



Universiteit
Leiden
The Netherlands

Latency, energy, and schedulability of real-time embedded systems

Liu, D.; Liu D.

Citation

Liu, D. (2017, September 6). *Latency, energy, and schedulability of real-time embedded systems*. Retrieved from <https://hdl.handle.net/1887/54951>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/54951>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/54951> holds various files of this Leiden University dissertation

Author: Liu, D.

Title: Latency, energy, and schedulability of real-time embedded systems

Issue Date: 2017-09-06

Chapter 5

Energy Optimization for Real-Time Tasks

Di Liu, Jelena Spasic, Peng Wang, and Todor Stefanov,

"Energy-Efficient Scheduling of Real-Time Tasks on Heterogeneous Multicores Using Task Splitting",

"The 22nd IEEE International Conference on Embedded and Real-Time Computing Systems and

Applications", Daegu, South Korea, 2016, pp. 1-10.

As we discussed in Section 2.2.3, the semi-partitioned scheduling/task-splitting, e.g., [BDWZ12, GSYY10, JLBK13, BBA11], can achieve a good trade-off between global scheduling and partitioned scheduling in terms of schedulability, resource utilization and scheduling overhead. The advantage of the task-splitting technique has been extended to another dimension, namely, energy-efficient real-time scheduling. Lu and Guo [LG11] investigated to use the task-splitting technique given in [GSYY10] to energy-efficiently schedule real-time tasks under fixed-priority scheduling on homogeneous multicore systems. As presented in Section 1.1, the heterogeneous multicore systems are gradually replacing the homogeneous multicore systems in order to satisfy diverse performance requirements of different applications and at the same time reduce the energy consumption. However, there is no work investigating the task-splitting approach on heterogeneous multicore systems for the energy-efficient purpose. Motivated by this fact, in this chapter, we investigate how to adopt the task-splitting approach with dynamic priority scheduling to better utilize the resources on heterogeneous multicore systems for energy efficiency. We select the C=D approach [BDWZ12], which will be introduced in details later in Section 5.2.4, to split tasks among heterogeneous cores. We extend the C=D approach for heterogeneous multicore systems and propose an allocation algorithm to schedule real-time tasks with C=D task-splitting on heterogeneous multicore systems. Formally, our novel technical contributions are summarized as follows:

- We analyze the properties of the C=D task-splitting and extend it for heterogeneous multicore systems. We present a new definition, namely ‘*valid split*’, for the C=D task-splitting on heterogeneous multicore systems. This analysis is presented in Section 5.4;
- Based on the analysis of the C=D task-splitting and the characteristics of heterogeneous multicore systems, we propose an energy-efficient algorithm, called ASHM, to allocate and split real-time tasks on heterogeneous multicore systems. Algorithm ASHM is presented in Section 5.5;
- Since the existing methods to compute the minimum operational frequency for each core cannot work with the C=D approach, we propose a new approach based on Quick convergence Processor-demand Analysis (QPA) [ZB09] to compute the minimum frequency for each core in the system. This proposed approach is presented in Section 5.5.4

The extensive experiments on synthetic real-time tasks shows the effectiveness of our ASHM algorithm over the existing partitioned algorithms in terms of energy efficiency.

5.1 Related Work

Energy-efficient scheduling for real-time systems has been widely explored in the past two decades. Chen and Kuo in [CK07] comprehensively reviewed most of the papers addressing energy-efficient real-time scheduling problems before 2007. An updated review for energy-efficient real-time scheduling is provided by Bambagini et al. in [BMAB16]. We can see from [CK07, BMAB16] that most of the works consider homogeneous systems, whereas in this work we consider heterogeneous multicore systems which are more energy-efficient but more difficult to effectively schedule the tasks.

A few works consider heterogeneous systems. Chen and Thiele in [CT08] proposed a polynomial algorithm to energy efficiently schedule periodic tasks on heterogeneous systems but the systems they considered had only two cores. In contrast, we consider a more general system model where the system has two types of cores, and for each core type we can have any number of cores which can be seen on many real commercial processors. Chen *et al.* [CST09] developed two polynomial-time algorithms to energy efficiently allocate real-time tasks on a more general system model that can have different types of processors and different number of processors for each type like we consider in our work. However, in their work, they do not take voltage/frequency scaling (VFS) into account, whereas we consider VFS as a crucial

technique to improve the energy efficiency. With the consideration of VFS, we can further minimize the energy consumption of the heterogeneous system. In [HTC07], Huang *et al.* proposed an allocation algorithm to schedule frame-based real-time tasks on heterogeneous multicore systems, where a non-preemptive scheduling is considered. The main difference compared to our work is: (1) they consider frame-based real-time task model, whereas we consider the periodic task model which is more general; (2) the non-preemptive scheduling, they consider, is known to be NP-hard in the strong sense even on uniprocessor [JSM91]. In contrast, we consider preemptive scheduling.

Recently, more interests have risen for energy efficient real-time scheduling on single-ISA heterogeneous multicore systems. Liu *et al.* [LSCS15] consider an optimal cluster scheduling to schedule real-time tasks on cluster heterogeneous multicore systems. However, from practical perspective the optimal cluster scheduling suffers from a very high overhead caused by frequent context switching and task migration. When the overhead is taken into account, the achieved resource utilization may be quite low in practice [BBA11]. In contrast, the C=D task-splitting, we consider, has a limited number of migrations and on each core a normal EDF scheduler is used to schedule real-time tasks, hence it significantly reduces the context-switching and task migration overhead and makes it more practical for real implementation. Colin *et al.* [CKR14] and Elewi *et al.* [ESAS14] adopt the partitioned EDF scheduling to schedule real-time tasks on heterogeneous multicore systems, where both consider energy minimization as the objective. Due to the capacity loss of partitioned scheduling, the proposed approaches from [CKR14] and [ESAS14] do not fully utilize ‘LITTLE’ cores on a heterogeneous multicore system and thus possibly lose some opportunities to further reduce the energy consumption. Contrarily, in our work, we adopt the state-of-the-art C=D task-splitting approach to exploit the energy efficiency of a heterogeneous multicore system. Our experimental results on randomly generated task sets demonstrate the merit of the task-splitting on heterogeneous multicore systems.

