

1 ALLOCATION AND ANALYSIS

In this work, we consider partitioned scheduling. Each engine has its own scheduler and a separate ready-queue. Sub-tasks are allocated (partitioned) onto the available engines so that the system is schedulable. Partitioned scheduling allows to use the well-known single processor schedulability tests which make the analysis simpler and allows to reduce the overhead due to thread migration compared to global scheduling. The analysis presented here is modular, so engines may have different scheduling policies. In this paper, we restrict to preemptive-EDF.

Let assume n task specifications to be allocated onto m engines. Only concrete graphs can be allocated to engines. However, we often use *allocating task specification* to express the whole process of allocating one of its concrete digraphs.

1.1 Alternative patterns

Definition 1. Concrete task τ is schedulable if and only if all its sub-tasks at arrival j finish their execution within the time interval $[a_{ij}, a_{ij} + \mathcal{D}(i)]$, $\forall j \in N$ where a_{ij} is the j^{th} instance arrival time.

Definition 2. Task specification τ_i is schedulable, if:

$$\exists \tau_{id} \in \Omega_i, \tau_{id} \text{ is schedulable} \quad (1)$$

According to 2, it is sufficient to find one concrete task τ_d among all the concrete tasks of task specification τ , so it can be feasibly allocated. Thus, it is necessary to generate and test the allocation of all concrete tasks a given task specification. Algorithm shows how concrete tasks are generated.

Algorithm 1 generate_concrete_task

```

1: input:  $\tau$  Task
2:  $\text{alta\_node} = \text{select\_one\_alternative\_node}()$ 
3: if ( $\text{alta\_node} == \emptyset$ ) then
4:    $\text{gr\_list} = \text{explode\_node\_to\_graphs}()$ 
5:   for ( $\tau \in \text{gr\_list}$ ) do
6:      $\text{concrete\_task\_generation}(\text{task});$ 
7:   end for
8: else
9:    $\text{add\_to\_conc\_list}(\tau, \Omega)$ 
10: end if

```

Algorithm 1 takes as input one task specification, and selects one of its alternative nodes. Further, it explodes the nodes to several copies of the same graph by excluding all other alternative possibilities and calls back the same algorithm. The recursivity ends when a task specification does not contain any alternative control nodes, (becomes a concrete task).

Definition 3. Let $\tau_{d'}, \tau_{d''}$ be two concrete tasks of specification task τ

The order relation $>$ is defined as:

$$\tau_{d'} > \tau_{d''} \Rightarrow \sum_{v \in \tau_{d'}} C(v) \geq \sum_{v \in \tau_{d''}} C(v) \quad (2)$$

The order relation $>$ allows to sort concrete task according to their total execution time, thus allows preferring concrete tasks with less overhead. Further we will define

Algorithm 2 generate_task_separation

```

1: input:  $\tau$ : Task  $\triangleright$  specification or concrete
2: for ( $\text{tag} \in \text{tags}$ ) do
3:   for ( $v \in \mathcal{V}$ ) do
4:     if ( $v.\text{Tag} == \text{tag}$ ) then
5:        $\text{add\_to\_current\_separation}(v)$ 
6:     else
7:        $\text{add\_to\_current\_separation}(\emptyset)$ 
8:     end if
9:   end for
10: end for

```

relation order $>>$ to compare two concrete tasks based on the available engine number.

Order relations $>$, $>>$ will be used later later in Section 1.6 to select the concrete task to allocate for a given task specification when several concrete tasks can feasibility be allocated.

1.2 Graph separation

One concrete task may contain sub-tasks with different tags. Two sub-tasks with different tags are allocated onto different cores. We recall that in this paper we consider partitioned scheduling, analysis of each engine is independent. Hence, we need to *separe* vertices with the same tag before proceeding to its allocation. We call this operation separation.

