

Bringing Inter-Thread Cache Benefits to Federated Scheduling

Corey Tessler, Venkata P. Modekurthy, Nathan Fisher and Abusayeed Saifullah
Department of Computer Science, Wayne State University, Detroit, MI, USA

Abstract—Multiprocessor scheduling of hard real-time tasks modeled by directed acyclic graphs (DAGs) exploits the inherent parallelism presented by the model. For DAG tasks, a node represents a request to execute an object on one of the available processors. In one DAG task, there may be multiple execution requests for one object, each represented by a distinct node. These distinct execution requests offer an opportunity to reduce their combined cache overhead through coordinated scheduling of objects as threads within a parallel task. The goal of this work is to realize this opportunity by incorporating the cache-aware BUNDLE-scheduling algorithm into federated scheduling of sporadic DAG task sets.

This is the first work to incorporate instruction cache sharing into federated scheduling. The result is a modification of the DAG model named the *DAG with objects and threads* (DAG-OT). Under the DAG-OT model, descriptions of nodes explicitly include their underlying executable object and number of threads. When possible, nodes assigned the same executable object are collapsed into a single node; joining their threads when BUNDLE-scheduled. Compared to the DAG model, the DAG-OT model with cache-aware scheduling reduces the number of cores allocated to individual tasks by approximately 20 percent in the synthetic evaluation and up to 50 percent on a novel parallel computing platform implementation. By reducing the number of allocated cores, the DAG-OT model is able to schedule a subset of previously infeasible task sets.

I. INTRODUCTION

For hard real-time parallel tasks, where the total execution demand of a task may exceed its deadline, federated scheduling [9, 10, 26, 60] provides a method for executing each task across multiple cores and an accompanying analysis which determines if all tasks will always meet their deadlines. To analyze and schedule parallel tasks, each task is represented by a directed acyclic graph (DAG).

Nodes within a DAG represent the release and complete execution of an object upon one of the m identical cores of the system. Edges between nodes indicate precedence constraints between nodes; a node may not begin executing until all predecessors have completed. Associated with every node is a worst-case execution time (WCET) bounding any complete execution. Schedulability analysis of a task's DAG depends on the task's workload (sum of WCETs of all nodes) and critical path length (path through the DAG with greatest WCET).

Worst-case execution time calculation accounts for architecture features including cache memory. The variability in execution times due to cache memory has been well studied for uniprocessor single-threaded task sets in works such as [3, 4, 17, 25, 28, 40, 51, 58, 65]. Scheduling of tasks on multi-processor systems with cache memory has been

studied in works such as [14, 15, 18, 30, 44, 63, 64]. In most previous work on both multiprocessor and uniprocessor real-time systems, cache memory contributes primarily negatively to schedulability by increasing WCET values. Preemptions between jobs introduce cache-related preemption delays (CRPD) for both uniprocessor and multi-processor systems. Multiprocessor multi-threaded systems with shared caches are also affected by evictions from concurrent execution as well as cache coherency delays across cores [14, 15, 18, 44, 63, 64]. The method proposed in this work is the first to incorporate instruction cache reuse beneficially into real-time scheduling decisions for federated scheduling of multi-processor multi-threaded systems.

In the setting of uniprocessor multi-threaded task systems, the BUNDLE-based approaches [54–56] (referred to as BUNDLE throughout the rest of this work) treat cache memory positively, creating a benefit to schedulability. BUNDLE restricts the execution of the multiple threads of task on a single processor in a cache cognizant manner. This restricted execution allows the sharing of cached values to be quantified as the inter-thread cache benefit. The accompanying BUNDLE analysis incorporates the inter-thread cache benefit into a WCET function for each task. These WCET functions accept the number of threads released with each job and quantifies the total benefit of “bundling” the threads together.

BUNDLE is limited to single processor multi-threaded tasks. The inter-thread cache benefit applies exclusively to instruction caches. Furthermore, BUNDLE's scheduling and WCET analysis techniques are limited to a single executable object. As such, BUNDLE is not directly applicable to parallel DAG tasks utilized by federated or global multi-processor schedulers.

This work incorporates cache memory positively into multiprocessor parallel tasks by joining BUNDLE's analysis and scheduling techniques to those of federated scheduling. This is achieved by treating executable objects (nodes) of parallel DAG tasks as threads scheduled by BUNDLE. Each individual node of a DAG task represents a single thread of execution of the underlying object. Nodes sharing the same underlying object may be *collapsed* into a single node and the combined threads scheduled by BUNDLE.

The purpose of collapse is to increase schedulability by reducing the number of processors dedicated to high utilization tasks. However, collapse may not be arbitrarily applied to nodes of the same object. There are several challenges when considering collapsing two nodes, doing so may:

- Introduce a cycle to the DAG

- Produce an infeasible parallel task
- Increase the number of cores allocated to a task

To achieve the goal of reducing the number of cores allocated to high utilization tasks while carefully selecting which nodes to collapse the following contributions are made:

- A modification to the DAG model named DAG-OT, the first parallel task model to include inter-thread cache benefits
- The concepts of collapse and collapse candidacy
- An algorithm for collapsing nodes
- Heuristics for ordering collapse of nodes within a task
- An evaluation of synthetic tasks demonstrating the benefits of collapse and BUNDLE scheduling of nodes under a federated scheduler
- A feasibility study with a parallel DAG-OT task scheduler operating on physical hardware demonstrating the potential cache benefits

These contributions are made in the following sections. Section II expands upon the related work. Section III describes BUNDLE-based scheduling, the existing DAG model, and the proposed DAG-OT model. Section IV describes the collapse operation and its impact. Section V introduces the general algorithm for collapsing nodes. Section VI describes the proposed heuristics used to order candidates for collapse. Section VII describes the collapse of low utilization tasks. Section VIII describes the schedulability for tasks of the proposed model. Section IX describes the methods, metrics, and results of the synthetic evaluation. Section X presents the feasibility study and results. Section XI concludes the work.

II. RELATED WORK

Parallel hard real-time DAG tasks may be scheduled by federated [10, 26, 36, 50, 60] or global [11, 24, 46, 47] policies. Federated scheduling improves the analytical bounds of global scheduling by dedicating cores to tasks that require more than one core to meet their deadlines.

For multi-processor systems, the impact of cache memory focuses on shared caches. When caches are shared between cores, an object executing on one core may evict values placed there by another object on a distinct core – increasing execution times. There are several works dedicated to mitigating or providing bounds on the number of evictions with global scheduling policies, including [14, 15, 63, 64]. It should be noted the tasks in [63, 64] are not parallel tasks. Cache coherency and false sharing [18, 30, 44] is another source of execution time extension for parallel tasks running on a multi-processor system with shared caches.

In the setting of uniprocessor systems executing single-threaded tasks, cache memory has been well studied. WCET analysis accounting for cache reuse of single-threaded tasks and direct map caches is presented in [5, 38, 39] and expanded to set-associative caches in [29]. Addressing the impact of cache memory in a preemptive setting is the purpose of CRPD analysis [3, 23, 25, 32–35, 49, 51, 58].

Each of the CRPD analytical methods seek to accurately estimate the impact of preemptions upon the WCET bound

for single-threaded tasks. Other approaches seek to mitigate CRPD’s run-time impact. The PREM [6, 41, 62] model divides tasks into load and execute phases, preventing cross tasks cache interference. Explicit preemption point placement [12, 13, 25, 48, 61] limits when a job may preempt another based upon the cache impact it may have.

Caches receive a common treatment in the uniprocessor works related to CRPD and multi-processor works related to shared caches. Specifically, caches are seen from the negative perspective, exclusively detracting from schedulability analysis. The only works the authors’ are aware of that take a positive perspective are *persistent cache blocks* [42, 43] in the uniprocessor single-threaded setting and *cache spread* [16] in the multi-processor setting utilizing global scheduling.

The work proposed herein applies to the federated scheduling of hard real-time parallel DAG tasks on multi-processor systems, focusing on the impact of scheduling decisions in the presence of dedicated (not shared) caches. To the authors knowledge, there are no existing works that address this setting. Furthermore, the proposed approach treats instruction caches positively by decreasing the number of cores dedicated to high utilization tasks in an attempt to increase schedulability.

A positive perspective of caches is taken by BUNDLE [54–56] for multi-threaded tasks running on a single processor. This positive perspective is reflected in the WCET of a task which includes the *inter-thread cache benefit*: the speed-up one thread experiences by another placing values in the cache. The benefit is restricted to instruction caches when threads are scheduled by the BUNDLE scheduling algorithm.

BUNDLE analysis and scheduling are central to the combined scheduling approach proposed in Section III-D, depending upon the BUNDLE calculated WCET values for collapse decisions. As a product of BUNDLE analysis, each multi-threaded task is assigned a worst-case execution time and cache overhead (WCETO) function $c(\eta) : \mathbb{N}^+ \rightarrow \mathbb{R}$. Where $c(\eta)$ is an upper bound on the amount of time required to execute η threads by BUNDLE, which encapsulates the inter-thread cache benefit. The result is a strictly increasing concave function with respect to η . Summarily, for $\eta + 1$ threads, $c(\eta + 1) \leq c(\eta) + c(1)$. By combining BUNDLE and federated scheduling the concave property of BUNDLE analysis can be leveraged to increase the schedulability of parallel tasks.

III. SYSTEM MODEL

This work proposes changes to the parallel DAG model [27] of hard real-time tasks to support collapse operations. The decision to collapse nodes is a scheduling decision that depends upon the BUNDLE’d execution of combined threads upon a single core. The purpose of this section is to describe the BUNDLE model in Subsection III-A, summarize the DAG model in Subsection III-B and federated scheduling in Subsection III-C, describe the proposed model which combines federated and BUNDLE scheduling in Subsection III-D, and lastly illustrate the impact of collapse under the combined model.

A. BUNDLE

Tasks in the BUNDLE [54–56] model are represented by a tuple $\tau_i = (p_i, d_i, c_i(\eta_i), \eta_i, o_i)$. A task has the familiar minimum inter-arrival time p_i and relative deadline d_i . A task has an underlying executable object o_i , and a number of threads released per job η_i . With each job release of τ_i , η_i threads are simultaneously released and must complete before the relative deadline d_i . The task’s WCETO is given by $c_i(\eta_i)$ which provides an upper bound to complete all η_i threads when scheduled in BUNDLE’s cache cognizant manner.