A few works study the task migration/splitting approaches for energy-efficient real-time multicore system. Chen *et al.* [CHC⁺04] address the energy-efficient scheduling problem on homogeneous multicore systems with task migration, in which all tasks have the same release time and a common deadline. In our work, we consider a more general and widely-used periodic task model and instead of homogeneous multicore systems, we consider heterogeneous multicore systems which are more energy efficient. Lu and Guo [LG11] adopt the task-splitting approach proposed by Guan *et al.* [GSYY10] on homogeneous multicore systems to achieve energy efficiency. The main difference between [LG11] and our work is twofold: 1) they consider fixed priority scheduling, whereas dynamic priority scheduling, i.e., earliest deadline first (EDF) [LL73], is adopted in our work. It is known that EDF can achieve better re-

source utilization than fixed-priority scheduling; 2) they consider homogeneous multicore systems, whereas we target heterogeneous multicore systems which are more energy-efficient.

5.2 Background

In this section, we present the system model, task model, and energy model used in this work. Then, we give a brief description of the C=D task-splitting approach [BDWZ12].

5.2.1 System Model

We consider a heterogeneous multicore system M which consists of two types of cores, the ‘big’ core for performance and the ‘LITTLE’ core for low power. Throughout this chapter, we use PE and EE to denote a ‘big’ core and a ‘LITTLE’ core, respectively, like what we did in Chapter 4. We use M_{EE} and M_{PE} to denote the sets consisting of all EE cores and all PE cores, respectively. The power consumption of one core can be computed by the following equation,

$$P(f) = \alpha f^b + s \quad (5.1)$$

where α and $b \in [2, 3]$ are technology-based parameters [CK07], f is the operational frequency. For different types of cores, α and b are different. The first term of Equation (5.1) is the frequency-related power consumption, i.e., the dynamic power consumption. s denotes the power consumption unrelated to the frequency, i.e., the static power consumption. Each core executes independently from the others and has a discrete frequency set at which the core can run. Let $\bar{f}_j = \{f_1, \dots, f_l\}$ denote the frequency set of core j . Without loss of generality, we assume that the frequencies in the set are sorted in increasing order, i.e., $f_k < f_{k+1}$.

5.2.2 Task Model

The task model adopted in this work is similar to the one introduced in Section 2.2.1, but all tasks are assumed to start at time instant 0, i.e., $S_1 = S_2 = \dots = S_n = 0$. Moreover, since we have two type of cores, the WCET of each task may vary when executing on different types of cores. We slightly extend the model to have two WCETs as we did in Chapter 4.

- C_i^{EE} and C_i^{PE} are the worst-case execution times (WCETs) of task τ_i executing on an EE core and PE core at the maximum frequency, respectively;

Then, a task is characterized by a tuple of parameters $\tau_i = \{C_i^{EE}, C_i^{PE}, D_i, T_i\}$.

5.2.3 Energy Model

With the system and task models discussed above, we explain how to compute the energy consumption for the system. After all tasks are allocated to cores, the energy consumption for each core can be computed as follows:

$$E_j = hp \left(\alpha_j f_j^{b_j} \frac{f_{max}}{f_j} \sum_{\forall \tau_i \in \Gamma_j} \frac{C_i}{T_i} + s_j \right) \quad (5.2)$$

where Γ_j is the task set containing all tasks allocated to core j and hp is the hyper-period of task set Γ_j . The hyper-period is the least common multiple (*lcm*) of all tasks' periods. Every hyper-period has the same workload and thus we compute the energy consumption within one hyper-period. The energy consumption of the whole system is the summation of the energy consumption E_j of all cores.

5.2.4 C=D Task-Splitting

In this work, we adopt the C=D task-splitting to schedule real-time tasks on a heterogeneous multicore system. Burns *et al.* in [BDWZ12] propose the C=D approach to split real-time tasks on homogeneous systems. They use a preemptive earliest deadline first (EDF) scheduling [LL73] to schedule the tasks on each core. The tasks are first allocated to cores according to a certain allocation algorithm. If task τ_i cannot be integrally allocated to a core, the C=D approach splits unassigned task τ_i into two parts/subtasks, τ_i^1 and τ_i^2 . The split procedure is as follows:

- Find a processor x and then compute the maximum computation time C_i^1 for subtask τ_i^1 which ensures the schedulability of subtask τ_i^1 on processor x . For subtask τ_i^1 , its deadline D_i^1 is set to be equal to C_i^1 and its period T_i^1 is equivalent to its original period T_i , i.e., $\tau_i^1 = \{C_i^1, D_i^1 = C_i^1, T_i^1 = T_i\}$. Then, subtask τ_i^1 is allocated to processor x ;
- According to subtask τ_i^1 , we can obtain the second subtask τ_i^2 . The WCET C_i^2 of τ_i^2 is computed as $C_i^2 = C_i - C_i^1$, its deadline D_i^2 is computed as $D_i^2 = D_i - D_i^1$ and its period T_i^2 equals to its original period, $T_i^2 = T_i$, i.e., $\tau_i^2 = \{C_i^2 = C_i - C_i^1, D_i^2 = D_i - D_i^1, T_i^2 = T_i\}$. Subtask τ_i^2 is allocated to a processor which has enough space to schedule subtask τ_i^2 and is different from processor x on which subtask τ_i^1 is allocated.

In the remainder of this chapter, we call subtask τ_i^1 the first subtask and subtask τ_i^2 the second subtask.

The C=D task-splitting permits each core to have only one first subtask τ_i^1 . This means that the whole system has at most M split tasks, where M is the number of

cores. This task-splitting scheme can be realized by using task migration. τ_i^1 completes its execution on the allocated core. Then it migrates to the core where τ_i^2 is assigned and continues the execution of subtask τ_i^2 . From the experimental results in [JLBK13], the C=D task-splitting outperforms other existing semi-partition/task-splitting approaches in terms of schedulability.

