

[187] Andreas Weichslgartner, Deepak Gangadharan, Stefan Wildermann, Michael Glaß, and Jürgen Teich. 2014. DAARM: Design-time Application Analysis and Run-time Mapping for Predictable Execution in Many-core Systems. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 34:1–34:10.

[188] Stefan Wildermann, Felix Reimann, Daniel Ziener, and Jürgen Teich. 2011. Symbolic design space exploration for multi-mode reconfigurable systems. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 129–138.

[189] Dong Hyuk Woo and Hsien-Hsin S Lee. 2008. Extending Amdahl's law for energy-efficient computing in the many-core era. *Computer* 41, 12 (2008).

[190] Frédéric Worm, Paolo Ienne, Patrick Thiran, and Giovanni De Micheli. 2002. An adaptive low-power transmission scheme for on-chip networks. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 92–100.

[191] Xilinx. 2016. Zynq UltraScale+ EG MPSoC. (2016). https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html.

[192] Liping Xue, Ozcan ozturk, Feihui Li, Mahmut Kandemir, and I. Kolcu. 2006. Dynamic partitioning of processing and memory resources in embedded MPSoC architectures. In *Proceedings of IEEE Conference on Design, Automation and Test in Europe (DATE)*. 690–695.

[193] Peng Yang, Paul Marchal, Chun Wong, Stefaan Himpe, Francky Catthoor, Patrick David, Johan Vounckx, and Rudy Lauwereins. 2002. Managing dynamic concurrent tasks in embedded real-time multimedia systems. In *Proceedings of IEEE/ACM/IFIP Conference on Hardware/Software Codesign and System Synthesis (ISSS+CODES)*. 112–119.

[194] Chee Shin Yeo and Rajkumar Buyya. 2006. A Taxonomy of Market-based Resource Management Systems for Utility-driven Cluster Computing. *Softw. Pract. Exper.* 36, 13 (Nov. 2006), 1381–1419.

[195] C. Ykman-Couvreur, P. Avasare, G. Mariani, G. Palermo, C. Silvano, and V. Zaccaria. 2011. Linking run-time resource management of embedded multi-core platforms with automated design-time exploration. *IET Computers Digital Techniques* 2 (2011), 123 –135.

[196] Nicholas H. Zamora, Xiaoping Hu, and Radu Marculescu. 2007. System-level performance/power analysis for platform-based design of multimedia applications. *ACM Trans. Des. Autom. Electron. Syst.*, Article 2 (2007), 29 pages. Issue 1.

[197] Luyuan Zeng, Pengcheng Huang, and Lothar Thiele. 2016. Towards the design of fault-tolerant mixed-criticality systems on multicores. In *Proceedings of ACM Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES)*. 6:1–6:10.

[198] Xiuyi Zhou, Jun Yang, Yi Xu, Youtao Zhang, and Jianhua Zhao. 2010. Thermal-Aware Task Scheduling for 3D Multicore Processors. *IEEE Trans. Parallel Distrib. Syst.* (2010), 60–71.

[199] Yifan Zhu and Frank Mueller. 2005. Feedback EDF scheduling of real-time tasks exploiting dynamic voltage scaling. *Real-Time Systems* 31, 1-3 (2005), 33–63.

[200] Sergey Zhuravlev, Juan Carlos Saez, Sergey Blagodurov, Alexandra Fedorova, and Manuel Prieto. 2012. Survey of Scheduling Techniques for Addressing Shared Resources in Multicore Processors. *ACM Comput. Surv.* 45, 1, Article 4 (2012), 28 pages.