To produce the WCETO for a task, the control flow graph of the executable object is divided into conflict free regions: sub-graphs of the control flow graph where no two instructions map to the same cache block. Conflict free regions serve as input to the BUNDLE scheduling algorithm, which maximizes the number of threads executing over each region [55] in turn maximizing the inter-thread cache benefit.

Analysis of jobs scheduled by BUNDLE incorporates the inter-thread cache benefit in the task’s WCETO function $c(\eta)$ for $\eta \in \mathbb{N}^+$ threads. Every increase in η maximizes the contribution of an individual thread. The result is that $c(\eta)$ is a discrete concave function. Consider the addition of a single thread $c(\eta + 1)$ compared to the addition of two threads $c(\eta + 2)$. The WCETO increase of $c(\eta)$ to $c(\eta + 1)$ must be greater than or equal to the increase from $c(\eta + 1)$ to $c(\eta + 2)$: $c(\eta + 1) - c(\eta) \geq c(\eta + 2) - c(\eta + 1)$. If it were not, the increase of $c(\eta + 1)$ would not be maximal. Furthermore, if the increase of one thread were less than that of a second thread the bound of $c(\eta + 1)$ would be optimistic and unsafe. Thus, for any $\eta_a < \eta_b < \eta_c$ the point $(\eta_b, c(\eta_b))$ lies above the line defined by $(\eta_a, c(\eta_a))$ and $(\eta_c, c(\eta_c))$, therefore $c(\eta)$ is concave. Multi-threaded programs executed by BUNDLEP [55] illustrate the discrete concave growth described by the analysis.

B. DAG Model

Tasks in the parallel DAG model [27] are represented by a tuple $\tau_i = (T_i, D_i, G_i)$ of minimum inter arrival time T_i , implicit deadline $D_i = T_i$ and directed acyclic graph $G_i = (V_i, E_i)$. The set of n tasks is given by $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$. The set of all DAGs is denoted $\mathbb{G} = \{G_1, G_2, \dots, G_n\}$.

Within a DAG G_i , a node $v \in V_i$ represents the execution of a single thread. A thread executes on exactly one of the m cores of the target architecture (or distributed system). Each node is implicitly associated with an underlying executable object α_v : a set of machine instructions reachable from a single entry point. A worst-case execution c_v time is associated with every node v ; an upper bound on the execution time required to complete the thread without interruption on a single core. An edge $(u, v) \in E_i$ indicates an execution dependency between $u, v \in V_i$. For v to begin execution on any core, all immediate predecessors $\{u | (u, v) \in E_i\}$ must run to completion.

For simplicity of analysis, every DAG G_i must have exactly one source and sink node, $s, t \in V_i$ respectively. A source s has no incoming edges, $\nexists u | (u, s) \in E_i$. A sink t has no outgoing edges, $\nexists v | (t, v) \in E_i$. Without loss of generality, when a DAG contains multiple sources, the DAG is augmented

by adding an “empty source”: a single node with zero execution cost that is connected by outgoing edges to existing sources. Similarly, for a DAG with multiple sinks an “empty sink” is added with zero execution cost connected by incoming edges from the existing sinks.

Jobs of a task begin with one thread of s on one core. Jobs terminates when the single thread of t completes. During the execution of a job, up to m cores may execute any of the $v \in V$ threads in parallel. A task $\tau_i \in \tau$ generates a potentially infinite number of jobs, each arriving no less than T_i time units apart. All jobs of τ_i must complete within $D_i = T_i$ time units.

An example DAG task is shown in Figure 1. Accompanying each node is a single-threaded WCET. For u and v , their WCET values are $c_u = 20$ and $c_v = 10$ respectively. Edges illustrate the dependency order of execution, such as (s, v) precluding v from executing until s has completed.

For a DAG $G_i = (V_i, E_i)$, the length of a path through the graph is the sum of WCET values of all nodes along the path. The *critical path* λ_i of G_i , is a path from s to t with the greatest length L_i called the *critical path length*. If there are multiple paths with equal length L_i , only one is selected as the critical path. The *workload* of G_i is the sum of all WCET values $v \in V_i$. *Utilization* of the task τ_i is the ratio of its workload and minimum inter-arrival time.

| | |
|--|--------------------------------------|
| Critical Path Length of G_i | Workload of G_i |
| $L_i = \sum_{v \in \lambda_i} c_v$ (1) | $C_i = \sum_{v \in V_i} c_v$ (2) |
| Utilization of G_i | Utilization of τ |
| $u_i = C_i/T_i$ (3) | $U = \sum_{\tau_i \in \tau} u_i$ (4) |

TABLE I: Definitions for Parallel DAG Task Sets

In Figure 1, the critical path $\lambda = \langle s, u, t \rangle$ is highlighted. The calculated critical path length is $L = c_s + c_u + c_t = 60$ and workload $C = c_s + c_u + c_v + c_t = 70$.

C. Federated Scheduling

Federating scheduling [27] is a partitioned scheduling algorithm variant and analysis method developed for parallel DAG task sets. It divides the task set τ into two disjoint sets. Tasks with utilization greater than one are placed in the *high utilization task set* τ_{high} .

The *low utilization task set* τ_{low} contains the remainder of τ . Every task τ_i of τ_{high} is assigned m_i dedicated cores, where m_i is given by Equation 5. Only threads of τ_i may execute on the m_i cores dedicated to it. All jobs of a high utilization task τ_i scheduled on m_i cores are guaranteed to meet their deadlines [27].

$$m_i = \left\lceil \frac{C_i - L_i}{D_i - L_i} \right\rceil \quad (5)$$

Fig. 2: m_i for $\tau_i \in \tau_{high}$




Fig. 1: A DAG Task

The number of cores allocated to all high utilization tasks is denoted $m_{high} = \sum_{\tau_i \in \tau_{high}} m_i$. The remaining cores of low utilization tasks are denoted $m_{low} = m - m_{high}$. A task set τ is schedulable under federated scheduling if m_{low} is non-negative and all tasks of τ_{low} are partitioned on the m_{low} processors while meeting their deadlines when threads within jobs are scheduled sequentially.

Any greedy, work-conserving, parallel scheduler may be used to schedule a high utilization task $\tau_i \in \tau_{high}$ on its m_i dedicated cores. Low utilization tasks are treated as sequential tasks, executing at most one thread of a job at a time. Any multiprocessor scheduling algorithm (such as partitioned EDF) may be used to schedule all the low utilization tasks on the m_{low} allocated cores.

D. Proposed Model: DAG-OT

The model proposed in this work augments the existing DAG model by explicitly including the implicit number of threads and executable objects associated with every node $v \in V$. Doing so requires modification to the WCET of a node, converting the static value c_v to a function in terms of the number of threads executed. For clarity, the existing model is referred to as the directed acyclic graph model of parallel tasks or simply “the DAG model”, the proposed model is named the DAG with objects and threads or “the DAG-OT model”.

For a DAG in the DAG model $G_i = (V_i, E_i)$, two distinct nodes $u, v \in V_i$ represent the release of one thread of execution over their underlying executable objects α_u and α_v . There is no restriction on the relationship between α_u and α_v , they may be distinct or identical objects. The first proposed change to the DAG model is to explicitly include the executable object in a node’s description.

Similarly, for a node $v \in V_i$ in the DAG model, the execution of a single thread is bounded by a single WCET value c_v . The second proposed change to the DAG model is to explicitly include the number of threads η_v and present the WCET of a node as function in terms of the number of threads executed $c_v(\eta) : \mathbb{N}^+ \rightarrow \mathbb{R}^+$.

Combining the proposed changes, a node $v \in V_i$ in the DAG-OT model is represented by a tuple $v = \langle \alpha_v, c_v(\eta), \eta_v \rangle$. Figure 3 presents the differences between the DAG and DAG-OT models visually. A consistent illustrative shorthand is used in this work, with the order of nodes tuple’s preserved and the critical path highlighted in gray.




Fig. 3: From DAG to DAG-OT

Nodes of the DAG-OT model are compatible with nodes of the DAG model [27], where nodes from the DAG model can

be expressed as $v = \langle \alpha_v, c_v(\eta), \eta_v = 1 \rangle$ under DAG-OT. This is illustrated by Figures 3a and 3b, which are equivalent.

The motivation for including the executable object, threads, and WCET as a function in the description of a node is to satisfy the BUNDLE model and facilitate the combined scheduling technique. Combining the federated and BUNDLE scheduling techniques, each node is treated as single unit of execution to be BUNDLE scheduled upon one core. Each node requires $\eta_v \geq 1$ threads of object α_v to be executed by BUNDLE.

Under the DAG-OT model, when a node $v \in V_i$ is selected for execution all η_v threads of the object α_v are executed and scheduled by BUNDLE on one core. The total execution required to complete all threads is bounded by the WCETO function provided by BUNDLE analysis and associated with the node as $c_v(\eta)$.

Under federated scheduling (and in this work) DAG tasks execute on a parallel system with m identical cores. Requiring uniform cores ensures the validity of the WCET bound for each node regardless of which core the thread executes upon. Furthermore, this work requires each core to have identical cache configurations (hierarchy, size, etc.), memory architecture, and be timing-compositional [21]. Doing so guarantees the worst-case execution time and cache overhead (WCETO) of every node will be consistent across all cores. BUNDLE WCETO analysis is limited to the per-core dedicated instruction caches. Data caches, and cache memory shared between cores are not considered and are reserved for future work.

For the DAG-OT model, the definitions of critical path length and workload must be updated, given by Equations 6 and 7.

Definition 1 (DAG-OT Critical Path Length of G_i).

$$L_i = \sum_{v \in \lambda_i} c_v(\eta_v) \quad (6)$$

Definition 2 (DAG-OT Workload of G_i).

$$C_i = \sum_{v \in V_i} c_v(\eta_v) \quad (7)$$

E. Growth Factors

For simplicity of presentation and analysis the WCETO function $c_u(\eta)$ for a node u is described by a *growth factor* \mathbb{F}_u . A growth factor upper bounds the discrete concave WCETO of a node by a linear function using the single threaded ($c_u(1)$) value at the cost of increased pessimism. This simplification is leveraged in the evaluation.