Migration Overhead: Like [BDWZ12], in our work the migration overhead is assumed to be negligible. An extensive number of experiments on real hardware systems [BBA11] have shown that with cache coherence among cores the task migration overhead is at the similar order of magnitude as the normal context switching. The cache coherence hardware architecture, like CoreLink CCI-400 Cache Coherent Interconnect [ARM16], has been adopted by the big.LITTLE multicore systems to maintain the cache coherence between cores. Therefore, the migration overhead is accounted for in the WCET of a task.

5.3 Motivational Example

In this section, we use an example to motivate the application of the C=D task-splitting approach on heterogeneous multicore systems for energy efficiency purpose. For simplicity, assume that we have a multicore system with one PE core and one EE core. The PE core and EE core have different power parameters α and b (see Equation (5.1) and (5.2)). In this example, we use the parameter values from Table 4.3 given in Section 4.2.2 of Chapter 4. We recap it here for reference convenience.

Core type	$\alpha(W/Mhz^b)$	b	$s(W)$
PE	3.03×10^{-9}	2.621	0.155
EE	2.62×10^{-9}	2.12	0.027

Table 5.1: Power parameters for different core types

	$C^{PE}(ms)$	$C^{EE}(ms)$	$D(ms)$	$T(ms)$
τ_1	55	110	100	100
τ_2	20	40	100	100
τ_3	20	40	100	100
τ_4	15	30	100	100

Table 5.2: The original task set

	$C^{PE}(ms)$	$C^{EE}(ms)$	$D(ms)$	$T(ms)$
τ_4^1	10	20	20	100
τ_4^2	5	10	80	100

Table 5.3: Split subtasks

Suppose to have four tasks with the parameters given in Table 5.2. As far as the deadlines can be ensured, we strive to partition/allocate as many tasks as possible to the EE core in order to save energy consumption. However, since scheduling τ_1 on the EE core will violate the deadline guarantee, only τ_2 , τ_3 , and τ_4 are eligible to be scheduled on the EE core. But we cannot schedule τ_2 , τ_3 , and τ_4 together on the EE core, because a total utilization of $1.1 > 1$ leads to infeasibility. One task has to be scheduled on the PE core along with τ_1 . Then, we obtain a fully partitioned allocation for the given task set, where τ_1 and τ_4 are scheduled on the PE core and τ_2 and τ_3 are scheduled on the EE core. In contrast to the above fully partitioned allocation, we adopt the C=D task-splitting (explained in Section 5.2.4) to schedule the tasks on the multicore system. In the splitting case, τ_1 is scheduled on the PE core while τ_2 and τ_3 are scheduled on the EE core. But τ_4 is split into two subtasks, τ_4^1 and τ_4^2 , and then we schedule τ_4^1 on the EE core and τ_4^2 on the PE core. The parameters for the subtasks are shown in Table 5.3,

With the given allocation and the power parameters, we can compute a minimum frequency for each core such that the energy consumption can be minimized by using VFS while deadlines are still ensured. Table 5.4 shows the allocation, the minimum operational frequency of each core, and the energy consumption of the multicore system. We can see that the splitting approach saves energy consumption by 32% compared to the partitioned approach because it can effectively utilize the EE core to save energy and at the same time it can reduce the workload allocated to the PE core. As a result, the PE core in the splitting approach executes at a lower frequency compared to the partitioned approach.

Mapping	PE	EE	f^{PE}	f^{EE}	Energy(mJ)
Partitioned	τ_1, τ_4	τ_2, τ_3	1.4GHz	1.2GHz	5.42
Splitting	τ_1, τ_4^2	τ_2, τ_3, τ_4^1	1.2GHz	1.4GHz	3.69

Table 5.4: Energy consumption

From the example, we see the advantage of the C=D task-splitting approach on heterogeneous systems in terms of energy efficiency. In the subsequent sections, we will introduce our novel approach to exploit the C=D task-splitting on heterogeneous

multicore systems for minimizing the energy consumption.

5.4 C=D Task-Splitting on Heterogeneous Multiprocessor Systems

In [BDWZ12], the C=D task-splitting is devised for homogeneous multiprocessor systems. However, in our work, we target heterogeneous multicore systems [Mit15][Mit16] which have been emerging as an alternative of the conventional homogeneous multicore systems. In this section, we investigate how to adopt the C=D task-splitting on a heterogeneous system.

5.4.1 Task Splitting

	C^{PE}	C^{EE}	D	T
τ_1	60	120	100	100
τ_1^1	25	50	50	100
τ_1^2	35	70	50	100
τ_1^1	41	82	82	100
τ_1^2	19	38	18	100

Table 5.5: Let us assume that we split τ_1 into two subtasks τ_1^1 and τ_1^2 and allocate τ_1^1 and τ_1^2 to an EE core and a PE core, respectively. We assume that there is no constraint on the split. We give two different splits for τ_1 shown in rows 3,4 and 5,6. For the first split shown in rows 3,4, there is no problem to schedule the subtasks. However, for the second split, although the execution time on the EE core is maximized, it causes a deadline miss for subtask τ_1^2 due to $C^{PE} > D$, seen in the last row with red color.

Since, on heterogeneous multicore systems, a task's WCET is varying upon the allocated core, the splitting on the heterogeneous multicore system should pay more attention to the varying WCET and the relation between the obtained two subtasks. First, the deadline of the first subtask τ_i^1 is set according to where the first subtask is allocated. For instance, assume that a subtask τ_i^1 has its $C_i^{PE} = 5$ and $C_i^{EE} = 10$. If it is allocated to a PE core, its deadline D_i^1 equals to $C_i^{PE} = 5$, otherwise $D_i^1 = C_i^{EE} = 10$ if allocated to an EE core. Moreover, in some cases an improper split might cause a deadline miss for the second subtask τ_i^2 . The example given in Table 5.5 demonstrates this issue.

From the example, we observe the potential split issue on a heterogeneous multicore system. Thus, we give the following property to ensure that a proper split on

heterogeneous multicore systems is obtained:

Property 1. *On a heterogeneous multicore system, the following inequality must hold for a split task τ_i ,*

$$T_i - C_i^1 \geq C_i^2 \quad (5.3)$$

where C_i^1 and C_i^2 are the WCETs of subtasks τ_i^1 and τ_i^2 , depending on which type of core the subtasks have been allocated.

