

# Energy-aware Time- and Event-triggered KVM Nodes

Isser Kadusale, Gautam Gala and Gerhard Fohler  
Chair of Real-time Systems,  
University of Kaiserslautern-Landau (RPTU), Germany  
{kadusale,gala,fohler}@eit.uni-kl.de

**Abstract**—Industries are considering the adoption of cloud and edge computing for real-time applications due to current improvements in network latencies and the advent of Fog and Edge computing. Current cloud paradigms are not designed for real-time applications, as they neither provide low latencies/jitter nor the guarantees and determinism required by real-time applications. Experts estimate that data centers use 1% of global electricity for powering the equipment, and in turn, for dealing with the produced heat. Hence, energy consumption is a crucial metric in cloud technologies. Applying energy conservation techniques is not straightforward due to the increased scheduling overheads and application execution times. Inspired by slot shifting, we propose an algorithm to support energy-aware time-triggered execution of periodic real-time VMs while still providing the ability to execute aperiodic real-time and best-effort VMs in the slack of the time-triggered ones. The algorithm considers energy reduction techniques based on dynamic power management and dynamic voltage and frequency scaling. We implement our algorithm as an extension to the Linux kernel scheduler (for use with the KVM hypervisor) and evaluate it on a server-grade Intel Xeon node.

## I. INTRODUCTION

The usage of cloud technologies is continuously rising. The cloud offers better hardware efficiency and maintainability, higher availability, and more cost-effective scaling compared to traditional systems. For these reasons, cloud computing is picking up steam in safety-critical domains such as in railway [1], [2], [3], [4], automotive [5], and smart manufacturing [6]. However, safety-critical applications in these domains require timing guarantees that current cloud offerings cannot fulfill.

Time-triggered scheduling is widely used in the safety-critical domain because of its determinism and strong isolation properties. It can be used in cloud environments to help fulfill the timing demands of safety-critical applications. Specialized hypervisors, such as XtratuM [7], support time-triggered VMs but are proprietary and unsuitable for cloud computing. Previous work for time-triggered generic hypervisors, such as Tableau [8] and ARINC-653 scheduler for Xen [9], support time-triggered VMs in a private cloud. However, these approaches do not allow adding VMs to the existing scheduling table or have a high overhead table regeneration process. Moreover, previous work demonstrated that the KVM (with PREEMPT\_RT) hypervisor [10], in general, has considerably lower latency as compared to Xen, making KVM better suited to time-sensitive and safety-critical applications [11], [12], [13]. In addition, most time-triggered solutions undermine the benefits of cloud computing since they do not utilize the unused resources from time-triggered VMs.

Another aspect that existing time-triggered cloud hypervisors overlook is energy. High energy usage negatively impacts the environment, and energy consumption accounts for a significant portion of the costs of a data center [14]. Ignoring energy in computing systems also leads to more heat generation, increasing cooling requirements, and decreasing reliability and longevity. Therefore, it is essential to use energy-reduction techniques in computing systems. However, applying energy-reduction techniques to the real-time cloud adds another dimension to the problem, since they impact scheduling overheads and application execution times. These techniques must maintain timing guarantees, especially for safety-critical applications.

The slot shifting algorithm [15] combines time-triggered and event-triggered scheduling to get the benefits of both. It allows flexibility in a time-triggered schedule while maintaining its determinism and isolation properties. Slot shifting can accept and guarantee execution time to event-triggered tasks by keeping track of the available slack in the time-triggered schedule. This system of tracking the available slack could also be extended and used to ensure that energy-reduction techniques do not cause timing violations.

Inspired by the slot-shifting algorithm, we propose a novel energy-aware scheduling algorithm for cloud nodes to ensure that safety-critical VMs meet their requirements. Utilizing the concepts in slot shifting allows for having a time-triggered schedule with the flexibility of handling event-triggered tasks. In summary, our scheduler:

- supports time-triggered and event-triggered safety-critical/time-sensitive VMs.
- uses energy-reduction techniques to conserve energy without impacting the timing guarantees of safety-critical VMs.
- exploits the slack of the time-triggered VMs to execute best-effort/non-critical VMs.
- has low overheads when adding new VMs by appending them to the existing scheduling table at run-time instead of completely regenerating it.
- can be plugged into an RT-cloud resource orchestration solution such as those proposed in [11], [16].

Finally, we implemented our scheduler inside the Linux Kernel. Our approach keeps run-time overhead low by separating the TT task dispatching via a newly created scheduling class (on TT/managed cores) from the energy-aware slot-shifting algorithm decision-making. We execute the energy-

aware slot-shifting algorithm on non-TT/non-managed cores with a period equal to the slot length. We experimentally evaluated server-grade Dell hardware with an Intel 2<sup>nd</sup> generation Xeon processor (Cascade Lake).

The remainder of the paper is as follows: Section II presents related work. Section III introduces the used system models. Section IV and Section V describe the proposed scheduler and its implementation for the Linux kernel. Section VI presents our evaluation, and VII concludes the paper.

## II. RELATED WORK

*a) Joint scheduling of TT and ET RT tasks:* Many existing works have explored joint time-and event-triggered approaches, e.g., [15], [17], [18]. However, they do not integrate these approaches to combine TT and ET tasks with the Linux kernel scheduler. Fohler [15] presented the Slot-shifting algorithm for joint scheduling of TT and ET tasks. Schorr [17] presented a multiprocessor extension to the slot-shifting algorithm. We further extend the slot-shifting algorithm and propose two approaches to integrate energy awareness (using DPM and DVFS).

