Federated Scheduling for Stochastic Parallel Real-time Tasks

Jing Li, Kunal Agrawal, Christopher Gill, and Chenyang Lu
Department of Computer Science and Engineering
Washington University in St. Louis
{li.jing}@wustl.edu, {kunal, cdgill, lu}@cse.wustl.edu

Abstract—Federated scheduling is a strategy to schedule parallel real-time tasks: It allocates a dedicated cluster of cores to each high-utilization task (utilization ≥ 1); It uses a multiprocessor scheduling algorithm to schedule and execute all low-utilization tasks sequentially, on a shared cluster of the remaining cores. Prior work has shown that federated scheduling has the best known capacity augmentation bound of 2 for parallel tasks with implicit deadlines. In this paper, we explore the soft real-time performance of federated scheduling and address average-case workloads instead of worst-case ones. In particular, we consider stochastic tasks — tasks for which execution time and critical-path length are random variables. In this case, we use bounded expected tardiness as the schedulability criterion. We define a stochastic capacity augmentation bound and prove that federated scheduling algorithms guarantee the same bound of 2 for stochastic tasks. We present three federated mapping algorithms with different complexities for core allocation. All of them guarantee bounded expected tardiness and provide the same capacity augmentation bound. In practice, however, we expect them to provide different performance, both in terms of the task sets they can schedule and the actual tardiness they guarantee. Therefore, we present numerical evaluations using randomly generated task sets to examine the practical differences between the three algorithms.

Index Terms—soft real-time scheduling, parallel scheduling, federated scheduling, stochastic capacity augmentation bound

I. INTRODUCTION

Multi-core processors are changing the computing landscape as the number of cores per chip continues to increase. There has been extensive prior research on **multiprocessor scheduling**, where each task is sequential and can only use one core at a time. In contrast, **parallel scheduling** allows *intra-task parallelism* where individual tasks are themselves parallel programs. Therefore, unlike multiprocessor scheduling, individual tasks can run faster with an increasing number of cores. Parallel real-time systems are particularly well suited for computationally intensive tasks with tight deadlines; one recent example [1] showed that parallelism provides demonstrably better system performance for an autonomous vehicle.

As for real-time performance and guarantees, different applications have different requirements and characteristics. Some have a strict constraint that deadlines must be met and **hard real-time** guarantees should be provided; others have **soft real-time** constraints and can tolerate tardiness as long as it can be bounded. In this paper, we focus on *soft real-time scheduling*. In particular, we consider **stochastic tasks**, which may have large variability in their parameters; the worst-case execution time could be orders of magnitude larger than the mean execution time. In this case, the system may suffer significant utilization loss if we use only worst-case execution time while analyzing schedulability. Instead, we could use the mean and variance of task parameters for analysis. For these tasks, our schedulers guarantee that the expected tardiness of tasks is bounded.

In this paper, we analyze the **federated scheduling strategy** for parallel real-time tasks with implicit deadlines. In this strategy, each *high-utilization task* (utilization ≥ 1) is allocated a dedicated cluster (set) of cores and all the *low-utilization tasks* share a cluster composed of the remaining cores. The federated strategy works in two stages: Given a task set τ, a *mapping* algorithm either *admits* a task set and outputs a *core assignment* — which consists of a cluster for each high-utilization task and a final cluster for all low-utilization tasks — or declares that it cannot guarantee schedulability of the task set. Therefore, a mapping algorithm is automatically a schedulability test. After the mapping is done, the scheduling is straightforward. A *greedy* (work-conserving) scheduler is used to schedule each high-utilization task on its dedicated cluster. Since low-utilization tasks can be run sequentially, a multiprocessor scheduling algorithm, *global earliest deadline first* (*global EDF* or *GEDF*), is used to schedule all low-utilization tasks on their shared cluster. Notably, this federated strategy does not require any task decomposition or transformation; therefore, the internal structure of the DAGs need not be known in advance in order to use this strategy.

For stochastic parallel real-time tasks with soft real-time constraints, we choose to analyze the federated scheduling because of two reasons. First, prior work has shown that federated scheduling has the best known capacity augmentation bound of 2 for hard real-time parallel tasks [2]. A scheduling strategy $\mathcal{S}$ provides a (hard) capacity augmentation bound of b if, given a task set τ and a machine with m cores, $\mathcal{S}$ guarantees no deadline miss if (1) the total utilization of the tasks in τ is at most m/b and (2) the critical-path length of each task is at most $1/b$ of its deadline. Capacity augmentation bound is an extension of the utilization bound from sequential tasks to parallel real-time tasks. Similarly, it indicates how much load a system can schedule. Note that to guarantee hard real-time schedulability, worst-case values are used for task parameters. As the federated scheduling

has the best capacity augmentation bound in the worst case, it is interesting to analyze its performance for the average case. Second, the federated scheduling is a generalization of partitioned scheduling to parallel tasks. It assigns dedicated cores to each high-utilization task, which makes each high-utilization task being isolated from any interference from other tasks. Therefore, we could analyze each of them individually and directly apply result from single server queueing theory to bound the expected tardiness.

In this paper, we analyze the soft real-time performance of federated scheduling for stochastic tasks based on average-case workload. We define a **bounded-tardiness stochastic capacity bound** (**stochastic capacity bound** for short) that uses expected values for utilization and critical-path length to provide bounded expected tardiness. In contrast to using a hard real-time bound, it allows the system to be over-utilized in the worst case, as long as in average total workload is less than the bound. The contributions of this paper are:

1) Stochastic federated scheduling for implicit-deadline stochastic parallel tasks that provides a soft real-time guarantee of bounded expected tardiness on uniform multicores. We prove that federated scheduling provides a stochastic capacity bound of 2 for general DAG tasks. To our knowledge, this is the first result for stochastic parallel tasks. We also describe the procedure for calculating the corresponding (upper bound on) expected tardiness for all these algorithms.