Definition 3 (Growth Factor \mathbb{F}). For a node $u \in V$ of a DAG $G_i = (V, E)$, the growth factor of u is a value $\mathbb{F}_u \in (0, 1]$ that satisfies Equation 8 for all $\eta_u \geq 1$.

$$c_u(\eta_u) \leq c(1) + \mathbb{F}_u \cdot (\eta_u - 1) \cdot c_u(1) \quad (8)$$

An example for a node u , associated $c_u(\eta_u)$, and growth factor $\mathbb{F}_u = .5$ is shown in Figure 4. The values of $c_u(\eta_u)$ are 10, 15, 17, 18, 19 for $\eta_u \in [1, 5]$. While any growth factor greater than .5 would satisfy the definition, the minimum was selected for illustrative purposes.




Fig. 4: Example Growth Factor

IV. COLLAPSING NODES

To bring the inter-thread cache benefit to the DAG-OT model, this work proposes the concept of *collapsing* nodes. Under the DAG-OT model, two nodes $u, v \in V_i$ which execute the same object $\alpha_u = \alpha_v$ may potentially be combined into a single node are referred to as *candidates*. Candidates feature prominently in the fork-join [31] and MapReduce [19] parallel task models; which are restricted instances of parallel DAG tasks. A general DAG task may include fork-join or MapReduce sub-graphs as part of the task’s graph. Collapsing two nodes into a single node turns two distinct execution requests executing on (possibly) distinct cores, into a single request to execute the combined threads on one core using BUNDLE scheduling. By virtue of BUNDLE’s analysis incorporating the inter-thread cache benefit, the WCETO of the combined node may be less than the sums of the individual nodes.

Definition 4 (Candidate for Collapse). *For a DAG $G_i = (V, E)$ and nodes $u, v \in V$, u and v are candidates for collapse if and only if they share an executable object $\alpha_u = \alpha_v$.*

To illustrate, consider Figure 5. Nodes u and v share the same executable object α_1 . If the WCETO of one thread scheduled by BUNDLE on one core is 10 and two is 12, two nodes executing on separate cores will require 20 total cycles. Collapsing u and v , requiring the two threads to be scheduled in a cache cognizant manner on one core by BUNDLE reduces the workload (and potentially the critical path length) by 8.




Fig. 5: Node Collapse

Collapse restricts the execution of threads and cores. In Figure 5 pre-collapse u and v may have executed on distinct cores. Post-collapse the combined threads of u and v must execute on the same core scheduled by BUNDLE. To differentiate between pre and post-collapse values a “hat” will be used for the latter. In Figure 5, before collapse u and v each execute one

thread. Collapsing the two nodes into \hat{u} joins the two threads $\eta_{\hat{u}} = 2 = \eta_u + \eta_v$. The pre-collapse workload is $C_i = 43$ and post-collapse workload $\hat{C}_i = 35$. The reduction in workload is due to the concave WCETO function $c_u(\eta) = c_v(\eta) = c_{\hat{u}}(\eta)$, where $c_u(1) = 10$ and $c_u(2) = 12$.

Formally, the collapse operation is defined as follows.

Definition 5 (Collapse $\hat{u} \leftarrow u \bowtie v$). *For pre-collapse nodes $u, v \in V$ of $G_i = (V, E)$, collapsing u and v (denoted $u \bowtie v$) into \hat{u} , resulting in a new DAG named $\hat{G}_i = (\hat{V}, \hat{E})$ where:*

$$\eta_{\hat{u}} \leftarrow \eta_u + \eta_v \quad (9)$$

$$\alpha_{\hat{u}} \leftarrow \alpha_u \quad (10)$$

$$c_{\hat{u}} \leftarrow c_u \quad (11)$$

$$\hat{V} \leftarrow \hat{u} \cup V \setminus \{u, v\} \quad (12)$$

$$\hat{E} \leftarrow \{(\hat{u}, y) | (u, y) \in E \vee (v, y) \in E\} \quad (13)$$

$$\cup \{(x, \hat{u}) | (x, u) \in E \vee (x, v) \in E\}$$

$$\cup E \setminus \{(x, y) | x \in \{u, v\} \vee y \in \{u, v\}\}$$

Equation 9 joins the threads of u and v to \hat{u} . Equation 10 and 11 assigns the executable object and WCETO function shared between u and v to \hat{u} . Equation 12 takes the nodes from G_i and copies them to \hat{G}_i , removing u and v . Similarly, Equation 13 copies the edges of G_i to \hat{G}_i while removing incoming and outgoing edges of u and v replacing them with incoming and outgoing edges of \hat{u} .

A. Infeasibility and the Impact of Collapse

Collapsing nodes may reduce the critical path length L_i . This is illustrated by Figure 6, where the pre-collapse critical path length is $L_i = 50$. After collapsing $\hat{u} \leftarrow u \bowtie v$, the critical path length of \hat{G}_i is $\hat{L}_i = 40$.




Fig. 6: Critical Path Reduction


Observation 1 (Critical Path Reduction). *For a DAG $G_i = (V, E)$ and candidate nodes $u, v \in V$, the collapse of $u \bowtie v$ into \hat{u} may reduce the critical path length in \hat{G}_i : $\hat{L}_i \leq L_i$.*

Under the DAG model, a task τ_i is infeasible (for any number of dedicated cores m_i) if the critical path length is greater than the deadline, i.e., $L_i > D_i$. A task τ_i deemed infeasible due to critical path length and period under the DAG model ($L_i > D_i$) may become feasible (and possibly schedulable) under the DAG-OT model by collapse and Observation 1 ($\hat{L}_i \leq D_i$). Thus the $L_i > D_i$ infeasibility test does not apply pre-collapse to the DAG-OT model. However, for any post-collapse \hat{G}_i of τ_i if $\hat{L}_i > D_i$ the task set is unschedulable under DAG-OT.

Observation 2 (Critical Path Extension). *For a DAG $G_i = (V, E)$ and candidates nodes $u, v \in V$, the collapse*

of $u \bowtie v$ into \hat{u} may extend the critical path length in \hat{G}_i : $\hat{L}_i \geq L_i$.

In contrast to Observation 1, collapse may extend the critical path length. This can occur when one of the candidate nodes $u, v \in V$ lies on the pre-collapse critical path and the other does not. In Figure 7, u lies on the pre-collapse critical path. Collapsing $\hat{u} \leftarrow u \bowtie v$ increases the critical path length \hat{L}_i compared to L_i by $c_u(\eta_u + \eta_v) - c_u(\eta_u)$.



(a) Pre-Collapse $L_i = 34$ (b) Post-Collapse $\hat{L}_i = 38$


Fig. 7: Critical Path Extension

Observation 3 (Workload Decrease). For a DAG $G_i = (V, E)$ and candidate nodes $u, v \in V$, the collapse of $u \bowtie v$ into \hat{u} reduces the workload, i.e., $\hat{C}_i \leq C_i$.

For candidates $u, v \in V$, their contribution to the workload of C_i is $c_u(\eta_u) + c_v(\eta_v)$. The contribution of $\hat{u} \leftarrow u \bowtie v$ to \hat{C}_i is $c_{\hat{u}}(\eta_{\hat{u}}) = c_u(\eta_u + \eta_v)$. Since, $c_u(\eta)$ is a concave function, $c_u(\eta_u + \eta_v) \leq c_u(\eta_u) + c_v(\eta_v)$ and $\hat{C}_i \leq C_i$.

Observation 4 (Collapse Occlusion). For a DAG $G_i = (V, E)$, candidates (u, v) and (x, y) , the collapse of $u \bowtie v$ may prevent the collapse of $x \bowtie y$.

Collapsing one candidate (u, v) may preclude the collapse of another. For example, consider Figure 8. By collapsing (u, v) the pair (x, y) cannot be collapsed – doing so would introduce a cycle into the DAG.



(a) Pre-Collapse (b) Post-Collapse

Fig. 8: Collapse of (u, v) before (x, y)


| C_i | L_i | m_i | $u \bowtie v$ | \hat{C}_i | \hat{L}_i | \hat{m}_i |
|-------|-------|-------|---------------|-------------|-------------|-------------|
| 52 | 32 | [2.5] | \rightarrow | 50 | 33 | [2.42] |

TABLE II: Collapse of u and v from Figure 8

Given a deadline $D_i = 40$ the result of collapsing (u, v) with respect to the workload, critical path length, and dedicated cores are summarized in Table II.

Observation 5 (Alternate Collapse may Decrease \hat{m}). For a DAG $G_i = (V, E)$, candidates (u, v) and (x, y) , the collapse of $u \bowtie v$ which occludes $x \bowtie y$ and resulting allocation of

cores denoted $\hat{m}_{(u \bowtie v)}$ may be greater than the allocation of cores due to collapsing $x \bowtie y$, i.e., $\hat{m}_{(x \bowtie y)} < \hat{m}_{(u \bowtie v)}$.



(a) Pre-Collapse (b) Post-Collapse

Fig. 9: Collapse of (x, y) before (u, v)

Continuing the example, collapsing (x, y) precludes the collapse of (u, v) . Collapsing (x, y) instead of (u, v) is shown in Figure 9. The impact upon the workload and critical path length of $x \bowtie y$ differs from that of $u \bowtie v$ and ultimately a difference in m .

| C_i | L_i | m_i | $x \bowtie y$ | \hat{C}_i | \hat{L}_i | \hat{m}_i |
|-------|-------|-------|---------------|-------------|-------------|-------------|
| 52 | 32 | [2.5] | \rightarrow | 49 | 29 | [1.81] |

TABLE III: Collapse of x and y from Figure 9

Table III illustrates the impact of ordering of collapse with respect to m , where collapsing $x \bowtie y$ in place of $u \bowtie v$ yields a smaller number of dedicated cores m .

B. Beneficial Collapse

By Observations 1-5, collapsing any individual candidate may increase or decrease the number of cores allocated to a task. A collapse may increase or decrease the critical path length creating an infeasible task set or introduce a cycle into the graph. This subsection defines which collapses are *beneficial*.