*b) Dynamic Voltage and Frequency Scaling (DVFS) techniques:* Early research on energy-aware real-time scheduling explored the use of Dynamic Voltage and Frequency Scaling (DVFS) techniques which adjust the voltage and frequency of the processor to save energy [19]. DVFS involves identifying the slack in a schedule and using that information to allow the execution of tasks with a lower processor frequency without violating their deadlines. Previous works (e.g., [20]) proposed offline algorithms determining the optimal execution frequency for tasks in a task set, assuming that the tasks execute until their WCET. Dynamic approaches, such as [21], [22], take advantage of the unused execution time of tasks finishing earlier than their WCETs.

*c) Dynamic Power Management (DPM) techniques:* The power consumption of a processor has a dynamic and static component. The dynamic power consumption is due to circuit switching activity, while the static power consumption is due to leakage current.

Unlike DVFS, Dynamic Power Management (DPM) techniques reduce static power consumption by switching the processor to a lower power mode. However, while in a low-power mode, the processor cannot execute any tasks, and transitioning to and from a low-power mode incurs a time and energy overhead. Processors typically have several low-power modes, each with a different time and energy overhead. In general, modes with lower power consumption have higher overheads. Energy-aware real-time scheduling algorithms [23], [24] consolidate idle periods so that the processor can be put in a lower power mode and minimize transition overheads.

*d) Energy-saving real-time approaches for clouds:* Chen et al. [25] proposed an energy-aware scheduling scheme for scheduling jobs in a virtualized cloud. The proposed method dynamically creates, deletes, and consolidates VMs to meet real-time requirements while saving energy. Zhu et al. [26] presented an energy-aware scheduling architecture for real-time tasks in virtualized clouds. They employ a rolling horizon for waiting tasks, allowing rescheduling for better system schedulability and less energy consumption. These

works, however, only handle aperiodic tasks. Sun and Cho [27] proposed a real-time scheduling algorithm for cloud computing servers. They use flow networks to schedule periodic tasks, ensuring the system meets the deadlines while consolidating idle periods to maximize energy savings using DPM. The proposed scheduler does not consider aperiodic tasks. In addition to only being limited to a single type of task, these schedulers cannot offer the same determinism that a time-triggered scheduler offers.

*e) Energy-saving techniques in Linux/KVM:* Linux/KVM offers dynamic CPU frequency scaling (DVFS) via the CPUFreq subsystem. Linux frequency scaling governors implement the algorithms to compute the desired CPU frequency. At the same time, the scaling drivers interact with the CPU directly to enact the frequencies computed by the current governor.

Linux/KVM offers dynamic power management (DPM) that turns off the power or switches system components to a low-power state when inactive/idle for a pre-defined minimum amount of time.

Unfortunately, these energy-saving features in the Linux kernel need to be often disabled for use with real-time systems as they impact the kernel latency and determinism.

*f) Energy-aware Time-triggered hypervisors/OSes:* Proprietary real-time hypervisors, such as XtratuM [7], often support energy-aware (DPM) table-driven scheduling but are not suitable for use in cloud environments as they have limited hardware and existing software support, allow only a handful of guest operating systems and have other limitations such as limited CPU models for VMs.

Xen with ARINC 653 scheduler [9] and Tableau [8] extension to Xen hypervisor introduce support for table-driven scheduling of VMs but are not energy-aware. Moreover, authors in [11] demonstrated that Xen, in general, has higher overheads than Linux (+PREEMPT\_RT patch) in conjunction with KVM, and therefore, KVM is better suited for low latency real-time safety-critical applications.

Very recently, Karachatzis et al. [28] presented an implementation and evaluation of a kernel-level time-triggered scheduling approach for Linux processes. However, the approach turns off energy-saving features in the Linux kernel to reduce the kernel's latency and improve determinism. In addition to energy awareness in Linux/KVM, our approach supports the guaranteed execution of ET RT VMs by appending them to the scheduling table at run-time (without needing to regenerate the table).

Gala et al. [29] proposed a new TT scheduling class for the Linux kernel and integrated slot-shifting to support guaranteed execution of ET RT tasks. We build upon this work to support TT execution VMs (in conjunction with KVM). Additionally, we integrate our new slot-shifting extension to support energy awareness while providing flexibility.

## III. MODELS

We consider a time-triggered approach and assume a timeline that is divided into equal lengths called slots. A slot is a time interval,  $[t, t + 1)$ , where  $t$  is the slot length. Slots are

synchronized between all cores, and scheduling decisions take effect at the start of each slot.

We adopt the same model for VMs and their mapping to periodic tasks as in Tableau [8] as well as their mapping to periodic real-time tasks. Each VM consists of one or more vCPUs and each vCPU has a reserved utilization,  $U$ , and maximum scheduling latency,  $L$ . Each vCPU is then mapped to a periodic task,  $\tau_i$ , with a corresponding worst-case execution time (WCET),  $C_i$ , and period,  $T_i$ , and where the relative deadline is equal to the period. Since a task can be scheduled immediately at the start or right at the end of a period, the longest gap between two invocations of a task is  $2(T - C)$ , or equivalently  $2T(1 - U)$ . This would be the maximum scheduling latency experienced by the vCPU. Therefore, in order to meet the latency requirement, a suitable period,  $T$ , must be selected that satisfies Equation 1.

$$T \leq \frac{L}{2(1 - U)} \quad (1)$$

With the period,  $T$ , selected, the WCET can be calculated as  $C = UT$  to meet the required utilization. We assume that the specified utilization,  $U$ , and by extension the resulting WCET,  $C$ , takes into account possible interference from other cores.

A single invocation of a task,  $\tau_i$ , is referred to as a job,  $j_k^i$ . Each job has a remaining execution time,  $c_k^i$ , a release time,  $r_k^i$ , and an absolute deadline,  $d_k^i$ . For periodic tasks, a job is released at every period,  $T_i$ , (i.e.  $r_{k+1}^i = r_k^i + T_i$ ) with an absolute deadline of  $d_k^i = r_k^i + T_i$ . Each occurrence of an aperiodic task is also a job with the mentioned properties.