This property is to ensure enough space to execute the second subtask τ_i^2 on a heterogeneous system. We can see that for subtask τ_i^2 it must have,

$$D_i^2 \geq C_i^2 \quad (5.4)$$

Since $D_i^2 = D_i - D_i^1 = T_i - D_i^1$ and $D_i^1 = C_i^1$, see Section 5.2.4, we obtain

$$T_i - C_i^1 \geq C_i^2 \quad (5.5)$$

Thus, the property is observed. Based on this property, we give the following definition,

Definition 5.4.1 (valid split). If two subtasks τ_i^1 and τ_i^2 obtained by splitting task τ_i satisfy Property 1, we call such split a *valid split*.

If the split is not a *valid split*, then the second subtask cannot meet its deadline.

5.4.2 Subtask Allocation

In Section 5.4.1, we discussed how to find a *valid split* for a task on a heterogeneous multicore system. Here, we continue to discuss the allocation of subtasks. Before proceeding to the discussion, we distinguish tasks in two categories and give their definitions as follows,

Definition 5.4.2. If a task can be integrally scheduled on an EE core, we call such task an **eligible task** (E-task).

Definition 5.4.3. If a task **cannot** be integrally scheduled on an EE core, we call such task a **non-eligible task** (NE-task).

If we look at the motivational example in Section 5.3-Table 5.2, τ_2 , τ_3 , and τ_4 are E-tasks and τ_1 is NE-task. Now, we discuss the possible allocation destinations for these two categories of tasks.

E-task

When an E-task is selected to be split, any split is a *valid split* regardless of which type of core the subtasks are allocated. Therefore, for an E-task, the two subtasks can be allocated to any type of core, as long as the schedulability of the system is ensured. Thus, we can have three possible combinations to allocate the two subtasks of an E-task:

- Allocate the two subtasks to two EE cores;
- Allocate the two subtasks to one EE core and one PE core; and
- Allocate the two subtasks to two PE cores.

NE-task

When a NE-task is about to be split, we need to ensure that the obtained split is a valid split by satisfying Property 1. For a NE-task, we cannot allocate the two subtasks to two EE cores, because Property 1 will be violated and then it leads to an invalid split. Excluding the invalid combination, we have two possible combinations to allocate the two subtasks of a NE-task:

- Allocate the two subtasks to one EE core and one PE core; and
- Allocate the two subtasks to two PE cores.

With the above possible allocation destinations for the two categories of tasks, in the next section, we will use this information to devise an energy-efficient allocation strategy for each category of tasks.

5.5 Allocation and Split on Heterogeneous Multicore Systems (ASHM)

In [CT08], Chen and Thiele have shown that allocating real-time tasks onto two different processors is an NP-hard problem. Their problem is just a subset of our problem, so our problem is also an NP-hard problem. Hence, we propose a heuristic algorithm to energy-efficiently schedule real-time tasks on heterogeneous multicore systems with task-splitting. We call this algorithm ASHM. ASHM first handles all E-tasks and then all NE-tasks. For the sake of clarity, we first explain the different parts in the ASHM algorithm and after that we explain the whole ASHM algorithm. Before proceeding to the detailed discussion, we introduce the following property for the core with first subtask τ_i^1 allocated on it,

Property 2. *A core must run at the maximum frequency if first subtask τ_i^1 of a split task τ_i is assigned to it.*

It is trivial to see this property because the first subtask of a split task has its WCET equal to the deadline. Scaling down the frequency leads to a deadline miss. This property is useful to determine the allocation of the subtasks.

5.5.1 Allocation and splitting of E-tasks

ASHM first starts to allocate and split E-tasks. The procedure to allocate and split E-tasks is summarized as follows:

1. Use a bin-packing algorithm, first-fit-decreasing (FFD) [CGJ97], to integrally allocate E-tasks to EE cores;
2. Split unallocated E-tasks on the platform. For a given unallocated E-task τ_i , we use the following allocation and splitting order,
 - (a) Split τ_i among two EE cores. If it fails, try step b);
 - (b) Split τ_i among one EE core and one PE core. If fails, try step c);
 - (c) Allocate τ_i integrally to one PE core. If it fails, try step d);
 - (d) Split τ_i among two PE cores. If it fails, the system is unschedulable on the platform with $M = \{M_{EE}, M_{PE}\}$.

For the first step, we use FFD to integrally allocate EE tasks to EE cores because FFD is proven to be the resource efficient bin-packing algorithm [AY03]. By using FFD we could leave some EE cores with a lot of free capacity. This could later benefit the NE-tasks for energy saving.

After some E-tasks are integrally allocated to EE cores, we might have some E-tasks left unallocated. The next step is to split and allocate them on the system. The allocation and split order summarized above prioritizes the EE cores to explore the energy-efficient potential on the EE cores. Therefore, we first try to allocate the subtasks of a split E-task to two EE cores. If the task cannot be split among two EE cores, this means that there is no enough space on EE cores. So, we try one EE core and one PE core. Since, a PE core consumes much more power than an EE core and Property 2 indicates the maximum frequency requirement, it is not favorable to allocate the first subtask to a PE core. Therefore, we constrain ourself to allocate the first subtask to an EE core and the second subtask to a PE core. For the selection of the PE cores, we use the approach proposed in [CKR14] which selects the core with the smallest energy cost contribution to the whole system when the task is allocated to it. If the

combination of one EE core and one PE core still fails, we need to find an allocation among PE cores.

On PE cores, we first try to integrally allocate the E-task to one PE core because if we split an E-task among two PE cores, Property 2 requires that one PE core must execute at the maximum frequency which leads to a very high power consumption. Hence, we prefer to integrally allocate the E-task to one PE core than split it among two PE cores. We also use the approach from [CKR14] to select the energy-efficient core for the task. If it still fails, we try the final step to split it on two PE cores in order to ensure its schedulability.