2) We propose three different mapping algorithms for stochastic tasks: All these algorithms satisfy the same stochastic capacity augmentation bound and provide bounded tardiness. The three algorithms differ in their calculation for core allocation. They have increasing computation complexity (from linear-time to pseudo polynomial time) and also have increasing schedulability performance or decreasing upper bound on expected tardiness.

3) We conduct numerical evaluations using randomly generated task sets to understand the efficacy of the different stochastic mapping algorithms.

The outline of the paper is as follows: Section II discusses related work. Section III defines our stochastic task model and stochastic capacity augmentation bound. Section IV presents the stochastic federated scheduling strategy and expected tardiness calculation, and proves that expected tardiness is bounded. Section V presents the three different mapping algorithms for core allocation; Section VI proves that all of these algorithms provide a capacity augmentation bound of 2; Section VII presents the numerical experiments to compare the three mapping algorithm. Finally, we conclude in Section VIII.

II. RELATED WORK

Real-time multiprocessor scheduling for tasks with *worst-case* task parameters has been studied extensively [3, 4]. In particular, for implicit deadline hard real-time tasks, the best known utilization bound is $\approx 50\%$ using partitioned fixed priority scheduling [5] and partitioned EDF [6]; this trivially implies a capacity augmentation bound of 2. In comparison,

GEDF has a capacity augmentation bound of $2 - \frac{1}{m} + \epsilon$ for small ϵ [7, 8].

For parallel tasks with hard real-time constraints and worst-case task parameters, early work considered idealized models for tasks such as moldable and malleable tasks [9–12]. A commonly considered model is the **parallel synchronous model**, which is a subcategory of the **directed acyclic graph (DAG)** model. Many strategies for parallel synchronous tasks decompose parallel tasks into sets of sequential tasks [1, 13–16]. Without decomposition, researchers have studied both synchronous tasks [17] and general DAG tasks [18–22]. For hard real-time tasks with worst-case parameters, the best capacity augmentation bound known for general DAGs is 2 using federated scheduling (a partition-like strategy) without decomposition [2]; 2.6 using GEDF without decomposition [2]; 3.73 for rate-monotonic with and without decompostion [1, 2]; and 3.42 for a more restricted class of synchronous tasks [13].

Most prior work on bounded tardiness (and other soft real-time guarantees) considers sequential tasks with worst-case parameters [23]. For these tasks, an earliest-pseudo-deadline-first scheduler [24] and GEDF [25, 26] both provide bounded tardiness with no utilization loss; these results were generalized to many global schedulers [27]. Lateness guarantees also have been studied for GEDF-like scheduling [28]. For parallel tasks, Liu et al. [18] for the first time provide a soft real-time response time analysis for GEDF.

For stochastic analysis, there is some prior work on sequential stochastic tasks. For a resource reservation scheduler, a lower bound on the probability of deadline misses was derived in [29]. For multiprocessor scheduling, [30] shows that GEDF guarantees bounded tardiness to sequential tasks if the total expected utilization is smaller than the number of cores. We use this result directly in our algorithms and analysis to guarantee bounded tardiness to low-utilization tasks. There also has been some work on stochastic analysis of a system via Markov processes or approximation [31, 32]. We are not aware of any work that considers stochastic parallel tasks.

There has been significant work on purely parallel systems, which are generally built to execute single parallel programs on pre-allocated cores to maximize throughput. Examples include parallel languages and runtime systems, such as the Cilk family [33, 34], OpenMP [35], and Intel's Thread Building Blocks [36]. While multiple tasks on a single platform have been considered in the context of fairness in resource allocation [37], none of this work considers real-time constraints.

III. STOCHASTIC PARALLEL TASK MODEL

In this section, we formalize the **stochastic task model** in which execution time and critical-path length are described using probabilistic distributions, which is consistent with the task model for sequential tasks in existing work on stochastic real-time analysis [30]. We also define the capacity augmentation bound for stochastic tasks with soft real-time tardiness constraint. Throughout this paper, we use the calligraphic letters to represent random variables.

Fig. 1: Example parallel DAG job $\tau_{1,1}$ of task τ_1 with work (total execution time) $c_{1,1} = 12$ and critical-path length $l_{1,1} = 10$. Each node of the job is drawn as a node with execution time written in the center. The directed lines indicate dependences between node. The critical-path is the longest path in the DAG, which is shown by the dotted line.

A parallel task is represented by a directed acyclic graph (DAG), such that each node is a sequential segment of computational work (also called a subtask) and each edge is a dependence between two subtasks. A node is **ready** to execute when all of its predecessors have finished. A task set τ consists of n tasks $\tau = \{\tau_1, \tau_2, ..., \tau_n\}$, each of which can be a parallel or a sequential task. An example of a parallel DAG job is shown in Figure 1.

Like ordinary real-time tasks, stochastic tasks have a fixed relative deadline D_i ($= P_i$, the period, for implicit deadline tasks). However, each stochastic task is described using its **stochastic work** C_i — total execution time on 1 core, and **stochastic critical-path length** $\mathcal{L}_i$ — execution time when it is running on a machine with an in infinite number of cores. Note that both C_i and $\mathcal{L}_i$ are random variables.

In this paper, the internal structure of each DAG task is not assumed or used to derive the schedulability analysis. The execution time of each node from the same task could vary in different jobs (execution instances), which would result in varying execution times and critical-path lengths. Moreover, the internal structure of each job from the same task could also be different each time. For example, a parallel for-loop in a program could have different numbers of iterations given different inputs, resulting in a different DAG structure.