The remaining execution time of a job is initialized as the WCET, and the value is decreased by  $f_{slot}$  every time the job is executed in a slot. The normalized frequency,  $f_{slot}$ , is the ratio of the frequency of the core during the slot execution relative to the maximum frequency (i.e.  $f_{slot} = F_{slot}/F_{max}$ ).

#### IV. SCHEDULER

We assume a time-triggered schedule has been constructed for each core with information on the release time, deadline, and WCET of the VMs and take this as input for the scheduler. Our method is agnostic to the algorithm used to create the schedule. We then use the same concepts presented in slot shifting [15] to provide flexibility in a time-triggered schedule.

A table is created from the input schedule of each core. The table is then annotated with information about the flexibility and leeway available in the schedule. An example of how this is done is shown in Figure 1.

The entire scheduling table is divided into groups of consecutive slots referred to as a capacity interval ( $I$ ), or simply an interval, as shown in Figure 1b. The division is based on the jobs' release times and deadlines. Every unique deadline defines the end of an interval. Each job with the same deadline as the end of an interval belongs to that interval. The start of the interval is the maximum of either the end of the previous interval or the earliest release time of the jobs belonging to the interval. Consecutive slots that do not have a job ready for execution are grouped into a single empty interval.

Each interval has an associated spare capacity value which represents execution time in the interval that is not committed

to the jobs in the schedule. This is calculated from the length of the interval, the WCETs of the jobs belonging to the interval, and the spare capacity of the succeeding interval as shown in Equation 2. If the spare capacity of an interval is negative, this means that it borrows slots from one or more previous intervals to execute jobs that belong to it. An example of the spare capacity calculation is illustrated in Figure 1c.

$$sc(I_i) = |I_i| - \sum_{T_j \in I_i} C_j + \min(sc(I_{j+1}), 0) \quad (2)$$

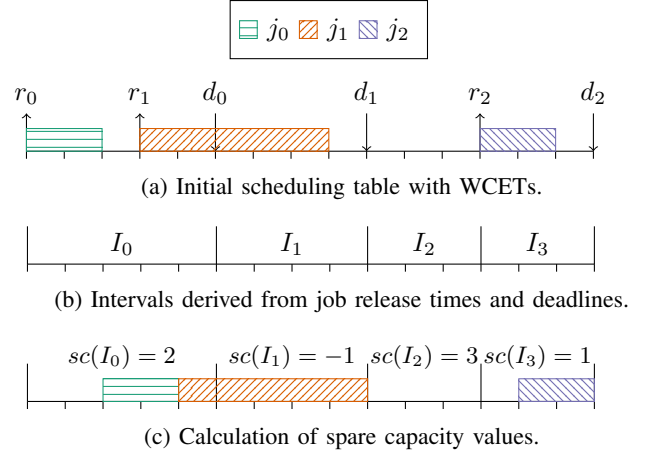


Fig. 1: Construction of annotated scheduling table.

The spare capacity values are regularly updated during runtime, as shown in Algorithm 1. At every slot, execution time is lost for the current interval, so the spare capacity of the current interval is decremented by 1. If a job,  $j$ , is executed for a slot, this also frees up execution time for the job's interval,  $I_j$ . If the job is executed at maximum frequency, a full slot is freed from the job's interval, so its spare capacity,  $sc(I_j)$ , is increased by 1. At less than maximum frequency, only part of a slot's execution time is freed — the same amount that the job's remaining execution time is decreased by,  $f_{slot}$ . However, this extra execution time can only be used by the job itself (to execute at less than maximum frequency) since only one job is executed at each slot. We refer to this freed partial slot as the *reserved spare capacity* of the job,  $reserved\_sc(j)$ , and it accumulates every time a job is executed at less than maximum frequency. When the reserved spare capacity of a job reaches 1 or greater, this means a full slot has been freed, which can now be used by other jobs. This is then converted to the spare capacity of the job's interval by incrementing the interval's spare capacity by 1 and decrementing the job's reserved spare capacity by 1. If the interval borrows spare capacity from one or more previous intervals, the spare capacity of the previous intervals that it borrows from is also incremented.

During runtime, the scheduler makes a scheduling decision at every slot. Jobs with earlier deadlines are given higher priority. If there are no jobs to run based on the scheduling table, best-effort/non-critical VMs can be executed.

##### A. Adding VMs and handling Event-Triggered Tasks

Slot shifting has a method for integrating new jobs into the scheduling table called the guarantee algorithm. Regenerating new scheduling tables for when a new VM needs to be added can be a relatively long process. To quickly execute new VMs,

---

**Algorithm 1** Updating spare capacity values

---

```
function ssmod_timer_cb(struct hrtimer * timer)
  sc(Icurr) := sc(Icurr) - 1
  reserved_sc(j) := reserved_sc(j) + fslot
  if reserved_sc(j) ≥ 1 then
    reserved_sc(j) := reserved_sc(j) - 1
    prev_sc := sc(Ij)
    sc(Ij) := sc(Ij) + 1
    Ii := Icurr
    while prev_sc < 0 do
      Ii := Ii-1
      prev_sc := sc(Ii)
      sc(Ii) := sc(Ii) + 1
    end while
  end if
end function
```

---

the guarantee algorithm is used to integrate the jobs of a new VM to the existing schedule.

There are also applications whose executions are unpredictable and intermittent, where it would be more suitable to schedule them on demand rather than regularly. These can be modelled as aperiodic jobs and integrated into the schedule as needed with the acceptance test and guarantee algorithm.

The guarantee algorithm first conducts an acceptance test, which checks if the job can be added to the schedule and be completed within its deadline. This is done by tallying up the spare capacities up to the deadline of the job. If there is enough spare capacity for the worst-case execution time of the job, then it can be accepted. If there isn't enough spare capacity, the scheduler attempts to delegate the job to another core using the negotiation-based acceptance test global algorithm described in [17].