Beneficial collapse depends on the Definition 6 of *improving* the allocation of cores. Improving the number of allocated cores balances the concepts of reducing the number of cores allocated to a feasible task, avoiding the creation of an infeasible task, and (possibly) creating feasible tasks from infeasible ones.

Definition 6 (Improved Core Allocation). For a given number of cores allocated to a task m_i , \hat{m}_i is an improvement upon m_i denoted $\hat{m}_i \ll m_i$ if and only if:

- 1.) $m_i > 0 \Rightarrow 0 < \hat{m}_i \leq m_i$
- 2.) $m_i \leq 0 \Rightarrow \hat{m}_i \geq m_i$

When m_i is greater than zero, an \hat{m}_i less than m_i and greater than zero is an improvement, reducing the number of cores allocated to the task. When $m_i < 0$, the critical path length has exceeded the deadline $L_i > D_i$. Such a task is not feasible under the DAG model. For m_i less than zero, a \hat{m}_i greater than m_i is an improvement; an increase over m_i may result in a schedulable task under DAG-OT.

Improvement of m_i does not include the ceiling described by Equation 5. This is due to the difference in context of m_i under the DAG model compared to DAG-OT. For the DAG model, m_i is calculated once and an integer number of cores are assigned to the task τ_i for schedulability analysis. For the DAG-OT model, m_i is recalculated after every collapse

operation. Only when collapse operations have ceased is the final integer ceiling of \hat{m}_i assigned to τ_i for schedulability analysis. The treatment of m_i (and \hat{m}_i) as a real number rather than an integer is consistent throughout this work.

Beneficial collapse, given by Definition 7 includes the improvement of core allocation as one of three conditions. The first condition maintains the integrity of the analysis, a beneficial collapse may not introduce a cycle into the graph which the critical path length calculation depends upon.

Definition 7 (Beneficial Collapse). *For a task τ_i , DAG $G_i = (V, E)$, and candidate nodes $u, v \in V$ the collapse of $u \bowtie v$ which results in \hat{G}_i is beneficial if and only if:*

- 1) \hat{G}_i contains no cycles
- 2) $L_i \leq D_i \Rightarrow \hat{L}_i \leq D_i$
- 3) $\hat{m}_i \ll m_i$

Condition 2 of beneficial collapse definition protects against collapse increasing the critical path length L_i beyond the deadline D_i , which would create an unschedulable task. Protection does not prevent unschedulable tasks becoming schedulable by collapse, due to the post-collapse critical path length being bounded by the deadline only if the pre-collapse critical path length was also less than the deadline.

C. Optimal Collapse

The primary goal of this work is to improve the schedulability of a task set by reducing the number of cores reserved for high utilization tasks. Defining optimality with respect to the number of cores assigned to a task matches the goal of this work.

Definition 8 (Optimal Collapse of a Task). *The optimal collapse of a DAG G is a DAG \hat{G} with the least positive \hat{m} obtainable by collapsing candidates of G .*

Currently, the complexity class of selecting the optimal set of candidates to collapse for a single task is unknown and remains an open problem. Observations 1-5 along with Definitions 6 and 7 illustrate the difficulties of identifying candidates that are beneficial to collapse. The only known method to compute the optimal collapse of a task requires the exploration of all possible combinations of candidates. Since there may be V^2 candidates per task, exploring all possible combinations has a time complexity of $\mathbb{O}(2^{V^2})$. Generating the optimal formulation and finding an optimal collapse of a task are both potentially intractable problems and reserved for future work. As a practical alternative, heuristics for ordering candidates for collapse are proposed in Section VI.

V. COLLAPSING HIGH UTILIZATION TASKS

Due to the intractability of optimal collapse for a task, this work proposes an intuitive heuristic presented in Algorithm 1. It collapses beneficial candidates (Definition 7), attempting to reduce the number of cores allocated to a high utilization task.

Reduction begins by identifying the potential candidates for collapse on Line 2. Candidacy follows Definition 4. Calculating the complete set of candidates is of complexity $\mathbb{O}(V^2)$. The set

Algorithm 1 DAG-OT Dedicated Core Reduction Algorithm

```

1: procedure DAGOT-REDUCE( $G_i$ )
2:    $A \leftarrow$  CANDIDATES( $G_i$ )
3:    $A \leftarrow$  ORDER( $A$ )
4:   while  $|A| \neq 0$  do
5:      $(u, v) \leftarrow$  FIRST( $A$ )
6:      $A \leftarrow A \setminus (u, v)$ 
7:     if BENEFIT( $G_i, u, v$ ) then
8:       COLLAPSE( $G_i, u, v$ )
9:     end if
10:  end while
11: end procedure

```

of candidates is prioritized for collapse consideration by ORDER. Ordering heuristics are proposed in Section VI. Each proposed heuristic is of equal or lesser computational complexity than the while loop (and its contents) beginning on Line 4.

Only candidates that benefit the task set are collapsed, improving (Definition 6) the number of cores allocated to a task without introducing a cycle into the DAG. The time complexity of checking for a cycle in \hat{G}_i by a depth first search (DFS) is $\mathbb{O}(V + E)$. The time complexity of calculating \hat{L}_i of a DAG by topological sort is also $\mathbb{O}(V + E)$. Determining if the number of allocated cores satisfy the definition of improvement is an $\mathbb{O}(1)$ operation, and collapse is an $\mathbb{O}(E)$ operation. Iterating over $\mathbb{O}(V^2)$ possible candidates, time complexity of Algorithm 1 is $\mathbb{O}(V^3 + V^2E)$.

During each iteration of the while loop on Line 4 of the DAGOT-REDUCE Algorithm 1 the current state of the DAG G_i serves as input and \hat{G}_i is the output. A subsequent iteration of the loop consumes the previous \hat{G}_i value as input when considering the next candidate for collapse.

VI. CANDIDATE ORDERING

In this work, two heuristics for collapse ordering are proposed: “greatest benefit”, orders the candidates by descending workload savings; “least penalty”, orders candidates by increasing longest path extension.

A. Greatest Benefit

For the greatest benefit heuristic, intuition suggests that collapsing nodes that most reduce the total workload C_i will also reduce the number of cores m_i maximally. The difference in workload is represented by Δ in Equation 14. There is a one time cost to calculate Δ for all candidates in A and order the set. This operation is of $\mathbb{O}(V \lg V)$ complexity. Employing the greatest benefit heuristic, Algorithm 1 is then $\mathbb{O}(V \lg V + V^3 + V^2E) = \mathbb{O}(V^3 + V^2E)$ complex.

$$\Delta = c_u(\eta_u) + c_v(\eta_v) - c_u(\eta_u + \eta_v) \quad (14)$$

B. Least Penalty

For the least penalty heuristic, the intuitive reasoning is collapsing nodes that impact the critical path length by the smallest amount (possibly negative) may permit more nodes to be collapsed overall. Penalties γ are calculated once by Equation 15 for every candidate pair. The set of candidates A are ordered by increasing penalty for use in Algorithm 1.

$$\gamma = \hat{L}_i - L_i \quad (15)$$

Penalty calculation requires a topological sort for every candidate to find \hat{L}_i with complexity $\mathbb{O}(V + E)$, for $\mathbb{O}(V^2)$ candidates. Sorting the candidates by penalty is $\mathbb{O}(V \lg V)$ complex. Therefore, the initial penalty ordering complexity is $\mathbb{O}(V^3 + V^2E + V \lg V)$. The complexity of Algorithm 1 utilizing the least penalty heuristic is then $\mathbb{O}(V^3 + V^2E + V \lg V + V^3 + V^2E) = \mathbb{O}(V^3 + V^2E)$.

Penalty calculations apply to a single DAG $G_i = (V, E)$ instance. Collapsing two nodes $u, v \in V$ may impact the critical path length, i.e. $\hat{L}_i \neq L_i$. The penalty of collapse depends on the critical path length, the collapse of $u \bowtie v$ may impact the penalty γ of other candidates. Penalties are not adjusted after collapsing nodes for the least penalty heuristic in favor of maintaining the $\mathbb{O}(V^3 + V^2E)$ complexity of Algorithm 1.

VII. COLLAPSING LOW UTILIZATION TASKS

Previous sections have focused on reducing m_i for high utilization tasks. Low utilization tasks may also incorporate the inter-thread cache benefit through collapse. To incorporate the benefit, a non-preemptive scheduler is required due to BUNDLE’s lack of preemptive schedulability analysis.

A low utilization DAG task $\tau_i \in \tau_{low}$ requires no more than one core $m_i = 1$ to meet all deadlines. Therefore, τ_i may be *serialized*. To serialize τ_i a topological sort of G_i is performed and nodes are executed on the single processor in sort order. Figure 10 illustrates the serialization of a task τ_i .




Fig. 10: Serializing a Task τ_i

Before a low utilization task is serialized all beneficial (Definition 7) candidates $u, v \in V_i$ collapsed. For a serialized task τ_i , the workload bounds the critical path length $C_i \geq L_i$. A serialized task is infeasible if $C_i > D_i$. Since the workload is only reduced by collapse, collapse preceding serialization cannot convert a feasible task into an infeasible one.

Similar to high utilization tasks, the complexity of serializing low utilization tasks depends on the number of candidates $\mathbb{O}(V^2)$, a DFS to check for cycles $\mathbb{O}(V + E)$, and a topological sort to order execution $\mathbb{O}(V + E)$. The total complexity of the operation is $\mathbb{O}(V^2 \cdot (V + E) + (V + E)) = \mathbb{O}(V^3 + V^2E)$.

Another concern shared with high utilization tasks is the order of collapse. For simplicity, collapse ordering is defined for the entire task set and shared between high and low utilization tasks. Whichever heuristic is selected for high utilization tasks is also selected for low utilization tasks for all tasks $\tau_i \in \tau$.

Every collapsed and serialized low utilization task $\tau_i \in \tau_{low}$ is scheduled non-preemptively, lest the inter-thread cache benefit of scheduling individual threads of nodes via BUNDLE be lost. To preserve the benefit of collapse and BUNDLE, low utilization jobs are scheduled by non-preemptive EDF.