Definition 4. Let $\tau_d = (\mathcal{V}, E)$ denote a concrete task. The set $\mathcal{S}(\tau_d) = \{\tau_{d_1}^{\text{tag}^1}, \tau_{d_2}^{\text{tag}^2}, \dots, \tau_{d_s}^{\text{tag}^s}\}$ is a separation of concrete task τ_d if and only if:

- Every separation $\tau_{d_s}^{\text{tag}^s}$ is an isomorphe graph regarding τ_d
- The number of tags of each separation is equal to **one**
- The union of the same vertex in all separations is equal to the original vertex. Thus, each vertex is either empty or is equal to the original vertex.

Each concrete task generates at least as separations as tags in the target architecture. Algorithm 2 shows how the minimum number of task separations is generated. The algorithm acts as a filter on each tag. The vertices having a different tag regarding the current one are set to empty. An empty vertex is a vertex having an execution time that equals $C(\emptyset) = 0$, otherwise the vertex is added to the current task separation.

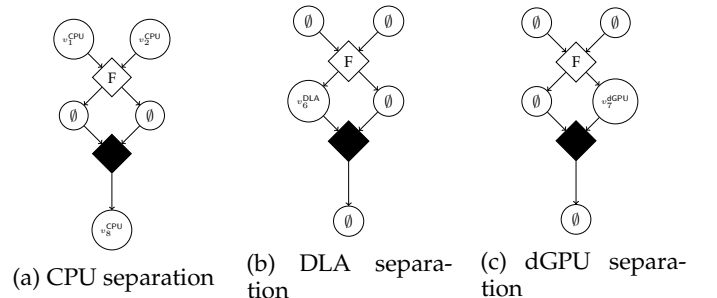


Fig. 1: minimal set of sepefication for concrete task at Figure

Figure 1 shows the generated task separations using Algorithm 2. As the concrete task has three tags, three task separations are generated. The first one contains only vertices having CPU tag, respectively the second the DLA and the third one the GPU. Each of the task separations will be allocated on one or multiple engines having the same tag.

Definition 5. Let $\tau_{d'}$, $\tau_{d''}$ be two concrete tasks of specification task τ . Let \mathbf{t}_\downarrow denote the list of ordered tag in an non-decreasing order of number of cores.

The order relation \gg is defined as:

$$\tau_{d'} \gg \tau_{d''} \Rightarrow \sum_{v \in \tau_{d'}(\mathbf{t}_\downarrow(0))} C(v) > \sum_{v \in \tau_{d''}(\mathbf{t}_\downarrow(0))} C(v) \quad (3)$$

1.3 Artificial deadlines and offsets

The respect of a real-time task constraints depends on the task allocation which is defined as the problem of allocating sub-tasks onto different engines. As these sub-tasks share buffers, they are forced to respect an execution order. Analyzing the behavior of dependant sub-tasks is complex. To reduce allocation complexity, lot of works have proposed to insert artificial offsets and deadlines to sub-tasks so they can be analyzed separately with respect to their initial execution order. We use similar techniques because the analysis is more complex in the presence of highly heterogeneous engines such as those we are interested in. Deadline assignment techniques are applied to concrete tasks because applying it to task specification may enforce some sub-task to have tight deadline because the critical path of a concrete path is shorter or equal to the critical path of its task specification.

Most of the deadline assignment techniques are based on the computation of the execution time of the critical path. A path $P_x = \{v_1, v_2, \dots, v_l\}$ is a sequence of sub-tasks of task τ such that :

$$\forall(v_l, v_{l+1}) \in P_x, \exists e(v_l, v_{l+1}) \in E$$

Definition 6. Let P_x denote the path x of task τ and Π denotes the set of all paths of task τ .

The critical path \mathcal{P} of task τ is define as:

$$\mathcal{P} = P_x \in \Pi, \max_x \left\{ \sum_{v \in P_x} C(v) \right\} \quad (4)$$

Informally, the critical path is defined as the path having the highest execution time. Let assume $\mathcal{P} = \{v_1, v_2, \dots, v_l\}$ denote the critical path of concrete task τ_d . We define slack time $S_{\mathcal{P}}$ of task τ_d as the following:

$$S(\mathcal{P}) = \mathcal{D}(\tau_d) - \sum_{v \in \mathcal{P}} C(v) \quad (5)$$