After a job is accepted, the new job is added to the scheduling table. A new interval is created (if needed) and spare capacity values are updated. This guarantees that the new job's worst-case execution time is considered as other jobs are added or the energy-aware techniques described in the next sections use up the spare capacity.

### B. Energy-aware slot shifting extension

The following sections outline our extension to the slot shifting algorithm to use energy-reduction techniques. We describe two separate algorithms, one using DPM (EASS-DPM) and another using DVFS (EASS-DVFS).

1) *EASS-DPM*: This algorithm attempts to combine idle slots to maximize energy savings by going to a deeper sleep state and minimizing the number of sleep transitions.

The algorithm executes all jobs that are able to execute until completion, and then postpones the execution of the upcoming scheduled jobs as much as possible. When there are no jobs ready to execute in the upcoming slot, the idle period duration is calculated based on the spare capacity of the current interval and possibly succeeding intervals. The algorithm for calculating the idle duration is show in Algorithm 2.

At minimum, the core can be idle for the duration of the spare capacity of the current interval,  $I_{curr}$ . If there are no jobs in the current interval that have remaining execution times, the

---

**Algorithm 2** Low Power Duration

---

```
function low_power_duration
  duration := sc(Icurr)
  if rem_exec(Icurr) == 0 then
    Ii := Icurr+1
    duration := duration + max(0, sc(Ii))
    while empty(Ii) do
      Ii := Ii+1
      duration := duration + max(0, sc(Ii))
    end while
  end if
  return duration
end function
```

---

---

**Algorithm 3** Available Spare Capacity

---

```
function available_spare_capacity
  spare_capacity := max(0, sc(Ij)) + reserved_sc(jk)
  Ii := Icurr
  while Ii! = Ij ∧ rem_exec(Ii) == 0 ∧ sc(Ii) > 0 do
    spare_capacity := spare_capacity + sc(Ii)
    Ii := Ii+1
  end while
  return spare_capacity
end function
```

---

core can also be idle for the spare capacities of subsequent intervals up to and including the first interval that has a job that has remaining execution time. This results in the job with the next closest deadline to be executed as late as possible with enough slots for its WCET.

2) *EASS-DVFS*: This algorithm attempts to use the available spare capacity to execute the current task at the lowest frequency without violating any deadlines.

At each slot, the frequency,  $f_{slot}$ , to run the core with is calculated. If there is no job to execute in the upcoming slot, the core frequency is set to minimum. If there is a job,  $j_k$ , to execute, the available spare capacity that the job can use is calculated as shown in Algorithm 3. This consists of the job's reserved spare capacity, the spare capacity of the job's interval, and, if the job's interval is not the current interval, the spare capacities of the current and subsequent intervals that have spare capacity and no longer have jobs to execute.

With the available spare capacity, we then calculate the frequency,  $f_{slot}$ , that the job,  $j_k$ , will be executed at using Equation 3. This spreads the remaining execution time,  $c_k$ , of the job to use up the available spare capacity.

$$f_{slot} = \frac{c_k}{c_k + \text{Available Spare Capacity}} \quad (3)$$

## V. LINUX IMPLEMENTATION

Our work builds upon TT scheduling class for the Linux kernel proposed in [29]. We implemented our scheduler extensions on the Linux kernel (v5.19.9) with the PREEMPT\_RT patch. The available CPU cores are divided into managed cores and non-managed cores. Processes that are managed by the scheduler are executed on the managed cores, while other processes may only run on the non-managed cores. To enforce this, the kernel command line parameters *isolcpus* and *irqaffinity* were used.

The implementation for the energy-aware scheduler consists of the following parts: a Linux scheduling class, a Linux kernel module, and CPUFreq and CPUIdle governors. An overview of the implementation and how it interacts with the core Linux kernel is shown in Figure 2.

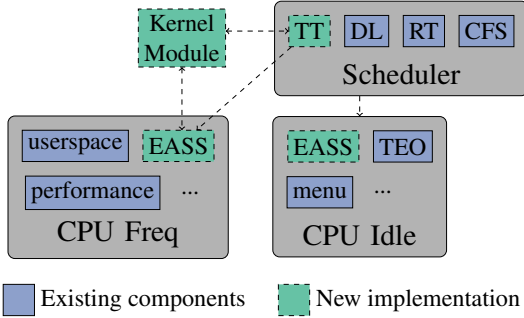


Fig. 2: Linux implementation overview.

Scheduling decisions in the Linux kernel are handled by scheduling classes. Within the core Linux scheduler’s `__schedule()` function, the scheduling classes are polled one by one, by order of priority, if they have a process to schedule. If higher priority scheduling classes return a process to be scheduled, this will be scheduled over processes belonging to a lower priority scheduling class. A new scheduling class was implemented for the energy-aware scheduler and was given the highest priority. This scheduling class performs two main functions: triggering the rescheduling of a core at the start of each slot, and returning the correct process to be scheduled in the upcoming slot.

Most of the scheduling algorithm logic is implemented in the kernel module so that the scheduling class can be used with other scheduling algorithms without needing to recompile the kernel. The kernel module also implements a `sysfs` interface that lets users control the scheduler, as well as a `debugfs` interface for collecting logs and statistics.

The CPUIdle subsystem of the Linux kernel handles switching to low power states. When a core goes idle, the subsystem is invoked to select and switch to the appropriate low power state. The CPUIdle subsystem has CPUIdle governors which handles the logic for selecting the low power state. A new CPUIdle governor was created which is used for the EASS-DPM scheduler.

The CPUFreq subsystem of the Linux kernel handles setting the core frequencies of the CPU. Like the CPUIdle subsystem, it also has CPUFreq governors which handles the logic for the frequency switching. A new CPUFreq governor was created to be used for the EASS-DVFS scheduler. Its function is to translate the calculated normalized frequency to a frequency that’s available on the platform, as well as interface with the CPUFreq subsystem to change the core frequency.