Algorithm 6 presents the pseudo-code to allocate and split E-tasks, called EAS, following the procedure explained above. EAS takes as inputs task set Γ_E consisting of all E-tasks and the heterogeneous multicore platform consisting of EE core set M_{EE} and PE core set M_{PE} and outputs the allocation of all E-tasks. At Line 1, we first use FFD to allocate E-tasks to EE cores integrally. If there are some unallocated E-tasks, we follow the steps introduced above to split unallocated E-tasks among two EE cores or one EE core and one PE core - see Line 3-10. We use function $\text{mpwr}()$ to represent the core selection approach from [CKR14], where the inputs of $\text{mpwr}()$ are a core set and a task and the output is a core which can schedule the task and has the smallest contribution to the energy consumption. However, if the task is not allocated successfully, we have to try to allocate or split the task among PE cores - see Line 11-24. From Line 12-14, the integral allocation on one PE core is first tried. If it fails, from Line 15-24 EAS splits τ_i among two PE cores. Function SPLIT in Algorithm 6 finds the first subtask τ_i^1 with the maximum WCET which is schedulable on core x and also gives the corresponding τ_i^2 . We will explain SPLIT in details later in Section 5.5.3.

5.5.2 Allocation and Splitting of NE-tasks

After all E-tasks are allocated, we proceed towards allocating and splitting NE-tasks on the system. The procedure to allocate and split NE-task τ_i is summarized as follows:

1. Split τ_i among one EE core and one PE core. If it fails, try step 2);
2. Allocate τ_i integrally onto one PE core. If it fails, try step 3);
3. Split τ_i among two PE cores. If it fails, it is unschedulable.

Since, after the allocation of E-tasks, EE cores might have some free space to execute parts of NE-tasks, we first try to split a NE-task among one EE core and one PE core in order to utilize EE cores for energy saving. Since the first subtask needs a

Algorithm 6: E-task Allocation and Split (EAS)

Input: All E-tasks Γ_E and the heterogeneous multicore platform $M = \{M_{EE}, M_{PE}\}$

Output: Allocation for all E-tasks

```

1   $M_{EE} \leftarrow$  using FFD to allocate tasks from  $\Gamma_E$ 
2   $\Gamma_{un} \leftarrow$  unallocated tasks from  $\Gamma_E$ 
3  for  $\forall \tau_i \in \Gamma_{un}$  in order of decreasing  $U$  do
4      for  $\forall x \in M_{EE}$  in order of increasing  $U$  do
5           $\tau_i^1, \tau_i^2 = \text{SPLIT}(\tau_i, x)$ 
6          if  $\tau_i^1 \neq \emptyset$  then
7               $x \leftarrow \tau_i^1$ 
8               $y \leftarrow \text{mpwr}(M = \{M_{EE}, M_{PE}\}, \tau_i^2)$ 
9              if  $y = \emptyset$  then
10                  $x \leftarrow x - \tau_i^1$ 
11         if  $\tau_i$  is not allocated successfully then
12              $x \leftarrow \text{mpwr}(M_{PE}, \tau_i)$ 
13             if  $x \neq \emptyset$  then
14                  $x \leftarrow \tau_i$ 
15             else
16                 for  $\forall x \in M_{PE}$  in order of decreasing  $U$  do
17                      $\tau_i^1, \tau_i^2 = \text{SPLIT}(\tau_i, x)$ 
18                     if  $\tau_i^1 \neq \emptyset$  then
19                          $x \leftarrow \tau_i^1$ 
20                          $y \leftarrow \text{mpwr}(M_{PE}, \tau_i^2)$ 
21                         if  $y = \emptyset$  then
22                              $x \leftarrow x - \tau_i^1$ 
23                         else
24                              $y \leftarrow \tau_i^2$ 
25         if  $\tau_i$  is not allocated successfully then
26             return Unscheduleable
27 return Allocation of  $\forall \tau_i \in \Gamma_E$ 
    
```

maximum operational frequency (Property 2), we constrain the first subtask to the EE core and allocate the second subtask to a PE core for ensuring the schedulability. However, when we maximize the execution time of the first subtask on an EE core, it might bring a negative effect on the second subtask. Maximizing the execution of the first subtask will reduce the slack time for the second subtask, i.e., $D_i^2 - C_i^2$. As a consequence, the reduced slack time leaves a little space to scale down the frequency of the PE core which might compromise the energy saving from the EE core. Hence, in order to provide an energy-efficient split, we set the following constraint for splitting a NE-task on one EE core and one PE core.

$$\frac{C_i^2}{D_i^2} \leq \frac{C_i}{T_i} \quad (5.6)$$

Constraint (5.6) can guarantee that after the split the slack ratio of the second subtask is not smaller than before. Therefore, it would not require to run at a higher frequency. If the task cannot be split on one EE core and one PE core, we integrally allocate NE-task τ_i to one PE core. For the integral allocation, we try to allocate task τ_i to the PE core given by function `mpwr()`. If task τ_i cannot be allocated to a PE core, then we split it among two PE cores in order to ensure its schedulability.

Algorithm 7 presents the pseudo-code to allocate and split NE-tasks, where we call this algorithm NEAS. The inputs for NEAS are all NE-tasks and the platform. From Line 2-13, NEAS splits task τ_i among one EE core and one PE core. For this combination NEAS selects the EE core with the smallest utilization and the PE core given by function `mpwr()` to split task τ_i in order to save the energy consumption. At Line 5-7 constraint (5.6) is checked. If the combination of one EE core and one PE core fails to allocate task τ_i , then, from Line 14-17, NEAS tries to integrally allocate task τ_i to one PE core which can schedule τ_i and has the minimum contribution to the energy consumption. If it does not successfully allocate τ_i to one PE core, NEAS splits τ_i among two PE cores from Line 18-24. In this case, it finds the PE core with the largest utilization to schedule the first subtask τ_i^1 . Because τ_i^1 requires the maximum frequency to guarantee the schedulability and the PE core with the largest utilization should execute at a high frequency compared to others, allocating τ_i^1 to the PE core would not increase the frequency too much which in turn does not lead to a lot of extra energy consumption for the task allocated to the PE core. For τ_i^2 , we still use function `mpwr()` to find the candidate core. If splitting among two PE cores fails, NEAS returns a failure.