Each low utilization task $\tau_i \in \tau_{low}$ is assigned to exactly one of the m_{low} cores by the *Worst-Fit* [8]¹ heuristic. Worst-Fit assigns each task $\tau_i \in \tau_{low}$ to a per-core task set on a core m_k when including τ_i will not create an infeasible per-core task set determined by [22]. Once assigned, jobs of τ_i will execute only upon its assigned processor.

VIII. SCHEDULING AND SCHEDULABILITY ANALYSIS

Federated scheduling and schedulability analysis [27] of parallel DAG task sets may be considered for separate task sets: high and low utilization tasks. For any high utilization task $\tau_i \in \tau_{high}$, any greedy non-preemptive scheduler may be used to select which node to execute upon one of the m_i cores dedicated to the task. For low utilization tasks τ_{low} , a preemptive or non-preemptive multi-core scheduling algorithm may be used to execute nodes upon the m_{low} cores.

Schedulability analysis of high utilization tasks follows from two conditions. The first is the requirement that a task $\tau_i \in \tau_{high}$ must have a critical path length less than its deadline $L_i < D_i$. The second is that τ_i has m_i cores allocated as calculated by Equation 5. If there are an insufficient number of cores in the system to satisfy all high utilization tasks i.e. $m < m_{high} = \sum_{\tau_i \in \tau_{high}} m_i$, the task set is unschedulable.

In this work, low utilization tasks are scheduled by partitioned EDF to the $m_{low} = m - m_{high}$ cores. For DAG tasks, that may be preemptive or non-preemptive EDF. For DAG-OT tasks, it must be non-preemptive EDF or the benefits of BUNDLE scheduling cannot be guaranteed. In partitioned EDF, tasks are assigned to cores. In the preemptive and non-preemptive setting, tasks are assigned to cores by the *Worst-Fit* [8] heuristic. Under Worst-Fit partitioning, a task will not be assigned to a core if assigning it would violate the uniprocessor scheduling test. The uniprocessor non-preemptive schedulability test is taken from [22] and the preemptive schedulability test from [7].

In summary, a taskset is deemed schedulable if all high and low utilization tasks are schedulable. For high utilization tasks, there must be a sufficient number of m_{high} cores and all critical paths are less than deadlines. For low utilization tasks, all tasks must be partitioned on the m_{low} cores according to Worst-Fit without violating the appropriate per-core schedulability test [22] or [7].

IX. EVALUATION

Evaluation of the approach proposed in this work focuses on two metrics: *schedulability ratios* and the *reduction of dedicated cores* to high utilization tasks. No existing approach to federated scheduling tasks incorporating the positive impact of instruction caches is (currently) known. To illustrate the potential of inter-thread cache benefits to DAG tasks under federated scheduling [27], high utilization tasks are scheduled by any non-preemptive work-conserving algorithm on the cores dedicated to the individual tasks. Low utilization tasks are assigned to cores by the Worst-Fit [8] partitioning algorithm

¹Any non-preemptive EDF schedulability test based task assignment to cores could be chosen.

and scheduled by non-preemptive EDF. In addition to the non-preemptive EDF scheduling of low utilization tasks, a comparison to federated scheduling using preemptive EDF of low utilization tasks is made. For preemptive EDF, it is assumed that preemptions have no preemption cost. As the proposed approach uses non-preemptive scheduling for scheduling low utilization tasks, this assumption only benefits the previous federated scheduling schemes which require preemption.

To permit larger scale testing, if any schedulability test for a task set takes more than 10 minutes to complete, then that task set is deemed unschedulable for the given test. For fairness across the heuristics, such a task set is deemed unschedulable for all heuristic collapse methods.

The existing schedulability analysis approaches are compared to collapse by DAGOT-REDUCE using the proposed heuristics. The proposed heuristics are also compared against an “arbitrary” (random) ordering to highlight each heuristic’s impact. Table IV summarizes the existing and proposed approaches used in the evaluation along with their notation. The approaches are enumerated by their inclusion of collapse and their use of non-preemptive EDF (EDF-NP) or preemptive EDF (EDF-P) for low utilization tasks.

| Approach | EDF-NP | EDF-P |
|---------------------------|--------|-------|
| Baseline (No Collapse) | B-NP | B-P |
| Collapse Arbitrary | OT-A | ∅ |
| Collapse Greatest Benefit | OT-G | ∅ |
| Collapse Least Penalty | OT-L | ∅ |

TABLE IV: Federated Schedulability Tests to Compare

Synthetic task sets are provided to each of the schedulability tests. Generation of the synthetic DAG tasks takes the form of a pipeline, where individual tasks are synthesized and then combined to make task sets. To allow a comparison to be made between the baseline and collapsed tasks, tasks are generated for the baseline first and then collapsed. DAGOT-REDUCE modifies the structure of DAG tasks, as well as the critical path length and total demand. Due collapse related changes, tasks that were trivially infeasible (i.e., $L_i > D_i$) may become feasible. As such, existing approaches [59] to task set generation which exclude trivially infeasible tasks are unsuitable for this work.




Fig. 11: Task Set Generation Pipeline

Figure 11 describes the pipeline in coarsest detail. Individual tasks are generated, and filtered. Due to length restrictions, the full details of the task set generation pipeline are available from [57], as is the framework on github [52, 53].

Summarily, tasks are generated by creating a representative set of 90 DAG task graphs of $\{16, 32, 64\}$ nodes, with a variable edge probability between each pair of nodes. For each

graph structure, nine graphs are generated by parameterizing the number of executable objects $\{4,8,16\}$ and their growth factors $\{0.2, 0.6, 1.0\}$. For each task, six new tasks are generated with deadlines calculated using a target utilization $\{0.25, 0.50, 2.0, 4.0, 8.0, 16.0\}$. In total, 4,860 tasks are generated.

Filtration of the 4,860 tasks removes only those tasks which are trivially infeasible ($L_i > D_i$) for the baseline DAG task and for all collapsed DAG-OT tasks. It should be noted that any post-collapse DAG-OT task $\hat{\tau}_i$ which is trivially infeasible could not have originated from a pre-collapse trivially infeasible DAG task τ_i . The filtered tasks are then duplicated, once per collapse ordering, before being assembled into task sets.

Table V enumerates the total number of task sets created by target task set utilization U , cores of the architecture c , and number of task sets assembled per utilization and core specification N . The total reflects the total number of DAG task sets assembled, it does not reflect the equivalent DAG-OT task sets (resulting from collapse by each of the heuristics).

| Parameter | Range |
|-----------|---|
| U | $\{0.5, 1, 2, 4, 8, 12, 16, 20, 24, 28, 32, 36\}$ |
| c | $\{4, 8, 12, 16, 20, 24, 28, 32\}$ |
| N | 1000 |
| Total | $N \cdot c \cdot U = 96,000$ |

TABLE V: Task Set Assembly Parameters

A. Evaluation Metrics

A schedulability ratio is calculated for each of the schedulability tests. For the DAG-OT schedulability tests, the number of cores saved $m_{i,saved}$ per task τ_i is calculated by Equation 16 where pre-collapse $m_{i,high}$ comes from Equation 5 and $\hat{m}_{i,high}$ after Algorithm 1 has terminated.

$$m_{i,saved} = m_{i,high} - \hat{m}_{i,high} \quad (16)$$

For a task set τ , the change in number of cores allocated to high utilization tasks is given by Equation 17.

$$\Delta_m = \sum_{\tau_i \in \tau} m_{i,saved} \quad (17)$$

For a DAG-OT task $\hat{\tau}_i$ collapsed from a DAG task τ_i , the workload reduction and critical path length extension of collapse are computed by Equations 18 and 19 respectively.

$$\Delta_C = C_i - \hat{C}_i \quad (18)$$

$$\Delta_L = \hat{L}_i - L_i \quad (19)$$

B. Results

Figure 12 summarizes the schedulability results. In the title ‘4’ indicates the utilization interval the column summarizes. For the histograms labeled ‘0’, the utilization schedulability ratio is for task sets from with utilization $[0, 4)$. The height of the bar is the average schedulability ratio over the interval. From this summary data, it is clear that collapse improves the schedulability of federated scheduled DAG tasks where collapsed task sets have a higher schedulability ratios.




Fig. 12: Mean Schedulability

Furthermore the gains in schedulability from collapsing outweigh any deleterious effects of the non-preemptive scheduling requirement for DAG-OT. This can be observed through the higher schedulability ratios for collapsed task sets compared to the uncollapsed *fully preemptive* low utilization task sets of B-P. The fully preemptive scheduler incurs no penalty for preemptions between low utilization tasks.

Requiring consideration for trivially infeasible tasks where the critical path length exceeds the deadline ($L_i > D_i$), constraints found in other works for task set formulation are prohibited. For example, in [45] the minimum period for an arbitrary period task τ_i is $T_i = L_i + 2\frac{C_i}{m}$. Due to implicit deadline $T_i = D_i$ tasks, no arbitrary period task in [45] will require more than $\frac{m}{2}$ cores. In this work, no such constraint is possible resulting in tasks requiring up to V_i cores. Consequentially, higher utilization task sets assembled from tasks with core allocations upper bounded by the number of nodes in a task are more likely to be unschedulable. Thus, schedulability ratios for utilizations over twenty are near zero.

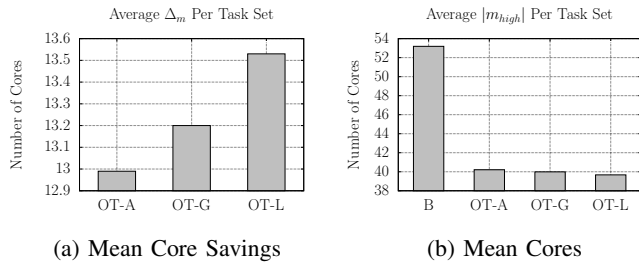


Fig. 13: Mean Cores and Savings

It is unclear from Figure 12 which of the collapse heuristics is the most desirable. For different utilization intervals, the heuristic with highest schedulability ratio may differ. Figures 13-15 present the impact of collapse for metrics other than schedulability. Both the preemptive (B-P) and non-preemptive (B-NP) uncollapsed baseline methods share the same task sets and therefore the same metrics. The label B represents both uncollapsed baselines in the figures. Figure 13 focuses on the central purpose of collapse: to reduce the number of cores assigned to high utilization tasks. The least penalty heuristic (OT-L) performs better than greatest benefit (OT-G). With arbitrary collapse ordering (OT-A) performing below the heuristics. For these task sets, the OT-L heuristic provides an approximately 20% reduction in dedicated cores, greater than arbitrary or OT-G ordering for collapse.