### A. Runtime Operation

The runtime operation of the scheduler is driven by the `hrtimer` interrupts of the scheduling class and the kernel module. These interrupts execute relative to the start of every slot, as shown in Figure 3.

The kernel module sets up an `hrtimer` interrupt on a non-managed core, referred to as the module core, which runs

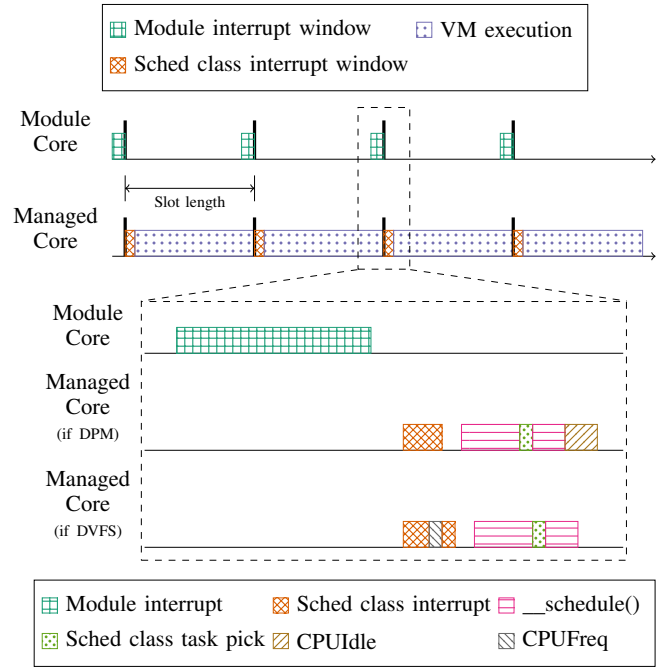


Fig. 3: Interrupt execution on each slot.

before the start of every slot. This goes through the scheduling algorithm for each core. It handles moving jobs to different states, accepting new jobs by running the guarantee algorithm, maintaining the spare capacity values of the intervals, and selecting which job to run on each core which it then passes to the scheduling class. With EASS-DPM, if the core will be idle in the upcoming slot, the idle period duration is calculated and passed to the scheduling class. For EASS-DVFS, the ideal normalized frequency is calculated by the module. This value is passed to the CPUFreq Governor, which finds the lowest available frequency that’s higher than or equal to the calculated frequency and stores it for the upcoming slot.

The scheduling class has `hrtimer` interrupts on each managed core, which fires at the start of each slot. The interrupt checks the scheduling decision made by the kernel module and marks the core for rescheduling if necessary. This, then, starts the scheduling procedure of the Linux core scheduler. The interrupt also tells the CPUFreq governor to change the core frequency in the EASS-DVFS version. With EASS-DPM and if the core is set to be idle for longer than two slots, the next interrupt for that core is delayed until the end of the idle period. Since all other processes and interrupts are configured to run on the housekeeping core, the CPUIdle governor can use the functions made available by the CPUIdle subsystem to determine the time until the next timer interrupt and select the most appropriate low power state. There is no direct communication between the CPUIdle governor and the rest of the EASS-DPM scheduler.

## VI. EVALUATION

The scheduler was evaluated on a Dell R640 server with an Intel Xeon Gold 5218 processor. The processor has 16 cores with a max turbo frequency of 3.9 GHz, but this was limited to the base frequency of 2.3 GHz. Each core’s frequency can be set independently within a range of 1.0 GHz and 2.3 GHz, with increments of 100 MHz. For our tests, CPU0 was set as

a non-managed core while the other cores, CPU1 to CPU15, were managed by the scheduler.

We used the UUnifast algorithm [30] to generate the VM specifications for the experiments. We specify parameters for the algorithm such as the total target utilization, WCET range, and period range. Its output is a set of files containing the properties of the jobs of the VMs. When running the experiment, this output is parsed and fed to the kernel module’s sysfs interface.

The performance of the scheduler in terms of power consumption and scheduler overhead was evaluated. We tested three versions of the scheduler: a base slot shifting scheduler without any energy-saving techniques (BSS), the EASS-DPM scheduler, and the EASS-DVFS. For BSS and EASS-DVFS, a CPUIdle governor that doesn’t put the cores in a low power state was used. The userspace CPUFreq governor was used for BSS and EASS-DPM, and all cores were set to run at maximum frequency.

### A. Power consumption

To evaluate performance in terms of energy-savings, we generated several test cases with the parameters shown in Table I. VMs that are already integrated into the schedule have a total utilization ranging from 20% to 80% with increments of 10%. New jobs that need to be added to the schedule during runtime have a total utilization of 10%, 20%, or 50%. Ten test cases were generated for each combination of the utilization of the scheduled VMs and the utilization of new jobs. Each schedule was 1800 to 2200 slots long and was tested with 3 different slot lengths: 1ms, 3ms, and 10ms. Each test case was run 3 times for each slot length.

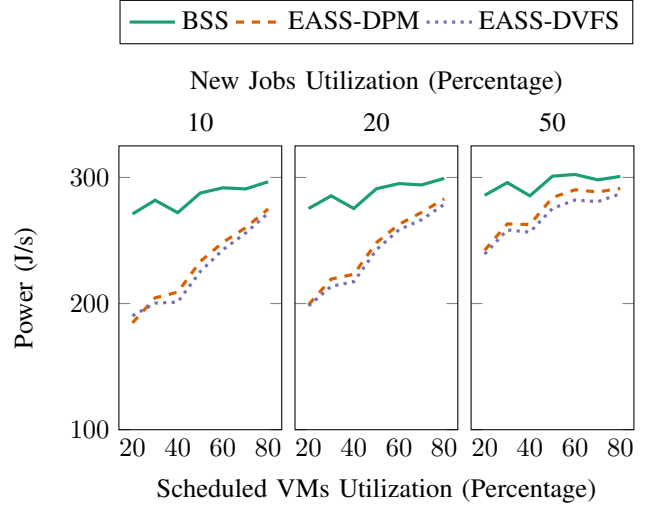
Parameter	Scheduled VMs	New Jobs
WCET	[1,15]	[10,15]
Period	[15,50]	[10,15]
Utilization	[20 - 80]%	[10,20,50]%

TABLE I: Parameters of test cases in the energy-usage test.