We assume that the expectations $\mathrm{E}[C_i]$ and $\mathrm{E}[\mathcal{L}_i]$ of these random variables are known. Given $\mathrm{E}[C_i]$ and $\mathrm{E}[\mathcal{L}_i]$, we can calculate the *expected utilization* of a stochastic task τ_i as $\mathrm{E}[\mathcal{U}_i] = \mathrm{E}[C_i]/D_i$, and the total expected utilization of the entire task set as $\sum_i \mathrm{E}[\mathcal{U}_i]$.

We now specify a few additional parameters that are needed only if we wish to calculate an upper bound on the tardiness itself or to optimize this tardiness using our third (ILP-based) mapping algorithm. First, for all tasks, we must know the standard deviations δ_{C_i} and $\delta_{\mathcal{L}_i}$ of the execution time and the critical-path length. Second, for low-utilization tasks, we need the finite worst-case execution time $\widehat{c}_i$ for calculating tardiness. Finally, for high-utilization tasks, we need the covariance $\sigma(C_i, \mathcal{L}_i)$ between work and critical-path length.

The exact distributions of C_i and $\mathcal{L}_i$ are not explicitly required in all three schedulability tests. Our linear-time algorithm can calculate mappings that provide bounded tardiness using just these parameters. With the distributions, another algorithm can generate potentially better mappings.

In addition, for analysis purposes, we define some job specific parameters: $c_{i,j}$ is the actual execution time of job j of task i and $l_{i,j}$ is its actual critical-path length; these are drawn from distributions C_i and $\mathcal{L}_i$ respectively. For example, the job $\tau_{1,1}$ in Figure 1 has a total execution time of $c_{1,1} = 12$, which is the sum of the execution time of all nodes. Its critical-path length is $l_{1,1} = 10$, which is the sum of the nodes on the critical-path indicated by the dotted line.

We say that the *release time* of job j of task i is $r_{i,j}$ and its *response time* (or *completion time*) is $t_{i,j}$. Tardiness $T_{i,j}$ of job $\tau_{i,j}$ is defined as $\max(0, t_{i,j} - D_i)$. Tardiness $\mathcal{T}_i$ of a task τ_i is also a random variable; $\mathrm{E}[\mathcal{T}_i]$ is its expected value.

We now define the capacity augmentation bound for stochastic tasks. In particular, we consider the schedulability criterion of **bounded expected tardiness**; that is, a task set τ is deemed *schedulable* by a scheduling algorithm $\mathcal{S}$ if the expected tardiness of each task is guaranteed to be bounded under $\mathcal{S}$.

Definition 1. *A scheduling algorithm $\mathcal{S}$ provides a **stochastic capacity augmentation bound** of b if, given m cores, $\mathcal{S}$ can guarantee bounded expected tardiness to any task set τ as long as it satisfies the following **conditions:***

$$\text{Total available cores, } m \geq b \sum \mathrm{E}[\mathcal{U}_i] \qquad (1)$$

$$\text{For each task, } D_i \geq b(\mathrm{E}[\mathcal{L}_i] + \epsilon_i) \qquad (2)$$

where ϵ_i is 0 if the variances of C_i and $\mathcal{L}_i$ are 0, and is an arbitrarily small positive constant otherwise.

Note that when C_i and $\mathcal{L}_i$ are deterministic, the variance of C_i and $\mathcal{L}_i$ is 0, so $\epsilon_i = 0$ and the definition of stochastic capacity augmentation bound reduces to the definition for hard real-time constraints based on worst-case task parameters.

IV. STOCHASTIC FEDERATED SCHEDULING GUARANTEES BOUNDED TARDINESS

In this section, we first describe stochastic federated scheduling; Then we prove that if federated scheduling algorithm can produce a valid mapping, then it guarantees bounded expected tardiness; Finally, we calculate the expected tardiness.

A. Stochastic Federated Scheduling Strategy

Just like the corresponding federated scheduling strategy for hard real-time tasks, the stochastic federated scheduling strategy classifies tasks into two sets: τ_{high} contains all **high-utilization tasks** — tasks with expected utilization at least 1 ($\mathrm{E}[\mathcal{U}_i] \geq 1$), and τ_{low} contains all the remaining **low-utilization tasks**. The federated scheduling strategy works in two stages:

1) Given a task set τ, a **mapping algorithm** either **admits** τ and outputs a core assignment, or declares that it cannot guarantee schedulability of τ. Different mapping algorithms differ in the assignment of n_i dedicated cores to each high-utilization task τ_i, but $n_i > \frac{\mathrm{E}[C_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$ is always required. All low-utilization tasks share the remaining $n_{\text{low}} = m - \sum_{\tau_i \in \tau_{high}} n_i$ cores. Each mapping algorithm only admits a task set if $n_{\text{low}} > \sum_{\tau_i \in \tau_{\text{low}}} \mathrm{E}[\mathcal{U}_i]$ always holds.

2) Once the mapping is done, the scheduling is straightforward. The high-utilization tasks are scheduled on their dedicated cores using a greedy (work-conserving) scheduler. The low-utilization tasks are scheduled and executed sequentially on the remaining cluster of cores via a GEDF scheduler.

Note that we chose GEDF to schedule low-utilization tasks, because of an existing result that shows that GEDF provides bounded tardiness to sequential stochastic tasks [30]; we can apply this result directly to low-utilization tasks since they are executed sequentially by our federated scheduler. Other multiprocessor scheduling algorithms can be used only if they provide guarantees of bounded tardiness for sequential tasks.