The least penalty (OT-L) heuristic seeks to collapse those nodes with the smallest increase to the critical path length before others. Surprisingly, Figure 14 shows the least penalty ordering of collapse may not have the intended effect. For OT-L, the average critical path length is greater than greatest benefit

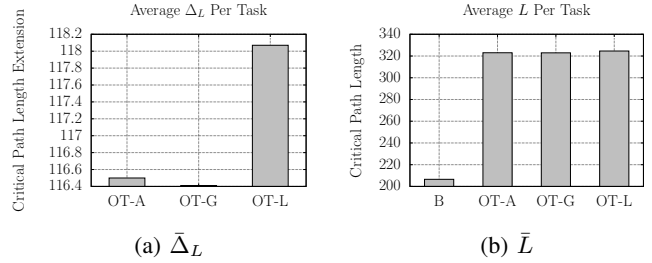


Fig. 14: Mean Critical Path Lengths and Extensions

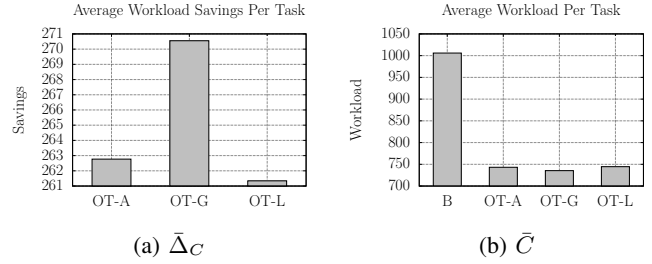


Fig. 15: Mean Workloads and Savings

(OT-G) or arbitrary ordering (OT-A); although it remains within 2 percent of OT-G and OT-A.

Figure 15 illustrates the benefits of collapse to the workload for all orderings, with OT-G providing the greatest workload reduction of 27 percent; a 5 percent improvement over OT-L.

From the results of Figure 13, 14, and 15 the OT-G heuristic performs similarly to OT-L in terms of core savings, the primary purpose of collapse. However, OT-G outperforms OT-L in both workload savings and critical path length extension. This is due to the nature of critical path length extension in comparison to workload savings. With each collapse, there is potential for the critical path to shift from one set of nodes to another. If the critical path length shifts, the initial least penalty ordering may no longer be in descending critical path length extension order. Workload savings are not affected when the critical path shifts; thus greatest benefit provides more consistent behavior and overall better performance.

X. FEASIBILITY STUDY

To compliment the synthetic results, a feasibility study (FS) was developed using the TacleBench [20] benchmarks executing on Raspberry Pi 3 devices. The purpose of the FS is to verify the potential benefit of collapse and the concave growth of WCETO values from BUNDLE to parallel DAG tasks.

Existing parallel programming environments such as OpenMP and CilkPlus lack features for controlling thread scheduling and management with respect to cache behavior. Additionally, BUNDLE's analysis [55] is limited to MIPS processors. The BUNDLE scheduling algorithm presented in [55] requires the use of a MIPS simulator. Lacking an ideal environment and platform for the FS, Raspberry Pi 3 devices were selected for their cost and limited hardware components.

BUNDLE's analysis and scheduling algorithm have not been implemented for ARM processors or Raspberry Pi 3 devices. Existing WCET tools [1] for ARM processors do not

provide cache analysis. Lacking a WCET tool that provides adequate cache analysis, representative WCET and growth factor values are calculated based upon observed execution times. A representative WCET of a TacleBench benchmark is the maximum number of cycles from the observed worst-case response time (WCRT) from the set of multiple distinct executions upon a Raspberry Pi 3 device. Representative growth factors are calculated from the WCRT of η threads executed sequentially upon a single core – e.g. the benefit of three threads scheduled by BUNDLE is estimated by executing three threads one after another on one core. These representative values are not reliable since they are based upon observations rather than analysis and could not be used in an environment where deadlines must be kept.

In BUNDLE scheduling, executable objects are divided into conflict free regions. Thread execution is coordinated by region, thereby maximizing the inter-thread cache benefit. In comparison, sequential scheduling of threads, where each thread executes from the first instruction to the last, has (potentially) fewer inter-thread cache benefits [54, 55]. Thus, representative growth factor values from sequential execution produces greater (worse) values than BUNDLE analysis would provide. This over-estimation biases the results against the proposed approach.

The FS is comprised of a process server running on a general purpose computer and Raspberry Pi 3 devices, where each Pi device represents a single core. The process server schedules nodes of the DAG-OT task non-preemptively in a greedy fashion. Execution of a node is the sequential execution of a benchmark upon one of the Raspberry Pi’s representing a core. Figure 16 illustrates the computing platform architecture available for download [37].




Fig. 16: Experiment Architecture

Raspberry Pi 3 devices contain a Samsung four-core ARM Cortex A53 processor with 32KB of L1 instruction cache, 32KB of L1 data cache, and 512KB of L2 unified data and instruction cache [2]. All Raspberry Pi 3 devices utilize the Raspbian operating system running Linux kernel 4.18. To minimize interference from the operating system as well other processes one of the four cores is reserved exclusively for the execution of benchmarks.

The sequential execution of benchmarks may benefit from data values persisting in the cache, thereby decreasing growth factors and biasing the results towards the proposed approach. To address the potential bias an attempt at flushing the data cache is performed between benchmark executions. Algorithm 2 illustrates the method of sequential thread scheduling including data cache flushing. Note, the flush may be incomplete due to

the pseudo-random L1 and L2 replacement policy [2].

Each main function from the TacleBench suite is modified according to Algorithm 2, where m is a command line argument for the number of threads to execute. The READ_CYCLES() function reads the current cycle count of the processor. The OLD_MAIN() function is the TacleBench provided main function. The CLEAR_DATA_CACHE() function reads 512KB (the size of the L2 cache) of allocated memory in an attempt to flush the data cache.

Algorithm 2 Sequential Thread Execution

```

1: procedure MAIN( $m$ )
2:    $total \leftarrow 0$ 
3:   for  $i \leftarrow 1$  to  $m$  do
4:      $pre \leftarrow \text{READ\_CYCLES}()$ 
5:     OLD_MAIN()
6:      $post \leftarrow \text{READ\_CYCLES}()$ 
7:      $total \leftarrow total + post - pre$ 
8:     CLEAR_DATA_CACHE()
9:   end for
10: end procedure

```

Algorithm 2 executes a benchmark once per iteration of the for loop, completing m threads sequentially. After each benchmark execution, the response time is measured by taking the difference of the processor cycle count before and after execution. The difference is added to the *total* to compute the total number of cycles required to execute m threads of a benchmark. After every execution of a benchmark, CLEAR_DATA_CACHE() performs a partial flush of the data cache. The goal is to limit the contribution to the inter-thread cache benefit through the data cache.

The FS uses the sample input data from the TacleBench suite for every execution. Using the same input for each thread scheduled sequentially is an approximation of the inter-thread cache benefit BUNDLE scheduling would produce. Under these circumstances, the sequential execution provides a lower bound on the inter-thread cache benefit producing greater (worse) growth factors than BUNDLE scheduling would.

Each of the 42 benchmarks is executed on a dedicated core of a Raspberry Pi 3 for $m \in [1, 10]$ threads. For every m value, the benchmark is executed 100 times totaling 1,000 executions. The maximum *total* cycles calculated by Algorithm 2 from 100 executions of m threads is recorded as the representative WCET for m threads. From these WCET values, the minimum representative growth factor is calculated for every benchmark.

To verify the benefit of collapse proposed in this work, DAG tasks are constructed using the generation pipeline from Section IX with one modification: executable objects are one of the TacleBench benchmarks with WCET and growth factor values estimated by the repeated execution of Algorithm 2. Nodes within DAG tasks are assigned one thread of one benchmark. After assigning executable objects to nodes, the workload C_i , critical path length L_i , and dedicated cores m_i are calculated by Equations 2, 1, and 5, respectively. The DAG task is then converted to DAG-OT tasks and nodes are collapsed by each of the heuristics. The result is four tasks: one DAG requiring m_i cores, and three DAG-OT requiring $\hat{m}_i \leq m_i$ cores due to nodes being collapsed by the distinct heuristics.

The makespan of the four tasks is recorded by executing the tasks on the FS platform given the proper allocation of m_i or \hat{m}_i cores. Makespan values are compared to the task’s deadline, verifying schedulability and illustrating the core savings resulting from collapse.

A. Feasibility Study Results

Growth factors for the 42 benchmarks fall in the range of 0.3 to 7. Benchmarks with a representative growth factor greater than 1.0 are not collapsed, since they are not beneficial. A sample of benchmark values are provided in Table VI. The complete list of growth factors may be found in [57].

| Benchmark | fac | matrix1 | ndes |
|---------------|------|---------|------|
| Growth Factor | 0.42 | 0.84 | 1.38 |

TABLE VI: Sample TacleBench Benchmarks Growth Factors

| Pre Collapse | | | | |
|---------------|-------------|-------------|-----------|-------------|
| i | C_i | L_i | D_i | m_i |
| 1 | 6,168,224 | 4,287,924 | 5,248,928 | 3 |
| 2 | 4,616,294 | 3,448,417 | 6,347,882 | 2 |
| 3 | 5,310,666 | 3,614,573 | 3,953,663 | 4 |
| 4 | 6,684,846 | 4,149,946 | 4,448,542 | 4 |
| Post Collapse | | | | |
| i | \hat{C}_i | \hat{L}_i | D_i | \hat{m}_i |
| 1 | 6,087,061 | 5,195,855 | 5,248,928 | 2 |
| 2 | 3,888,725 | 3,888,725 | 6,347,882 | 1 |
| 3 | 5,018,733 | 3,853,186 | 3,953,663 | 3 |
| 4 | 6,342,401 | 3,883,653 | 4,448,542 | 3 |

TABLE VII: Pre and Post Collapse Metrics

A subset of 11 benchmarks were selected as the complete set of executable objects when generating tasks. These 11 were selected based on their growth factors which range from 0.3 to 2.63. Task graphs are generated with 32 nodes and edge probabilities in $\{0.01, 0.02, 0.03\}$. Every node is randomly assigned one thread of an executable object from the 11 heuristics. A total of 120 DAG tasks were generated and analyzed without collapse and as DAG-OT tasks collapsed by each of the benchmarks. Four of 120 were selected based on their results to illustrate the benefits of collapse. The tasks require two, three, or four cores. Each task’s core allocation is reduced by one to: one, two, and three cores respectively. Tasks requiring more cores were not considered due to the limited number of Raspberry Pi 3 devices available.