The CPU’s energy consumption was measured using the Linux kernel’s power capping framework [31]. The *energy\_uj* attribute of the CPU package was recorded at the start, and end of every schedule execution and the difference was divided by the duration of the schedule to get the power measurement.

Figure 4b shows the power percentage difference when using a slot length of 3ms and 10ms, with the difference in average power of each utilization combination as a data point. There is a decrease in power consumption as the slot length gets larger, but the difference is minute. This difference could be due to more frequent interrupts and context switches.

Figure 4a shows the average power consumption for each utilization combination. There is a significant decrease in power consumption when comparing the energy-aware schedulers to that of the base scheduler, with as much as a 31.85% and 29.81% decrease on the lowest utilization values for EASS-DPM and EASS-DVFS respectively. As can be expected, the difference decreases as utilization increases. At 50% total utilization, there was a decrease of roughly 23% for EASS-DPM and 28% for EASS-DVFS. This tapers down to a 3.16% difference at the highest utilization level for EASS-DPM and a 4.58% difference for EASS-DVFS.



(a) Average power consumption.



(b) Power difference compared to 1ms slot length.

Fig. 4: Power consumption results.

### B. Scheduling overhead

To measure the scheduler overhead, 50 test cases were generated with the parameters shown in Table II. The schedule durations for each test case ranged from 480 to 520 slots, with a slot length of 3ms. Each test case was run 5 times for each scheduler.

Parameter	Scheduled VMs	New Jobs
WCET	[1,15]	[10,15]
Period	[15,50]	[10,15]
Utilization	50%	50%

TABLE II: Parameters of the test cases in the overhead test.

Figure 5a shows the measurements from the start of the scheduling class interrupt up until the end of the `__schedule()` function. We see that for the majority of cases (i.e. within the 1st and 3rd quartile), EASS-DVFS has a higher switch duration when compared to BSS and EASS-DPM. On average, EASS-DVFS had a switch duration of  $2.74\mu s$  as compared to the  $1.44\mu s$  and  $1.52\mu s$  of BSS and EASS-DVFS respectively. One reason for this is that EASS-DVFS switches the core frequency during the scheduling class interrupt, which results in more overhead, as can be seen in Figure 5b. Another reason, is that EASS-DVFS lowers the frequency of the managed cores, which also slows down the scheduler.

It is also apparent in Figure 5 that, even though EASS-DPM had comparable values to the base scheduler in the average case, it had significantly higher values above the 3rd quartile. We’ve observed that these high values occur right after returning from a low power state. This can be seen in Figure

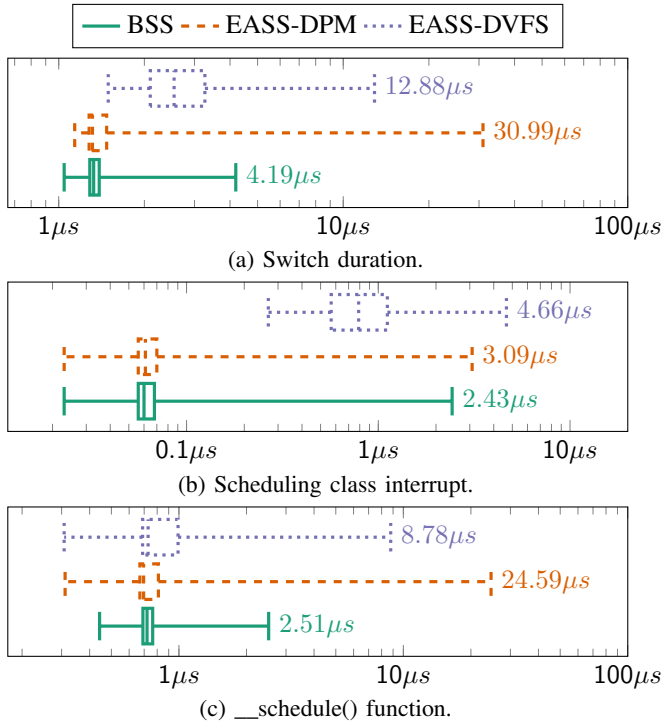


Fig. 5: Overhead measurements for managed cores.

6, which shows the overhead data for EASS-DPM but with the instances immediately after returning from a low power state separated from the rest of the data. Even with the highest observed value of  $30.99\mu s$ , this is still only about 1% of the slot length in this test.




Fig. 6: Overhead measurements for EASS-DPM on managed cores.

In Figure 7a we compare the module interrupt durations between the three versions of the scheduler. Note that the module interrupt runs on a non-managed core and, therefore, does not affect any of VMs managed by the scheduler.

The values from EASS-DPM are comparable to the base scheduler in the average case, and has significantly lower values below the 1st quartile. This is because the scheduler algorithm is not run on cores that are in a low power state. As shown in the per core measurements in Figure 7b, EASS-DPM is quite close to the base scheduler on a per core basis.