B. Mapping Algorithms Guarantee Bounded Tardiness

We first analyze high-utilization tasks. Since each of them has dedicated cores and does not suffer any interference from other tasks, we can analyze each task τ_i individually. We use the following result from queueing theory [38] which indicates that if the service time of jobs is less than the inter-arrival time, then the expected waiting time is bounded.

Lemma 1. [KING70] *For a D/G/1 queue, customers arrive with minimum inter-arrival time Y, and the service time $\mathcal{X}$ is a distribution with mean $\mathrm{E}\left[\mathcal{X}\right]$ and variance $\delta_{\mathcal{X}}^2$. If $\mathrm{E}\left[\mathcal{X}\right] < Y$, then the queue is stable and the expected waiting time $\mathcal{W}$ is bounded $\mathrm{E}\left[\mathcal{W}\right] \le \frac{\delta_{\mathcal{X}}^2}{2(Y-\mathrm{E}[\mathcal{X}])}$.*

In our context, for each high-utilization task, jobs are the customers; the inter-arrival time is $Y = D_i\, (= P_i)$; the response time $\mathcal{X} = t_{i,j}$ is the service time for job j of task τ_i. For a high-utilization job $\tau_{i,j}$, its tardiness $T_{i,j}$ depends on its response time $t_{i,j}$, the tardiness $T_{i,j-1}$ of previous job $\tau_{i,j-1}$ and deadline D_i. In particular, we have $T_{i,j} \le \max\{0, T_{i,j-1} + t_{i,j} - D_i\}$. Therefore, the waiting time $\mathcal{W}$ is a bound on the tardiness $\mathcal{T}$.

For a greedy scheduler on n_i cores, there are two straightforward lemmas (Lemma 1 and 2) derived in [22]. Using the two Lemmas, we can easily bound the finish time $t_{i,j}$.

Lemma 2. *If a job $J_{i,j}$ executes by itself under a greedy scheduler on n_i identical cores and it takes $t_{i,j}$ time to finish its execution, then $t_{i,j} \le (c_{i,j} + (n_i - 1)l_{i,j})/n_i$.*

Thus the response time for a job is bounded by $(c_{i,j} + (n_i - 1)l_{i,j})/n_i$. Using properties of mean and variance, we get

$$
\begin{aligned}
\mathrm{E}\left[\mathcal{X}\right] = \mathrm{E}\left[t_{i,j}\right] &\le (\mathrm{E}\left[\mathcal{C}_i\right] + (n_i - 1)\mathrm{E}\left[\mathcal{L}_i\right])/ni \quad (3)\\
\delta_{\mathcal{X}}^2 = \delta_{t_{i,j}}^2 &\le \delta_{\mathcal{L}_i}^2((n_i-1)/n_i)^2 + \delta_{\mathcal{C}_i}^2/n_i^2 \\
&\quad + 2\sigma(\mathcal{L}_i, \mathcal{C}_i)(n_i-1)/n_i^2 \quad (4)
\end{aligned}
$$

Note that Lemma 1 states that if $\mathrm{E}\left[\mathcal{X}\right] < Y$, then the queue is stable and the tardiness is bounded. Therefore, to prove the bounded expected tardiness of high-utilization task, we only need to prove $\mathrm{E}\left[\mathcal{X}\right] = (\mathrm{E}\left[\mathcal{C}_i\right] + (n_i-1)\mathrm{E}\left[\mathcal{L}_i\right])/n_i < D_i = Y$.

Theorem 1. *A mapping algorithm for stochastic federated scheduling guarantees bounded tardiness to high-utilization task τ_i, if the assigned number of cores $n_i > \frac{\mathrm{E}[\mathcal{C}_i]-\mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$.*

Proof: We first prove $(\mathrm{E}\left[\mathcal{C}_i\right] + (n_i - 1)\mathrm{E}\left[\mathcal{L}_i\right])/n_i < D_i$.

$$
\begin{aligned}
&D_i n_i - (n_i-1)\mathrm{E}\left[\mathcal{L}_i\right] = n_i(D_i - \mathrm{E}\left[\mathcal{L}_i\right]) + \mathrm{E}\left[\mathcal{L}_i\right]\\
> \; &\frac{\mathrm{E}\left[\mathcal{C}_i\right] - \mathrm{E}\left[\mathcal{L}_i\right]}{D_i - \mathrm{E}\left[\mathcal{L}_i\right]}(D_i - \mathrm{E}\left[\mathcal{L}_i\right]) + \mathrm{E}\left[\mathcal{L}_i\right]\\
> \; &\mathrm{E}\left[\mathcal{C}_i\right]
\end{aligned}
$$

Hence, $\mathrm{E}\left[\mathcal{X}\right] = \mathrm{E}\left[t_{i,j}\right] = (\mathrm{E}\left[\mathcal{C}_i\right] + (n_i - 1)\mathrm{E}\left[\mathcal{L}_i\right])/n_i < D_i = Y$ and by Lemma 1 the tardiness of τ_i is bounded. $\square$

In the stochastic federated scheduling strategy, $n_i > \frac{\mathrm{E}[\mathcal{C}_i]-\mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$ is always required for any mapping algorithm. We will show later that for all three proposed mapping algorithms, this is indeed satisfied for each high-utilization task.

Now we analyze the tardiness of low-utilization tasks, since they share n_{low} cores and are executed sequentially using GEDF scheduler. In [30], the following Lemma has been established.