5.5.3 The SPLIT function

In this section, we present the SPLIT function used in EAS and NEAS discussed above. Algorithm 8 presents the pseudo-code for SPLIT. The concept behind the SPLIT algo-

Algorithm 7: NE-task Allocation and Split (NEAS)

Input: All NE-tasks Γ_{NE} and the heterogeneous multicore platform

$$M = \{M_{EE}, M_{PE}\}$$

Output: Allocation for all NE-tasks

```

1  for  $\forall \tau_i \in \Gamma_{NE}$  in order of decreasing  $U$  do
2      for  $\forall x \in M_{EE}$  in order of increasing  $U$  do
3           $\tau_i^1, \tau_i^2 = \text{SPLIT}(\tau_i, x)$ 
4          if  $\tau_i^1 \neq \emptyset$  then
5              while  $\frac{C_i^2}{D_i^2} > \frac{C_i}{T_i}$  do
6                   $C_i^1 \leftarrow C_i^1 - 1$ 
7                  Recompute  $C_i^2$  according to the new  $C_i^1$  (see Section 5.2.4)
8                   $x \leftarrow \tau_i^1$ 
9                   $y \leftarrow \text{mpwr}(M_{PE}, \tau_i^2)$ 
10                 if  $y = \emptyset$  then
11                      $x \leftarrow x - \tau_i^1$ 
12                 else
13                      $y \leftarrow \tau_i^2$ 
14             if  $\tau_i$  is not allocated then
15                  $y \leftarrow \text{mpwr}(M_{PE}, \tau_i^2)$ 
16                 if  $y \neq \emptyset$  then
17                      $y \leftarrow \tau_i$ ; break
18             for  $\forall pe \in M_{PE}$  in order of increasing  $U$  do
19                  $\tau_i^1, \tau_i^2 = \text{SPLIT}(\tau_i, pe)$ 
20                 if  $\tau_i^1 \neq \emptyset$  then
21                      $x \leftarrow \tau_i^1$ 
22                      $y \leftarrow \text{mpwr}(M_{PE}, \tau_i^2)$ 
23                     if  $y = \emptyset$  then
24                          $pe \leftarrow pe - \tau_i^1$ 
25             if  $\tau_i$  is not allocated successfully then
26                 return Unschedulable
27 return Allocation of  $\forall \tau_i \in \Gamma_{NE}$ 
    
```

Algorithm 8: SPLIT

Input: τ_i and one processor x

Output: subtasks τ_i^1, τ_i^2

```

1   $C_i^1 = D_i^1 = (0.999 - U_x)T_i$ ;
2  Compute subtask  $\tau_i^2$  according to the parameters of  $\tau_i^1$  (see Section 5.2.4)
3  while  $C_i^2 > T_i - C_i^1$  and  $\tau_i$  is a NE-task and  $x$  is an EE core do
4       $C_i^1 \leftarrow C_i^1 - 1$ 
5      Recompute  $C_i^2$  according to the new  $C_i^1$  (see Section 5.2.4)
6   $\Gamma_x \leftarrow \Gamma_x + \tau_i^1$ ;
7  while True do
8      if  $C_i^1 < 1$  then
9          return  $\tau_i^1 = \tau_i^2 = \emptyset$ 
10     if  $QPA(\Gamma_x)$  reports unschedulable then
11          $t \leftarrow$  the failure point from QPA
12         while True do
13              $I = (t - \text{dbf}(\Gamma_x - \tau_i^1, t)) / \lfloor \frac{t + T_i - (C_i^1 - 1)}{T_i} \rfloor$ 
14             if  $I \neq C_i^1$  then
15                  $C_i^1 = C_i^1 - 1$ 
16             else
17                 Break;
18     else
19         Compute parameters for subtask  $\tau_i^2$  (see Section 5.2.4)
20     return  $\tau_i^1, \tau_i^2$ 

```

rithm is based on the approach proposed in [BDWZ12] and the properties of the C=D approach on heterogeneous multicore systems identified and discussed in Section 5.4. The inputs for SPLIT are a task τ_i and a core x while the output is two subtasks τ_i^1 and τ_i^2 . The objective of function SPLIT is to find the maximum WCET of τ_i^1 which can satisfy the schedulability on core x . The procedure is as follows:

- Initialize the parameters of subtasks. For τ_i^1 let $C_i^1 = D_i^1 = (0.999 - U_x)T_i$ (Line 1) and configure subtask τ_i^2 according to subtask τ_i^1 (Line 2), as explained in Section 5.2.4, where U_x denotes the total utilization of processor x ;
- If τ_i is a NE-task and x is an EE core (Line 3-5), ensure valid split according to Property 1;
- Use QPA [ZB09] to test whether subtask τ_i^1 can be allocated onto core x . If it is schedulable, return the subtasks τ_i^1 and τ_i^2 (Line 10, 18-20);
- If QPA reports ‘unschedulable’, recompute the WCET for subtask τ_i^1 . In this case, we use the recurrence approach from [BDWZ12] to make sure that

$$C_i^1 = (t - \text{dbf}(\Gamma_x - \tau_i^1, t)) / \lfloor \frac{t + T_i - (C_i^1 - 1)}{T_i} \rfloor \quad (5.7)$$

where t is the failure point returned by QPA, i.e., the time instance $\text{dbf}(\Gamma_x, t) > t$ and $\text{dbf}(\Gamma_x - \tau_i^1, t)$ represents the demand of tasks on core x excluding subtask τ_i^1 . The recurrence equation in Equation (5.7) computes a maximum value for C_i^1 such that $\text{dbf}(\Gamma_x, t) \leq t$ which ensures the schedulability of task set Γ_x at time instant t . If Equation (5.7) is satisfied, the recurrence procedure stops and returns C_i^1 for subtask τ_i^1 . Otherwise, it decrements C_i^1 by 1 and repeats the previous procedure (Line 10-17);

- Return failure if it cannot split task τ_i on core x (Line 8-9).