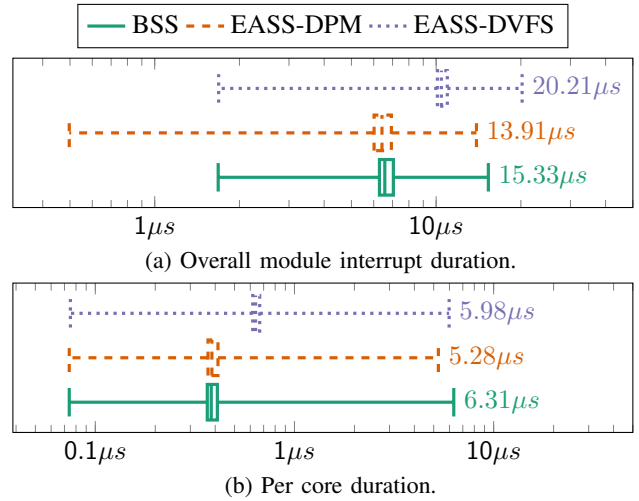


Fig. 7: Overhead measurements on the module core.

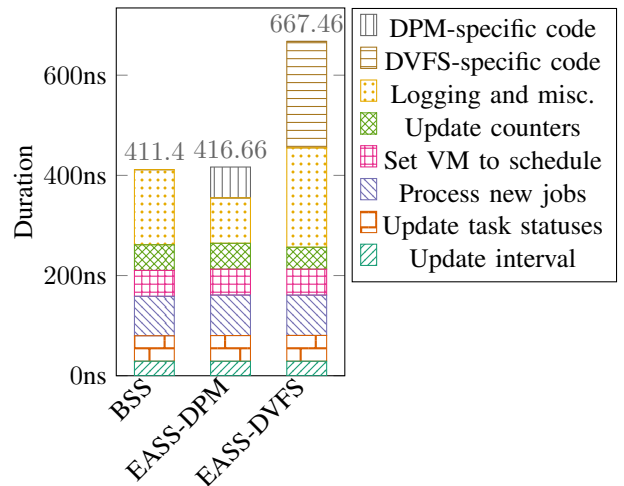


Fig. 8: Breakdown of average per core measurements.

In contrast, significantly higher values were measured for the EASS-DVFS. The added functionality of the DVFS algorithm amounts to a significant increase in the overhead measurements for the scheduler module, as can be seen in Figure 8, which shows a breakdown of the overhead values for different parts of the scheduler algorithm. This extra overhead for the EASS-DVFS includes calculating for the ideal frequency, as well as sending it to the CPUFreq governor to find a suitable frequency that's available on the CPU.

## VII. CONCLUSION

We presented an energy-aware time-triggered scheduler for cloud nodes based on Kernel Virtual Machine (KVM) in conjunction with Linux. The scheduler uses concepts from the slot shifting algorithm to provide flexibility at run-time for adding real-time Virtual Machines (VMs) to the offline scheduling table without requiring a high overhead table regeneration process. In addition, the scheduler also uses the slack in the schedule to apply energy-saving techniques without compromising the requirements of the real-time VMs.

We presented two versions of the scheduler as an extension to slot shifting: EASS-DPM and EASS-DVFS. We implemented both versions on Linux for use with KVM. We

performed experiments on server-grade hardware with Intel Xeon 2<sup>nd</sup> generation processors (Cascade Lake) to evaluate both the scheduler implementations for energy savings and scheduling overheads.

The evaluation shows a significant reduction in power consumption from the two energy-aware algorithms presented, with as much as a 31.85% difference in the lowest utilization level and around 23%-28% at 50% utilization. Both versions performed similarly in this aspect, with EASS-DVFS performing slightly better in most cases. In terms of overhead, we observed higher values with EASS-DVFS in the average case, but EASS-DPM had higher values in extreme cases. However, even in the most extreme case, we observed the overhead to be only 1% of the slot length.

In this paper, we assumed VMs always execute until their worst-case execution time. However, this is usually not the case in actual practice. In the future, we will account for the actual execution time of VMs to save energy further. Moreover, we plan to combine the DPM and DVFS scheduler versions to achieve even better performance in terms of energy.