Lemma 3. [Mills10] *If a set of sequential tasks τ_{low} is scheduled on n_{low} cores using GEDF and $n_{low} > \sum_{\tau_i \in \tau_{low}} \mathrm{E}\left[\mathcal{U}_i\right]$, then the expected tardiness of each task is bounded.*

Since the different mapping algorithms only admit a task set if $\mathrm{E}\left[\mathcal{U}_{\text{low}}\right] = \sum_{\tau_i \in \tau_{\text{low}}} \mathrm{E}\left[\mathcal{U}_i\right] < n_{\text{low}}$ and then schedule these tasks using GEDF, we can conclude that the expected tardiness of low-utilization tasks is also bounded.

Any task set that the mapping algorithm admits can be scheduled while guaranteeing bounded expected tardiness; hence, the mapping algorithm serves as a **schedulability test**.

C. Calculating Expected Tardiness

Here, we explain how the tardiness is calculated. Even though all the mapping algorithms provide bounded expected tardiness, the actual (upper bound on) tardiness can be different, because the corresponding core assignments (n_i for each high-utilization task and n_{low} for all low-utilization tasks) are different.

Note that from Section V, we can see that for the BASIC and FAIR mapping algorithms, the tardiness calculation is not necessary for producing core assignments. It is only needed in ILP mapping or to get the actual expected tardiness.

1) Tardiness of High-Utilization Tasks: For each high-utilization tasks with n_i assigned dedicated cores, by Lemma 1 and Inequality (4), the bounded expected tardiness is:

$$
\begin{aligned}
\mathrm{E}\left[\mathcal{T}_i\right] &\le \frac{\delta_{\mathcal{X}}^2}{2(Y - \mathrm{E}\left[\mathcal{X}\right])}\\
&\le \frac{\delta_{\mathcal{L}_i}^2(n_i-1)^2/n_i^2 + \delta_{\mathcal{C}_i}^2/n_i^2 + 2\sigma(\mathcal{L}_i,\mathcal{C}_i)(n_i-1)/n_i^2}{2(D_i - (\mathrm{E}\left[\mathcal{L}_i\right](n_i-1) + \mathrm{E}\left[\mathcal{C}_i\right])/n_i)} \quad (5)
\end{aligned}
$$

2) Tardiness of Low-Utilization Tasks: Since low-utilization tasks are executed sequentially using GEDF, we can use the linear-programming procedure described in [30] directly.

We first restate a couple of lemmas from [30] in our terminology. The first lemma bounds the tardiness of a hypothetical processor-sharing (**PS**) scheduler which always guarantees an execution rate of $\hat{u}_i$ (henceforth called the **PS rate allocation**) to each task τ_i.

Lemma 4. [Mills10] *For a given PS rate allocation such that* $\mathrm{E}\left[\mathcal{U}_i\right] \leq \hat{u}_i \leq 1$ *and* $\sum \mathrm{E}\left[\mathcal{U}_i\right] \leq n_{low}$, *the PS scheduler has a bounded tardiness* $\mathrm{E}\left[\mathcal{F}_i\right] \leq \frac{\delta_{\mathcal{C}_i}^2 / \hat{u}_i^2}{2(D_i - \mathrm{E}[\mathcal{C}_i]/\hat{u}_i)}$.

Using this tardiness bound, they then provide a bound on the tardiness provided by GEDF for low-utilization tasks.

Lemma 5. [Mills10] *For low-utilization tasks scheduled by a GEDF scheduler on* n_{low} *cores, the expected tardiness of each task* $\mathrm{E}\left[\mathcal{T}_i\right] \leq \mathrm{E}\left[\mathcal{F}_i\right] + \frac{\eta + n_{low} M}{n_{low} - v} + \hat{c}_i$, *where* $\mathrm{E}\left[\mathcal{F}_i\right]$ *is the expected tardiness of a hypothetical PS scheduler,* $\hat{c}_i$ *is the worst-case execution time of the task,* η *is the sum of the* $n_{low} - 1$ *largest* $\hat{c}_i$, M *is the maximum tardiness in PS, and* v *is the sum of* $n_{low} - 1$ *largest assigned* $\hat{u}_i$ *in PS.*

All the parameters except $\mathrm{E}\left[\mathcal{F}_i\right]$ are known or measurable (and bounded). In order to calculate $\mathrm{E}\left[\mathcal{F}_i\right]$, we must calculate the PS rate allocation $\hat{u}_i$ for each task τ_i.

As we will show in Section V, for the BASIC mapping, there exists a simple calculation of $\hat{u}_i$; while for FAIR and ILP mappings, the following linear program (LP) from [30] (can be derived using Lemma 4) is used to calculate the PS rate allocations.

$$\max \quad \zeta$$
$$\text{s.t.} \quad D_i \hat{u}_i - \frac{\delta_{\mathcal{C}_i}^2}{2}\zeta \geq \mathrm{E}\left[\mathcal{C}_i\right] \quad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] < 1$$
$$\sum_{i, \mathrm{E}[\mathcal{U}_i] < 1} \hat{u}_i \leq \hat{n}_{\text{low}}$$
$$u_i \leq \hat{u}_i \leq 1 \quad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] < 1$$

where $\zeta^{-1} \geq \max_i \left(\frac{\delta_{\mathcal{L}_i}^2}{2(\hat{u}_i D_i - \mathrm{E}[\mathcal{L}_i])}\right) = \max_i \mathrm{E}\left[\mathcal{F}_i\right]$. Therefore, solving the linear program provides the PS rate allocations $\hat{u}_i$ as well as a bound on the expected tardiness $\mathrm{E}\left[\mathcal{F}_i\right]$ of the PS scheduler. Given these values, we can calculate the tardiness of low-utilization tasks using Lemma 5.