Table VII presents the four selected tasks and the impact of collapse upon them when run on the POC platform. In this limited setting, each of the heuristics OT-G, OT-L, and OT-A collapsed the same set of candidate pairs. Thus, the critical path length and workload values were similar across all heuristics. The result is a core savings of 25 to 50 percent.

During execution on the POC, the makespan and workload are recorded. Given the similar performance of each heuristic,

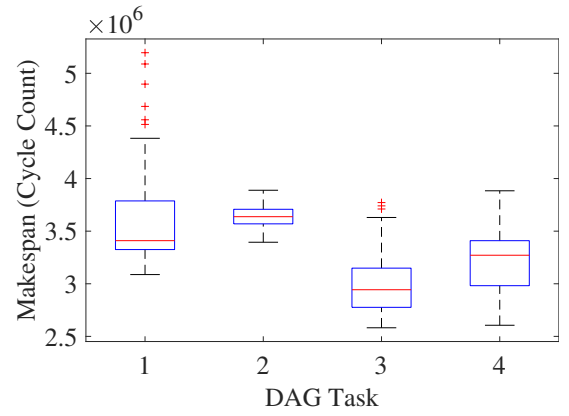


Fig. 17: Makespan Distribution for OT-G Collapsed Tasks

Figure 17 presents the makespan distribution of the 100 runs of each task collapsed by OT-G. Variation in makespan (and workload) may be attributed to interrupts, operating system interference, or the pseudo-random cache replacement policy. Combined with Table VIII, the average makespan falls within the 90th percentile. Given the distribution, the average values are presented for simplicity.

Table VIII provides the average makespan and workload savings for all tasks across each of the heuristics. Workload savings ranges from 2 to 16 percent. The results also verify schedulability of collapsed task with all makespan values falling below the deadlines in Table VII. In this limited setting including the negative effect of cache clearing between threads the savings in makespan, workload, and core allocation are encouraging for the method proposed herein.

| i | Makespan | | | $\bar{\Delta}_C$ |
|-----|-----------|-----------|-----------|------------------|
| | OT-G | OT-L | OT-A | |
| 1 | 4,531,262 | 5,195,855 | 4,640,880 | 2.03% |
| 2 | 3,888,725 | 3,858,028 | 3,942,390 | 16.43% |
| 3 | 3,853,186 | 3,600,027 | 3,835,659 | 6.81% |
| 4 | 3,883,653 | 4,213,339 | 4,436,712 | 5.12% |

TABLE VIII: Mean Makespan and Workload Savings

XI. CONCLUSION

This work proposes the DAG-OT model, joining the federated scheduling policy and analysis with BUNDLE thread-level scheduling and analysis through the proposed mechanism of collapsing candidate nodes of a DAG. The synthetic evaluation and proof of concept support the proposed mechanism, and heuristic algorithm for selecting and collapsing nodes; demonstrating the benefit of collapse to schedulability, workload, and total cores allocated to parallel DAG tasks.

There remains an open question of the complexity of optimal collapse of a task. Optimal collapse of all tasks within a task set remains undefined. The complexity of optimal collapse for a single task and all tasks of task set is reserved for future work. Future work also includes consideration for data caches, shared caches (evictions and false sharing), and permitting preemptions within BUNDLE scheduling.

XII. ACKNOWLEDGMENTS

The research presented in this paper was supported by the National Science Foundation under Grant Numbers CNS-1618185 and IIS-1724227. The authors would like to acknowledge the anonymous reviewers. Their thoughtful and careful feedback has improved this work, thank you.

REFERENCES

- [1] aiT Worst-Case Execution Time Analyzers. 2020. <https://www.absint.com/ait/index.htm>.
- [2] Arm cortex a53 technical reference manual, revision r0p4. 2020. <https://developer.arm.com/docs/ddi0500/g/level-2-memory-system/optional-integrated-l2-cache>.
- [3] Sebastian Altmeyer, Robert I Davis, and Claire Maiza. Improved cache related pre-emption delay aware response time analysis for fixed priority pre-emptive systems. *Real-Time Systems*, 48(5):499–526, 2012.
- [4] Sebastian Altmeyer and Claire Maiza Burguière. Cache-related preemption delay via useful cache blocks: Survey and redefinition. *Journal of Systems Architecture*, 57(7):707–719, August 2011.
- [5] R. Arnold, F. Mueller, D. Whalley, and M. Harmon. Bounding worst-case instruction cache performance. *Real-Time Systems Symposium, 1994., Proceedings.*, pages 172–181, Dec 1994.
- [6] S. Bak, G. Yao, R. Pellizzoni, and M. Caccamo. Memory-aware scheduling of multicore task sets for real-time systems. In *IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, pages 300–309, Aug 2012. doi:10.1109/RTCSA.2012.48.
- [7] S. K. Baruah, A. K. Mok, and L. E. Rosier. Preemptively scheduling hard-real-time sporadic tasks on one processor. In *IEEE Real-Time Systems Symposium*, pages 182–190, Dec 1990. doi:10.1109/REAL.1990.128746.
- [8] Sanjoy Baruah. Partitioned edf scheduling: a closer look. *Real-Time Systems*, 49(6):715–729, 2013.
- [9] Sanjoy Baruah. The federated scheduling of constrained-deadline sporadic dag task systems. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 1323–1328. EDA Consortium, 2015.
- [10] Sanjoy Baruah. Federated scheduling of sporadic dag task systems. In *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*, pages 179–186. IEEE, 2015.
- [11] Sanjoy Baruah, Vincenzo Bonifaci, and Alberto Marchetti-Spaccamela. The global edf scheduling of systems of conditional sporadic dag tasks. In *Real-Time Systems (ECRTS), 2015 27th Euromicro Conference on*, pages 222–231. IEEE, 2015.
- [12] M. Bertogna, O. Xhani, M. Marinoni, F. Esposito, and G. Buttazzo. Optimal selection of preemption points to minimize preemption overhead. In *Proceedings of the Euromicro Conference on Real-Time Systems*, pages 217–227, July 2011. doi:10.1109/ECRTS.2011.28.
- [13] R. Bril, S. Altmeyer, M. van den Heuvel, R. Davis, and M. Behnam. Fixed priority scheduling with pre-emption thresholds and cache-related pre-emption delays: integrated analysis and evaluation. *Real-Time Systems*, 53(4):403–466, July 2017.
- [14] J. M. Calandrino and J. H. Anderson. Cache-aware real-time scheduling on multicore platforms: Heuristics and a case study. In *2008 Euromicro Conference on Real-Time Systems*, pages 299–308, July 2008. doi:10.1109/ECRTS.2008.10.
- [15] J. M. Calandrino and J. H. Anderson. On the design and implementation of a cache-aware multicore real-time scheduler. In *2009 21st Euromicro Conference on Real-Time Systems*, pages 194–204, July 2009. doi:10.1109/ECRTS.2009.13.
- [16] John Michael Calandrino. *On the Design and Implementation of a Cache-aware Soft Real-time Scheduler for Multicore Platforms*. PhD thesis, The University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 2009.
- [17] Sudipta Chattopadhyay and Abhik Roychoudhury. Cache-related preemption delay analysis for multilevel noninclusive caches. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(5s):147, 2014.
- [18] Richard Cole and Vijaya Ramachandran. Analysis of randomized work stealing with false sharing. *Computing Research Repository - CORR*, 03 2011. doi:10.1109/IPDPS.2013.86.
- [19] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. In *OSDI'04: Sixth Symposium on Operating System Design and Implementation*, pages 137–150, San Francisco, CA, 2004.
- [20] Heiko Falk, Sebastian Altmeyer, Peter Hellinckx, Björn Lisper, Wolfgang Puffitsch, Christine Rochange, Martin Schoeberl, Rasmus Bo Sørensen, Peter Wägemann, and Simon Wegener. TACLeBench: A benchmark collection to support worst-case execution time research. In Martin Schoeberl, editor, *16th International Workshop on Worst-Case Execution Time Analysis (WCET 2016)*, volume 55 of *OpenAccess Series in Informatics (OASISs)*, pages 2:1–2:10, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [21] Sebastian Hahn, Jan Reineke, and Reinhard Wilhelm. Towards compositionality in execution time analysis: Definition and challenges. *ACM SIGBED Review*, 12(1):28–36, March 2015.
- [22] K. Jeffay, D. F. Stanat, and C. U. Martel. On non-preemptive scheduling of period and sporadic tasks. In *[1991] Proceedings Twelfth Real-Time Systems Symposium*, pages 129–139, Dec 1991. doi:10.1109/REAL.1991.160366.
- [23] Lei Ju, Samarjit Chakraborty, and Abhik Roychoudhury. Accounting for cache-related preemption delay in dynamic priority schedulability analysis. In *Proceedings of the conference on Design, automation and test in Europe*, pages 1623–1628. EDA Consortium, 2007.
- [24] Karthik Lakshmanan, Shinpei Kato, and Ragunathan Raj