## REFERENCES

- [1] "Product SEcurity for Cross domain REliable Dependable Automated Systems." [Online]. Available: <https://secredas-project.eu/>
- [2] "Infrastructure in the cloud," accessed: 2023-01-09. [Online]. Available: <https://www.mobility.siemens.com/global/en/portfolio/rail/automation/infrastructure-in-the-cloud.html>
- [3] R. Hametner, P. Tummelshammer, S. Resch, and W. Wernhar, "Cloud architecture for sil4 railway applications," *Thales Austria GmbH, Signal + Draht - SIGNALLING & DATA COMMUNICATION magazine, Eurailpress Archiv - DVV Media Group GmbH*, March 2022.
- [4] Thales, SYSGO, Fraunhofer IESE, University of Rostock, ESE, and DB Netz, "Research report - sil4 cloud," accessed: 2023-01-09. [Online]. Available: <https://digitale-schiene-deutschland.de/Downloads/Report\%20-%20SIL4\%20Cloud.pdf>
- [5] "Scalable open architecture for embedded edge (soafee)," accessed: 2023-01-09. [Online]. Available: <https://www.soafee.io/>
- [6] J. Eker, O. Angelsmark, and A. Sefidcon, "What is the real-time cloud and how do we get there?" 11 2020. [Online]. Available: <https://www.ericsson.com/en/blog/2020/11/what-is-real-time-cloud>
- [7] M. Masmano, I. Ripoll, A. Crespo, and J. Metge, "Xtratium: a hypervisor for safety critical embedded systems," in *11th Real-Time Linux Workshop*. Citeseer, 2009, pp. 263–272.
- [8] M. Vanga, A. Gujarati, and B. B. Brandenburg, "Tableau: a high-throughput and predictable VM scheduler for high-density workloads," in *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, R. Oliveira, P. Felber, and Y. C. Hu, Eds. ACM, 2018, pp. 28:1–28:16. [Online]. Available: <https://doi.org/10.1145/3190508.3190557>
- [9] N. Studer, "What is the arinc653 scheduler?" Dec 2013, accessed: 2023-01-17. [Online]. Available: <https://xenproject.org/2013/12/16/what-is-the-arinc653-scheduler/>
- [10] Jan 2023. [Online]. Available: <https://www.linux-kvm.org/>
- [11] G. Gala, G. Fohler, P. Tummelshammer, S. Resch, and R. Hametner, "Rt-cloud: Virtualization technologies and cloud computing for railway use-case," in *2021 IEEE 24th International Symposium on Real-Time Distributed Computing (ISORC)*. IEEE, 2021, pp. 105–113.
- [12] L. Abeni and D. Faggioli, "Using xen and kvm as real-time hypervisors," *Journal of Systems Architecture*, vol. 106, p. 101709, 2020.
- [13] T. Deshane, Z. Shepherd, J. N. Matthews, M. Ben-Yehuda, A. Shah, and B. Rao, "Quantitative comparison of xen and kvm," 2008.
- [14] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," pp. 68–73, 2008.
- [15] G. Fohler, "Joint scheduling of distributed complex periodic and hard aperiodic tasks in statically scheduled systems," in *Proceedings 16th IEEE Real-Time Systems Symposium*. IEEE, 1995, pp. 152–161.
- [16] G. Gala, J. Castillo, and G. Fohler, "Work-in-progress: Cloud computing for time-triggered safety-critical systems," in *2021 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2021, pp. 516–519.
- [17] S. Schorr, "Adaptive real-time scheduling and resource management on multicore architectures," Ph.D. dissertation, Technische Universität Kaiserslautern, 2015.
- [18] J. Real, S. Sáez, and A. Crespo, "A hierarchical architecture for time- and event-triggered real-time systems," *Journal of Systems Architecture*, vol. 101, p. 101652, 2019.
- [19] F. F. Yao, A. J. Demers, and S. Shenker, "A scheduling model for reduced CPU energy," in *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*. IEEE Computer Society, 1995, pp. 374–382. [Online]. Available: <https://doi.org/10.1109/SFCS.1995.492493>
- [20] H. Aydin, R. G. Melhem, D. Mossé, and P. Mejía-Alvarez, "Determining optimal processor speeds for periodic real-time tasks with different power characteristics," in *13th Euromicro Conference on Real-Time Systems (ECRTS 2001), 13-15 June 2001, Delft, The Netherlands, Proceedings*. IEEE Computer Society, 2001, pp. 225–232. [Online]. Available: <https://doi.org/10.1109/EMRTS.2001.934038>
- [21] C. Lee and K. G. Shin, "On-line dynamic voltage scaling for hard real-time systems using the EDF algorithm," in *Proceedings of the 25th IEEE Real-Time Systems Symposium (RTSS 2004), 5-8 December 2004, Lisbon, Portugal*. IEEE Computer Society, 2004, pp. 319–327. [Online]. Available: <https://doi.org/10.1109/REAL.2004.38>
- [22] M. Bambagini, F. Prosperi, M. Marinoni, and G. C. Buttazzo, "Energy management for tiny real-time kernels," in *International Conference on Energy Aware Computing, ICEAC 2011, Istanbul, Turkey, November 30 - December 2, 2011*. IEEE, 2011, pp. 1–6. [Online]. Available: <https://doi.org/10.1109/ICEAC.2011.6136687>
- [23] Y. Lee, K. P. Reddy, and C. M. Krishna, "Scheduling techniques for reducing leakage power in hard real-time systems," in *15th Euromicro Conference on Real-Time Systems (ECRTS 2003), 2-4 July 2003, Porto, Portugal, Proceedings*. IEEE Computer Society, 2003, pp. 105–112. [Online]. Available: <https://doi.org/10.1109/EMRTS.2003.1212733>
- [24] M. A. Awan and S. M. Petters, "Enhanced race-to-halt: A leakage-aware energy management approach for dynamic priority systems," in *23rd Euromicro Conference on Real-Time Systems, ECRTS 2011, Porto, Portugal, 5-8 July, 2011*, K. Arzén, Ed. IEEE Computer Society, 2011, pp. 92–101. [Online]. Available: <https://doi.org/10.1109/ECRTS.2011.17>
- [25] H. Chen, X. Zhu, J. Zhu, and J. Wang, "ERES: an energy-aware real-time elastic scheduling algorithm in clouds," in *10th IEEE International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing, HPCC/EUC 2013, Zhangjiajie, China, November 13-15, 2013*. IEEE, 2013, pp. 777–784. [Online]. Available: <https://doi.org/10.1109/HPCC.and.EUC.2013.113>
- [26] X. Zhu, L. T. Yang, H. Chen, J. Wang, S. Yin, and X. Liu, "Real-time tasks oriented energy-aware scheduling in virtualized clouds," *IEEE Trans. Cloud Comput.*, vol. 2, no. 2, pp. 168–180, 2014. [Online]. Available: <https://doi.org/10.1109/TCC.2014.2310452>
- [27] J. Sun and H. Cho, "A lightweight optimal scheduling algorithm for energy-efficient and real-time cloud services," *IEEE Access*, vol. 10, pp. 5697–5714, 2022.
- [28] P. Karachatzis, J. Ruh, and S. S. Craciunas, "An evaluation of time-triggered scheduling in the linux kernel," in *Proceedings of the 31st International Conference on Real-Time Networks and Systems*, ser. RTNS '23. ACM, 2023.
- [29] G. Gala, I. Kadusale, and G. Fohler, "Joint time- and event-triggered scheduling in the linux kernel," in *Operating Systems Platforms for Embedded Real-Time applications (OSPERT) workshop 2023 co-hosted with 35th Euromicro conference on Real-time systems*, 2023.
- [30] E. Bini and G. C. Buttazzo, "Measuring the performance of schedulability tests," *Real-Time Systems*, vol. 30, no. 1, pp. 129–154, 2005.
- [31] "Power capping framework," <https://www.kernel.org/doc/html/latest/power/powercap/powercap.html>, accessed: 2022-03-10.