V. MAPPING ALGORITHMS FOR STOCHASTIC FEDERATED SCHEDULING

We propose three mapping algorithms for stochastic federated scheduling. The three algorithms differ in their calculation of n_i for high-utilization tasks. They have increasing computational complexity and also have increasing schedulability performance or decreasing upper bound on expected tardiness: The first algorithm, **BASIC**, assigns cores based on utilization; The second algorithm, **FAIR**, assumes that the distributions of execution time and critical-path length are known and assigns cores based on the values with the same cumulative probability from task parameter distributions among all tasks; The last (ILP-Based) algorithm, (**ILP**), tries to minimize the maximum expected tardiness.

A. BASIC Federated Mapping Algorithm

For a high-utilization tasks τ_i, this mapping algorithm calculates n_i, the number of cores assigned to τ_i as follows:

$$n_i = \begin{cases} \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} \right\rceil & (\mathrm{E}\left[\mathcal{U}_i\right] > 1) \\ 2 & (\mathrm{E}\left[\mathcal{U}_i\right] = 1) \end{cases} \quad (6)$$

where $\alpha_i = D_i/b - \mathrm{E}\left[\mathcal{L}_i\right] > 0$ and $b = 2$.

The remaining $n_{\text{low}} = m - \sum_{\text{high}} n_i$ cores are assigned to the low-utilization tasks. The mapping algorithm admits a task set as long as $\mathrm{E}\left[\mathcal{U}_{\text{low}}\right] = \sum_{\text{low}} \mathrm{E}\left[\mathcal{U}_i\right] \leq n_{\text{low}}/b$ for $b = 2$.

Note that the major difference between this n_i and the one in [2] is the extra term α_i. α_i is used to accommodate the variation of execution time and critical-path length. We set this value of α_i to assign roughly the same number of cores relative to utilization. Hence, variances are not required to assign cores.

Bounded Tardiness (Schedulability Test): The tardiness can be bounded for any positive α_i since:
1) For $\mathrm{E}\left[\mathcal{U}_i\right] = 1$, $\frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} = 1$, so $n_i = 2 > \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$.
2) For $\mathrm{E}\left[\mathcal{U}_i\right] > 1$, since $D_i - \mathrm{E}\left[\mathcal{L}_i\right] > \alpha_i > 0$, we have

$$n_i \geq \tfrac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} > \tfrac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} > \tfrac{\mathrm{E}[\mathcal{C}_i]}{D_i} = \mathrm{E}\left[\mathcal{U}_i\right] > 1$$

3) For $\mathrm{E}\left[\mathcal{U}_i\right] < 1$, $\mathrm{E}\left[\mathcal{U}_{\text{low}}\right] \leq n_{\text{low}}/2 < n_{\text{low}}$.

By Theorem 1 and Lemma 3, the BASIC mapping can guarantee bounded tardiness for both high and low-utilization tasks. Therefore, the BASIC algorithm serves as a schedulability test that runs in linear time.

Tardiness calculation: Now we describe a faster and simpler method to calculate the upper bound on the expected tardiness of low-utilization tasks when using the BASIC mapping. This method relies on the requirement that $n_{\text{low}} \geq b \sum_{\text{low}} \mathrm{E}\left[\mathcal{U}_i\right]$ for $b = 2$. We can simply set PS rate allocation $\hat{u}_i = \min\left(b\mathrm{E}\left[\mathcal{U}_i\right], 1\right)$. This allocation satisfies the requirement in Lemma 4; therefore, the PS tardiness is

$$\mathrm{E}\left[\mathcal{F}_i\right] \leq \frac{\delta_{\mathcal{C}_i}^2}{2(\hat{u}_i^2 D_i - \hat{u}_i \mathrm{E}\left[\mathcal{C}_i\right])},$$

and by Lemma 5 the expected tardiness of low-utilization task under GEDF can be calculated directly as

$$\mathrm{E}\left[\mathcal{T}_i\right] \leq \frac{\delta_{\mathcal{C}_i}^2}{2(\hat{u}_i^2 D_i - \hat{u}_i \mathrm{E}\left[\mathcal{C}_i\right])} + \frac{\eta + n_{\text{low}} M}{n_{\text{low}} - v} + \hat{e}_i, \quad (7)$$

Unlike the FAIR and ILP algorithms, this tardiness calculation here does not require solving a linear program; it can be done in linear time.

B. FAIR Federated Mapping Algorithm

We now present the FAIR mapping, which admits more task sets than the BASIC one, while still providing the same theoretical guarantees. The FAIR mapping utilizes the distributions of execution time and critical-path length and assigns cores based on the values with the same cumulative probability from distributions among all tasks to provide fairness in core assignment. The schedulability test in FAIR still runs in linear time; however, the calculations for core assignment and expected tardiness are more complex, requiring near linear time and linear programming respectively.

We denote $\mathcal{C}_i(p)$ as the value c_i of random variable $\mathcal{C}_i$ when its cumulative distribution function (CDF) $\mathbb{F}_{\mathcal{C}_i}(c_i) = p$ (meaning that the probability that $\mathcal{C}_i \leq c_i$ is equal to p). We denote $\mathcal{L}_i(p)$ and $\mathcal{U}_i(p)$ similarly.

Note that when $p = 0.5$, $\mathcal{C}_i(p) = \mathrm{E}\left[\mathcal{C}_i\right]$ and $\mathcal{L}_i(p) = \mathrm{E}\left[\mathcal{L}_i\right]$. Additionally, $\mathcal{C}_i(p)$ and $\mathcal{L}_i(p)$ will increase when p increases.