Note that we use 0.999 instead of 1 to initialize a subtask at Line 1, because if using 1 would result in that QPA uses the hyper-period of all tasks as bound to test the schedulability. Then, QPA would be very complex and time-consuming.

5.5.4 Computing the minimum frequency

We use VFS to scale down the frequency of each core so that the energy consumption is further reduced. However, next to implicit deadline tasks (i.e., unsplit tasks), we might have some subtasks obtained by splitting on some cores which are constrained deadline tasks. In such case, we cannot simply use the utilization-based approach [CK07] [BMAB16] to compute the minimum frequency. Hence, we integrate the

Algorithm 9: Compute Minimum Frequency (CMF)

Input: core x and task set Γ_x

Output: the minimum operating frequency for core x

```

1 if  $x$  has a first subtask then
2   return  $f_{max}$ 
3 else
4   Compute a minimum achievable frequency  $f_{crit}$  based on  $U_x$ 
5    $\bar{f} \leftarrow \{\forall f_i | f_i \geq f_{crit}\}$  and sort  $\bar{f}$  in order of increasing frequency
6   for  $\forall f_i \in \bar{f}, i = \{1, 2, \dots, k\}$  do
7     if  $QPA(\Gamma_x, f_i)$  reports schedulable then
8       return  $f_i$ 
9   return  $f_{max}$ 

```

frequency into QPA [ZB09] to efficiently compute the minimum frequency for a core.

Algorithm 9 (CMF) presents the pseudo-code to compute the minimum operational frequency for each core. The inputs are one core x and a task set Γ_x which includes all tasks allocated to core x . The output is the minimum operational frequency for core x . If the core has first subtask τ_i^1 , its frequency will be set to the maximum frequency according to Property 2 - see Line 1-2. Otherwise, we compute a minimum operational frequency for the core from Line 4-8. First, we compute a frequency called f_{crit} based on utilization U_x of core x [CK07]. Frequency f_{crit} can be deemed as the lower bound of the operational frequency of core x . If the operational frequency is lower than f_{crit} , the system is not schedulable. Then, we select all frequencies from the core's frequency set which are greater than f_{crit} and let these frequencies form a frequency set \bar{f} sorted in order of increasing frequency - see Line 5. We start with the smallest frequency f_i in frequency set \bar{f} and use QPA to test whether the task set is schedulable at this frequency - see Line 7. If it is schedulable, CMF returns frequency f_i as the operational frequency. Otherwise, we take frequency f_{i+1} and use QPA to test whether the task set is schedulable at this frequency.

5.5.5 The ASHM Algorithm

Given all algorithms explained earlier, we present our complete Allocation and Split algorithm ASHM using the pseudo-code in Algorithm 10. We first divide all tasks into two task sets Γ_E and Γ_{NE} , one for all E-tasks Γ_E and another for all NE-tasks Γ_{NE} . Then, we use EAS (Algorithm 6) to allocate all E-tasks - see Line 2. If all E-

Algorithm 10: ASHM

Input: all tasks Γ and the platform $M = \{M_{EE}, M_{PE}\}$

Output: the allocation for all tasks and the minimum operational frequency
for each core on the platform

```

1  $\Gamma_E \leftarrow$  all E-tasks,  $\Gamma_{NE} \leftarrow$  all E-tasks
2  $M \leftarrow \text{EAS}(\Gamma_E, M)$ 
3  $M \leftarrow \text{NEAS}(\Gamma_{NE}, M)$ 
4 for  $\forall x \in M$  do
5    $f_x \leftarrow \text{CMF}(x, \Gamma_x)$ 

```

tasks are successfully allocated, we proceed to allocate all NE-tasks by using NEAS (Algorithm 7) - see Line 3. Finally, we apply CMF (Algorithm 9) to compute the minimum frequency for each core - see Line 4-5.

Complexity Analysis: In the worst case, EAS, NEAS, SPLIT and CMF are all pseudo-polynomial algorithms due to QPA. Although QPA has shown its efficiency in [ZB09], its complexity is still pseudo-polynomial in the worst case. This worst-case scenario happens when the utilization U equals to 1. However, in function SPLIT, we strive to avoid the worst-case scenario to occur by setting the utilization bound as 0.999 - see Line 1 of Algorithm 8. Therefore, in practice, our algorithms can be executed very efficiently.

5.6 Evaluation

In this section, we present extensive experimental results to show the effectiveness of our ASHM algorithm in terms of energy consumption compared to two widely-used bin-packing algorithms [CGJ97] and two existing related approaches [CKR14] [ESAS14]. We do not compare with the algorithm proposed in Chapter 4, because the approach proposed in Chapter 4 is very similar to [CKR14] when we consider the per-core VFS system. We do not compare with [CST09], because they do not take VFS into account. Therefore, our approach will always save more energy consumption than [CST09]. Since the authors in [CKR14] have shown that their approach outperforms the allocation approach proposed in [HTC07], we do not compare our ASHM to [HTC07].

5.6.1 Experimental Setup

Task Generation

To evaluate the effectiveness of ASHM, we adopt the widely-used random task generator based on *UUnifast-discard* [DB11]. *UUnifast-discard* enables the generation of unbiased task sets. It takes as inputs the number of tasks n and the total utilization U and generates utilization u_i for n tasks. The generation procedure is summarized as follows:

- For each task, utilization u_i is generated using *UUnifast-discard*;
- Period T_i is generated using a log-uniform distribution with a factor of 100 difference between the minimum and maximum possible task period. This presents a range of task periods from 10ms to 1s in real-time applications [DB11] [BDWZ12];
- C_i^{PE} is computed as $C_i^{PE} = u_i \cdot T_i$; and
- C_i^{EE} is computed as $C_i^{EE} = C_i^{PE} \cdot ce_i$, where ce_i is selected from a uniform random distribution in the range $[1.8, 2.3]$ which represents the variance of the execution time on different types of cores [Jef12].