Further, to assign artificial deadline, $S(\mathcal{P})$ is distributed on different vertices as following:

$$D(v) = C(v) + \text{calculate_share}(v, \mathcal{P}) \quad (6)$$

The calculate share function computes the slack for every single vertex. This slack can be shared according to two heuristics:

- **Fair distribution:** assigns vertex slack as the ratio of the original slack $s(\mathcal{P})$ by the number of vertices in the path ($|\mathcal{P}|$) (Equation ((7)))

$$\text{calculate_share}(v, \mathcal{P}) = \frac{S(\mathcal{P})}{|\mathcal{P}|} \quad (7)$$

- **Proportional distribution:** assigns single slack according to the contribution of the vertex execution time in the path length, the the share is computed as in Equation (8).

$$\text{calculate_share}(v, \mathcal{P}) = \frac{C(v)}{C(\mathcal{P})} \cdot S(\mathcal{P}) \quad (8)$$

Recurrently, other paths are evaluated according to an non-decreasing order of their vertice execution time. Already assigned deadlines are not assigned again and their respective slack is substracted from the original slack of each path $P_x \in \Pi$.

Further offsets of each vertex are computed as:

$$O(v) = \begin{cases} 0, & \text{preds}(v) = \emptyset \\ \max\{O(v_{pr}) + C(v_{pr})\}, & pr \in \text{preds}(v) \text{ Otherwise} \end{cases} \quad (9)$$

Definition 7. Vertex $v \in \mathcal{V}_\tau$ is feasible, if each arrival j finishes its execution within the interval bounded by its arrival time $a(v) = j \cdot T(\tau) + O(v)$ and its deadline $a(v) + D(v)$.

Theorem 1. A concrete task (resp. task separation) is feasible, if all vertices within the task are feasible.

Proof. It is easy to proof that exit-sub tasks have an absolute deadline not greater than the task absolute deadline, thus if exit sub-tasks are feasible (by assumption), thus task is feasible. \square

According to Theorem 1, the analysis of every sub-task can be done independently from each other, if they all respect their arrival time and deadlines.

Theorem 2. Let τ_d be a concrete digraph and $\mathcal{S}(\tau_d)$ be the set of its task separations.

If $\forall \tau_{d_s} \in \mathcal{S}(\tau_d)$ is feasible, τ is feasible

Proof. First, we would like to recall that vertices having an empty vertex as predecessor are forced to take its deadline and offset into account in their deadline/offset computation (as shown in Section 1.3).

If all task separation are feasible, it means that even with their blocking time due to an empty vertex, the exit nodes finish before the concrete task relative deadline $\mathcal{D}(i)$. Thus, the original concrete task is feasible. \square

1.4 Allocation of task specifications

1.5 Single core analysis

1.6 Schedulability analysis and on-line conditional graphs

Algorithm 3 Full algorithm

```

1: for  $\tau_i \in \mathcal{T}$  do
2:    $\Omega_i = \text{generate\_concrete\_task}(\tau_i)$ 
3:   for ( $\tau_{id} \in \Omega_i$ ) do
4:      $\text{assign\_deadlines\_offset}(\text{conc\_list})$ 
5:      $\mathcal{S}(\tau_{id}) = \text{generate\_task\_separation}(\Omega_i)$ 
6:     if (then  $\text{feasible\_sequential}(\mathcal{S}(\tau_{id}))$ )
7:        $\text{add\_to\_feasible\_list}(\text{f\_list})$ 
8:     end if
9:   end for
10:  if ( $|\text{f\_list}| > 0$ ) then
11:     $\text{sort}(\text{f\_list})$   $\triangleright$  Sort f_list using to ">" or ">>"
12:    Allocate f_list(0) to selected_engines
13:  else
14:     $\text{sort}(\Omega_i)$   $\triangleright$  sort according to order relations
15:     $(\tau_{id'}, \tau_{id''}) = \text{paralleliz\_concrete}(\Omega_i(0))$ 
16:    if (then  $\tau_{id'} \neq 0$ )
17:      allocate
18:         $\text{to\_selected\_cores}$  else
19:        return FAIL
20:    end if
21:

```