In the FAIR mapping, the number of cores assigned to high-utilization task τ_i (represented by $\hat{n}_i$) is calculated below.

$$\hat{n}_i(p) = \left\lfloor \frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} + 1 \right\rfloor \qquad (8)$$

$$= \begin{cases} \left\lceil \frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} \right\rceil & \left(\frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} \text{ is not integer} \right) \\ \frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} + 1 & \left(\frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} \text{ is integer} \right) \end{cases}$$

where p is the same probability for all tasks and $0.5 \leq p < 1$.

The FAIR mapping will admit a task set if $n_{\text{low}} = m - \sum_{\text{high}} \hat{n}_i(p) > \sum_{\text{low}} \mathrm{E}\left[\mathcal{U}_i(p)\right]$ for $p = 0.5$.

Bounded Tardiness (Schedulability Test): It is obvious that for $p = 0.5$, each $\hat{n}_i(p = 0.5) = \left\lfloor \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} + 1 \right\rfloor > \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$ for high-utilization task. Also, for all low-utilization tasks, $n_{\text{low}} > \sum_{\text{low}} \mathrm{E}\left[\mathcal{U}_i(p = 0.5)\right] = \sum_{\text{low}} \mathrm{E}\left[\mathcal{U}_i\right]$. By Theorem 1 and Lemma 3, FAIR guarantees bounded tardiness for all tasks and serves as a linear time schedulability test.

Dominance in Schedulability: In Section VI, we will show that $\hat{n}_i(p = 0.5) \leq n_i$ (of BASIC mapping) for any task τ_i and hence $\hat{n}_{low} \geq n_{low}$. Also, the FAIR algorithm allows $\mathrm{E}\left[\mathcal{U}_{low}\right]$ to be as high as $\hat{n}_{low}$ (instead of $n_{low}/2$ allowed by BASIC). Therefore, FAIR admits strictly more tasks than BASIC.

Core Allocation: $\hat{n}_i(p = 0.5)$, the *minimum core assignment*, is the minimum number of cores required to guarantee bounded tardiness for high-utilization tasks. However, directly using it will result in large tardiness for high-utilization tasks, because more cores are assigned to low-utilization tasks. To be fair to all tasks, the FAIR mapping further improves the *minimum core allocation* by increasing p until the largest $\hat{p}$ while still satisfying $n_{\text{low}} = m - \sum_{\text{high}} \hat{n}_i(\hat{p}) > \sum_{\text{low}} (\mathcal{C}_i(\hat{p})/D_i)$. In this way, FAIR in fact increases the core assignment and PS rate allocation for each task by the same amount according to the CDF of execution time and critical-path length. This ensures fairness among all tasks, because $\hat{p}$ is independent of τ_i. The complexity of this core assignment depends on the number of values of p tested until reaching $\hat{p}$. In practice, a binary search only needs at most 6 tests to find $\hat{p}$ with an accuracy of 0.01.

C. ILP-Based Federated Mapping Algorithm

We now present a third, ILP-Based, mapping algorithm for stochastic federated scheduling. This algorithm admits exactly the same task sets as FAIR (though it may find a different mapping for these task sets); therefore, it also provides the same theoretical guarantees. However, BASIC and FAIR make no attempt to balance maximum tardiness explicitly among high and low-utilization tasks. Compared to FAIR, the ILP mapping does not require the distributions of execution time and critical-path length. Instead, the stander deviations and covariance are used in ILP.

The ILP algorithm converts the mapping problem for high-utilization tasks into an *integer linear program* (ILP) that tries to minimize the maximum tardiness; when combined with the linear program for low-utilization tasks stated in Section IV-C2, the resulting mixed linear program indirectly tries to balance the tardiness among all tasks.

We convert Inequality (5) into a form similar to the expected tardiness of the PS schedule; that is, we define ζ_i where $\zeta^{-1} = \max_i \mathrm{E}\left[\mathcal{T}_i\right]$ and ζ is defined in terms of n_i. First, for task τ_i, let $\delta_i^2 = max\big(\delta_{\mathcal{L}_i}^2(m-1)^2/m, \delta_{\mathcal{C}_i}^2/2, \sigma(\mathcal{L}_i, \mathcal{C}_i)(m-1)/m\big)$. Note that δ_i^2 is bounded and can be calculated using only the expectation and variance of the task's execution time and critical-path length without knowing n_i. Now we use the fact that $2 \leq n_i \leq m$ for high-utilization task τ_i and see that

$$\begin{aligned} \delta_i^2 &\geq& \delta_{\mathcal{L}_i}^2 (m-1)^2/m = \delta_{\mathcal{L}_i}^2 (m-1)(1-1/m) \\ &\geq& \delta_{\mathcal{L}_i}^2 (n_i-1)(1-1/n_i) = \delta_{\mathcal{L}_i}^2 (n_i-1)^2/n_i \\ \delta_i^2 &\geq& \delta_{\mathcal{C}_i}^2/2 \geq \delta_{\mathcal{C}_i}^2/n_i \\ \delta_i^2 &\geq& \sigma(\mathcal{L}_i, \mathcal{C}_i)(m-1)/m = \sigma(\mathcal{L}_i, \mathcal{C}_i)(1-1/m) \\ &\geq& \sigma(\mathcal{L}_i, \mathcal{C}_i)(1-1/n_i) = \sigma(\mathcal{L}_i, \mathcal{C}_i)(n_i-1)/n_i. \end{aligned}$$

Now we calculate the upper bound on the variance of $\delta_{t_{i,j}}^2$ (from Inequality (4)) using δ_i^2

$$\begin{aligned} \delta_{t_{i,j}}^2 &=& \delta_{\mathcal{L}_i}^2 (n_i-1)^2/n_i^2 + \delta_{\mathcal{C}_i}^2/n_i^2 + 2\sigma(\mathcal{L}_i, \mathcal{C}_i)(n_i-1)/n_i^2 \\ &=& \frac{\delta_{\mathcal{L}_i}^2 (n_i-1)^2/n_i + \delta_{\mathcal{C}_i}^2/n_i + 2\sigma(\mathcal{L}_i, \mathcal{C}_i)(n_i-1)/n_i}{n_i} \\ &\leq& 4\delta_i^2/n_i \end{aligned}$$