Platforms

We have two types of cores (PE and EE) in the platforms and the core's power parameters are shown in Table 5.1 taken from [LSCS15]. In this experiment, we evaluate the effectiveness of our ASHM algorithm mainly on platforms with limited number of resources because on a platform with more resources our approach will always perform good, especially with more EE cores. Therefore, we conduct experiments on the following three limited platforms:

1. Platform 1: 2 PE cores and 2 EE cores
2. Platform 2: 2 PE cores and 3 EE cores
3. Platform 3: 3 PE cores and 2 EE cores

On the three platforms, we experiment with task sets with different U and a different number of tasks.

Comparison approaches

We compare our proposed ASHM algorithm with the following approaches in terms of energy consumption:

- FFD: Allocate E-tasks and NE-tasks to EE cores and PE cores, respectively, using FFD [CGJ97]. If E-tasks cannot be allocated to EE cores, then they are allocated to PE cores using FFD;
- WFD: Similar to FFD, but instead of FFD we use WFD [CGJ97] to allocate tasks;
- EFD: The allocation algorithm proposed in [ESAS14];
- m-pwr: The allocation algorithm proposed in [CKR14];

Comparison Metric

In the experimental results, we show the energy saving by using our ASHM compared to the above four reference approaches. The energy saving is computed as follows:

$$\text{Energy saving} = \frac{E_{ref} - E_{ASHM}}{E_{ref}} \cdot 100[\%] \quad (5.8)$$

where E_{ref} is the energy consumption of one of the four approaches given above and E_{ASHM} is the energy consumption of our proposed ASHM.

5.6.2 Experimental Results

All experimental results are plotted in Figure 5.1 to 5.6. For each point in the figures, we generate 100 random task sets and compute an average energy saving. Note that only when all reference approaches can schedule the generated task set we compute the energy saving using Equation (5.8). Our ASHM always can schedule more task sets than the other approaches because ASHM uses task-splitting. Since the schedulability advantage of the task-splitting approach has been reported in [BDWZ12], we do not compare the number of schedulable task sets in this work.

Impact of the Utilization

In this experiment, we fix the number of tasks for different platforms and then vary the total utilization to evaluate the effectiveness of ASHM. In order to have both NE-tasks and E-tasks in the generated task set, the number of tasks is fixed to 7 for all platforms. The results are plotted in Figure 5.1 to 5.3 where the y-axis is the energy saving computed using Equation (5.8) and the x-axis is the variable utilization. We can see that

our ASHM outperforms all allocation approaches in terms of energy efficiency. From the experimental results, we observe:

- The average energy saving by ASHM decreases as the total utilization increases. In the comparison between ASHM and WFD, EFD and m-pwr, this trend is easy to be observed although there is some variation due to the randomness of the generated task sets. The reason is that when we increase the total utilization, the slack space on the EE cores is reduced such that the task set cannot benefit from our ASHM too much. However, for FFD in Platform 1 and 3, see Figure 5.1 and 5.3, the energy saving increases until a point and then gradually decreases. The reason is that when we have task sets with a low utilization, FFD always tries to use the smallest number of cores to schedule tasks which might cause the PE cores to execute at a high frequency. The high frequency in turn leads to high energy consumption.
- ASHM saves more energy consumption on a platform with more EE cores, see Figure 5.2. The advantage of ASHM is to effectively utilize EE cores on the platform to achieve energy efficiency. More EE cores provide more space to split tasks and thus ASHM reduces more the energy consumption.

Impact of the Number of Tasks

In this experiment, we fix the utilization for different platforms and then vary the number of tasks to evaluate the effectiveness of ASHM. Since larger total utilization leads to smaller number of schedulable task sets, we fix the utilization to 2 for all platforms in order to compare our ASHM to the reference approaches on more schedulable task sets. We ensure that the number of tasks is greater than the number of cores, so we start with 4 tasks for Platform 1 and 5 tasks for Platform 2 and 3. The results are plotted in Figure 5.4 to 5.6.

Compared to the well-performed allocation approaches WFD and m-pwr, we can see that the energy saving is decreasing with the increasing number of tasks. The reason is that when the number of tasks increases with a fixed utilization, the tasks in the set become lighter, i.e., with a smaller utilization. Therefore, these tasks are easy to be allocated among the cores and then EE cores might be completely fulfilled or just have a little space for splitting of tasks. Therefore, ASHM cannot save too much in this case. However, as can be seen in Figure 5.4 and 5.6, compared to FFD, the energy saving by ASHM increases gradually. Since we have more tasks with a low utilization, FFD might allocate all tasks onto one core which will execute at a high frequency. However, since the dynamic power consumption still dominates the total power consumption, executing on two PE cores with lower frequencies is more energy-efficient than on one PE core with a high frequency.




Figure 5.1: Varying U on platform with 2 PE cores and 2 EE cores




Figure 5.2: Varying U on platform with 2 PE cores and 3 EE cores




Figure 5.3: Varying U on platform with 3 PE cores and 2 EE cores

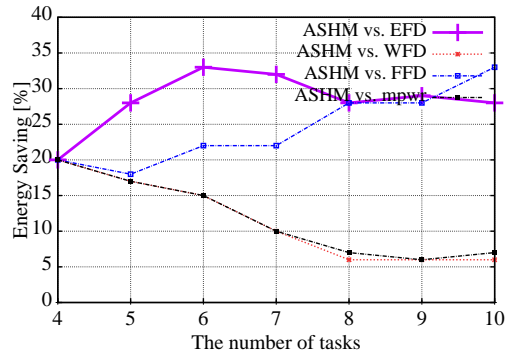


Figure 5.4: Varying the number of tasks on platform with 2 PE cores and 2 EE cores




Figure 5.5: Varying the number of tasks on platform with 2 PE cores and 3 EE cores




Figure 5.6: Varying the number of tasks on platform with 3 PE cores and 2 EE cores

5.7 Discussion

ASHM shows its effectiveness on per-core VFS systems via the experimental results in Section 5.6. We need to notice that a slight modification is capable of adapting ASHM to cluster heterogeneous multicore systems as considered in Chapter 4, but we leave this for future work. On the other hand, since the *HRT* scheduling framework seen in Section 2.3 can convert CSDF graphs into periodic task sets which can be fed as the input of the proposed ASHM, ASHM can also be applied by on a CSDF graph under the *HRT* scheduling.