- Rajkumar. Scheduling parallel real-time tasks on multi-core processors. In *2010 31st IEEE Real-Time Systems Symposium*, pages 259–268. IEEE, 2010.
- [25] Chang-Gun Lee, Joosun Hahn, Yang-Min Seo, Sang Lyul Min, Rhan Ha, Seongsoo Hong, Chang Yun Park, Minsuk Lee, and Chong Sang Kim. Analysis of cache-related preemption delay in fixed-priority preemptive scheduling. *IEEE Transactions on Computers*, 47(6):700–713, June 1998.
- [26] Jing Li, Kunal Agrawal, Christopher Gill, and Chenyang Lu. Federated scheduling for stochastic parallel real-time tasks. In *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2014 IEEE 20th International Conference on*, pages 1–10. IEEE, 2014.
- [27] Jing Li, Jian Jia Chen, Kunal Agrawal, Chenyang Lu, Chris Gill, and Abusayeed Saifullah. Analysis of federated and global scheduling for parallel real-time tasks. In *Real-Time Systems (ECRTS), 2014 26th Euromicro Conference on*, pages 85–96. IEEE, 2014.
- [28] Y. Li, V. Suhendra, Y. Liang, T. Mitra, and A. Roychoudhury. Timing analysis of concurrent programs running on shared cache multi-cores. In *Real-Time Systems Symposium, 2009, RTSS 2009. 30th IEEE*, pages 57–67, Dec 2009.
- [29] Y.-T. S. Li, S. Malik, and A. Wolfe. Cache modeling for real-time software: beyond direct mapped instruction caches. In *17th IEEE Real-Time Systems Symposium*, pages 254–263, Dec 1996.
- [30] Tongping Liu and Xu Liu. Cheetah: Detecting false sharing efficiently and effectively. In *Proceedings of the 2016 International Symposium on Code Generation and Optimization*, CGO '16, page 1–11, New York, NY, USA, 2016. Association for Computing Machinery. URL: <https://doi.org/10.1145/2854038.2854039>, doi:10.1145/2854038.2854039.
- [31] J. C. S. Lui, R. R. Muntz, and D. Towsley. Computing performance bounds of fork-join parallel programs under a multiprocessing environment. *IEEE Transactions on Parallel and Distributed Systems*, 9(3):295–311, March 1998. doi:10.1109/71.674321.
- [32] W. Lunniss, S. Altmeyer, C. Maiza, and R. I. Davis. Integrating cache related pre-emption delay analysis into edf scheduling. In *Real-Time and Embedded Technology and Applications Symposium (RTAS), 2013 IEEE 19th*, pages 75–84, April 2013. doi:10.1109/RTAS.2013.6531081.
- [33] Will Lunniss, Sebastian Altmeyer, Robert I Davis, et al. A comparison between fixed priority and edf scheduling accounting for cache related pre-emption delays. *LITES*, 1(1):01–1, 2014.
- [34] Will Lunniss, Sebastian Altmeyer, Giuseppe Lipari, and Robert I Davis. Accounting for cache related pre-emption delays in hierarchical scheduling. In *Proceedings of the 22nd International Conference on Real-Time Networks and Systems*, page 183. ACM, 2014.
- [35] Will Lunniss, Sebastian Altmeyer, Giuseppe Lipari, and Robert I Davis. Cache related pre-emption delays in hierarchical scheduling. *Real-Time Systems*, 52(2):201–238, 2016.
- [36] Alessandra Melani, Marko Bertogna, Vincenzo Bonifaci, Alberto Marchetti-Spaccamela, and Giorgio C Buttazzo. Response-time analysis of conditional dag tasks in multi-processor systems. In *Real-Time Systems (ECRTS), 2015 27th Euromicro Conference on*, pages 211–221. IEEE, 2015.
- [37] Venkata P. Modekurthy. Inter-thread cache benefit feasibility study. February, 2020. URL: <https://github.com/venkata-prashant/itcb-results/tree/1.0>.
- [38] Frank Mueller. *Static Cache Simulation and Its Applications*. Ph.d. dissertation, Florida State University, 1995.
- [39] Frank Mueller. Timing analysis for instruction caches. In *The Journal of Real-Time Systems* 18, pages 217–247, 2000.
- [40] Hemendra Singh Negi, Tulika Mitra, and Abhik Roychoudhury. Accurate estimation of cache-related preemption delay. In *Proceedings of the 1st IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, CODES+ISSS '03, pages 201–206, New York, NY, USA, 2003. ACM.
- [41] R. Pellizzoni, E. Betti, S. Bak, G. Yao, J. Criswell, M. Caccamo, and R. Kegley. A predictable execution model for cots-based embedded systems. In *IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 269–279, April 2011. doi:10.1109/RTAS.2011.33.
- [42] S. A. Rashid, G. Nelissen, S. Altmeyer, R. I. Davis, and E. Tovar. Integrated analysis of cache related preemption delays and cache persistence reload overheads. In *IEEE Real-Time Systems Symposium*, pages 188–198, Dec 2017. doi:10.1109/RTSS.2017.00025.
- [43] S. A. Rashid, G. Nelissen, D. Hardy, B. Akesson, I. Puaut, and E. Tovar. Cache-persistence-aware response-time analysis for fixed-priority preemptive systems. In *Euromicro Conference on Real-Time Systems*, pages 262–272, July 2016. doi:10.1109/ECRTS.2016.25.
- [44] Suntorn Sae-eung. Analysis of false cache line sharing effects on multicore cpus. *Master's Projects*, 01 2010.
- [45] A. Saifullah, D. Ferry, J. Li, K. Agrawal, C. Lu, and C. D. Gill. Parallel real-time scheduling of dags. *IEEE Transactions on Parallel and Distributed Systems*, 25(12):3242–3252, Dec 2014. doi:10.1109/TPDS.2013.2297919.
- [46] Abusayeed Saifullah, David Ferry, Jing Li, Kunal Agrawal, Chenyang Lu, and Christopher D Gill. Parallel real-time scheduling of dags. *IEEE Transactions on Parallel and Distributed Systems*, 25(12):3242–3252, 2014.
- [47] Abusayeed Saifullah, Jing Li, Kunal Agrawal, Chenyang Lu, and Christopher Gill. Multi-core real-time scheduling for generalized parallel task models. *Real-Time Systems*, 49(4):404–435, 2013.
- [48] J. Simonson and J.H. Patel. Use of preferred preemption points in cache based real-time systems. In *Proceed-*

- ings of *IEEE International Computer Performance and Dependability Symposium*, 1995.
- [49] J. Staschulat, S. Schliecker, and R. Ernst. Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In *Proceedings of the 2005 17th Euromicro Conference on Real-Time Systems (ECRTS)*, ECRTS '05, pages 41–48, July 2005. doi:10.1109/ECRTS.2005.26.
- [50] Jinghao Sun, Nan Guan, Xu Jiang, Shuangshuang Chang, Zhishan Guo, Qingxu Deng, and Wang Yi. A capacity augmentation bound for real-time constrained-deadline parallel tasks under gedf. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37(11):2200–2211, 2018.
- [51] Yudong Tan and Vincent Mooney. Timing analysis for preemptive multitasking real-time systems with caches. *ACM Transactions on Embedded Computing Systems*, 6(1), February 2007. doi:10.1145/1210268.1210275.
- [52] Corey Tessler. *libsched* (v1.2). February, 2020. URL: <https://github.com/ctessler/libsched/tree/v1.2>.
- [53] Corey Tessler. *Synthetic evaluation* (v1.0). February, 2020. URL: <https://github.com/ctessler/itcb-dag-eval/tree/v1.0>.
- [54] Corey Tessler and Nathan Fisher. *BUNDLE: Real-Time Multi-Threaded Scheduling to Reduce Cache Contention*. In *IEEE Real-Time Systems Symposium*, 2016.
- [55] Corey Tessler and Nathan Fisher. *Bundle: Prioritizing conflict free regions in multi-threaded programs to improve cache reuse*. In *2018 IEEE Real-Time Systems Symposium (RTSS)*, pages 325–337. IEEE, 2018.
- [56] Corey Tessler and Nathan Fisher. *NPM-BUNDLE: Non-Preemptive Multitask Scheduling for Jobs with BUNDLE-Based Thread-Level Scheduling*. In Sophie Quinton, editor, *31st Euromicro Conference on Real-Time Systems (ECRTS 2019)*, volume 133 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:23, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/10752>, doi:10.4230/LIPIcs.ECRTS.2019.15.
- [57] Corey Tessler, Venkata P. Modekurthy, Nathan Fisher, and Abusayeed Saifullah. *Bringing inter-thread cache benefits to federated scheduling – extended results and technical report*. 2020. <https://arxiv.org/abs/2002.12516>.
- [58] H. Tomiyama and N.D. Dutt. Program path analysis to bound cache-related preemption delay in preemptive real-time systems. In *Proceedings of the Eighth International Workshop on Hardware/Software Codesign (CODES)*, pages 67–71, May 2000.
- [59] Niklas Ueter. *Reservation-Based Federated Scheduling for Parallel Real-Time Tasks*. Master’s thesis, Technische Universität Dortmund, 2018.
- [60] Niklas Ueter, Georg von der Brüggen, Jian-Jia Chen, Jing Li, and Kunal Agrawal. Reservation-based federated scheduling for parallel real-time tasks. In *2018 IEEE Real-Time Systems Symposium (RTSS)*, pages 482–494. IEEE, 2018.
- [61] Y. Wang and M. Saksena. Scheduling fixed-priority tasks with preemption threshold. In *Proceedings of the International Conference on Real Time Computing Systems and Applications*, 1999.
- [62] B. C. Ward, J. L. Herman, C. J. Kenna, and J. H. Anderson. Making shared caches more predictable on multicore platforms. In *Euromicro Conference on Real-Time Systems*, pages 157–167, July 2013. doi:10.1109/ECRTS.2013.26.
- [63] J. Xiao, S. Altmeyer, and A. Pimentel. Schedulability analysis of non-preemptive real-time scheduling for multicore processors with shared caches. In *2017 IEEE Real-Time Systems Symposium (RTSS)*, pages 199–208, Dec 2017. doi:10.1109/RTSS.2017.00026.
- [64] M. Xu, L. T. X. Phan, H. Choi, and I. Lee. Analysis and implementation of global preemptive fixed-priority scheduling with dynamic cache allocation. In *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 1–12, April 2016. doi:10.1109/RTAS.2016.7461322.
- [65] Zhenkai Zhang and Xenofon Koutsoukos. Cache-related preemption delay analysis for multi-level inclusive caches. In *Proceedings of the 13th International Conference on Embedded Software*, page 16. ACM, 2016.