By Corollary 1, the expected tardiness is bounded by

$$\begin{aligned} \mathrm{E}\left[\mathcal{T}_i\right] &\leq& \frac{\delta_{\mathcal{X}}^2}{2(Y - \mathrm{E}\left[\mathcal{X}\right])} \\ &\leq& \frac{4\delta_i^2/n_i}{2(D_i - (\mathrm{E}\left[\mathcal{L}_i\right](n_i-1) + \mathrm{E}\left[\mathcal{C}_i\right])/ni)} \\ &\leq& \frac{2\delta_i^2}{n_i D_i - (\mathrm{E}\left[\mathcal{L}_i\right](n_i-1) + \mathrm{E}\left[\mathcal{C}_i\right])} \\ &=& \frac{2\delta_i^2}{n_i(D_i - \mathrm{E}\left[\mathcal{L}_i\right]) - (\mathrm{E}\left[\mathcal{C}_i\right] - \mathrm{E}\left[\mathcal{L}_i\right])} \qquad (9) \end{aligned}$$

Now we can set $\zeta^{-1} \geq \max_i \big(\frac{2\delta_i^2}{n_i(D_i - \mathrm{E}[\mathcal{L}_i]) - (\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i])}\big) \geq \max_i \mathrm{E}\left[\mathcal{T}_i\right]$ for high-utilization tasks and get inequality (11).

Combining this definition of ζ with the linear program in Section IV-C2, we get the following mixed linear program:

$$\begin{aligned} \max \quad & \zeta \\ \text{s.t.} \quad & D_i \hat{u}_i - \frac{\delta_{\mathcal{C}_i}^2}{2}\zeta \geq \mathrm{E}\left[\mathcal{C}_i\right] \qquad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] < 1 \;(10) \\ & (D_i - \mathrm{E}\left[\mathcal{L}_i\right])n_i - 2\delta_i^2 \zeta \geq \mathrm{E}\left[\mathcal{C}_i\right] - \mathrm{E}\left[\mathcal{L}_i\right] \\ & \qquad\qquad\qquad\qquad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] \geq 1 \;(11) \\ & \sum_{i, \mathrm{E}[\mathcal{U}_i]<1} \hat{u}_i + \sum_{i, \mathrm{E}[\mathcal{U}_i]\geq 1} n_i \leq m \qquad (12) \\ & u_i \leq \hat{u}_i \leq 1 \qquad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] < 1 \;(13) \\ & \hat{n}_i(p = 0.5) \leq n_i \qquad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] \geq 1 \;(14) \\ & n_i \text{ is integer} \qquad \forall i, \mathrm{E}\left[\mathcal{U}_i\right] \geq 1 \;(15) \end{aligned}$$

We solve this ILP to calculate: integral n_i— the number of cores assigned to high utilization task τ_i; fractional $\hat{u}_i$ — a valid PS rate allocation for low-utilization task τ_i; and ζ. Using the resulting n_i for high utilization tasks, we can calculate

$n_{\text{low}} = m - \sum_{\text{high}} n_i$, the number of cores assigned to low-utilization tasks.

Explanation of Constraints: Constraints (14) and (15) guarantee that each high-utilization task τ_i gets at least $\hat{n}_i(p = 0.5)$ dedicated cores; therefore Theorem 1 guarantees its bounded tardiness. Constraint (13) guarantees that the PS rate allocation is larger than the utilization of low-utilization tasks; therefore Lemma 4 guarantees bounded tardiness of these tasks. Constraint (12) guarantees that $n_{\text{low}} + n_{\text{high}} \leq m$. Finally, Constraint (10) is inherited from the LP in Section IV-C2.

Optimal Greedy Solution to the ILP: General ILP problems can be hard to solve. However, there is a unique property of the above ILP — ζ will decrease if at least one n_i or $\sum_{\text{low}} \hat{u}_i$ increases and the rest remain the same. Relying on this, we can easily see that a greedy algorithm — starting with the core assignment (n_i and $\hat{u}_i(p = 0.5)$) from the minimum core allocation of the FAIR mapping, iteratively increases the n_i or $\sum_{\text{low}} \hat{u}_i$ (a high utilization task or the sum of low utilization tasks) with largest tardiness by 1 until just before Constraint (12) would not hold — will successfully find the optimal solution to this ILP problem (provided that the LP in Section IV-C2 can directly calculate an optimal solution). By applying the greedy solution, we can reduce the mixed-ILP problem to an iterative LP problem. Obviously, the maximum number of iterations needed by the greedy algorithm is m.

Relationship to FAIR: The ILP mapping algorithm admits exactly the same task sets that FAIR admits: if FAIR admits a task set ($\hat{n}_i(p = 0.5)$ and $n_{\text{low}} = m - \sum_{\text{high}} \hat{n}_i(p = 0.5)$), then that mapping is a trivially feasible solution to the ILP since it satisfies all constraints for $\zeta = 0$. On the other hand, if the FAIR algorithm cannot find a solution, then there is no feasible solution to the ILP. Therefore, since FAIR provides a capacity augmentation bound of 2, so does this algorithm.

Faster Schedulability Test: As a consequence of the relationship with FAIR, we do not have to solve the ILP to check if the task set is schedulable using this ILP-based mapping; we can simply run FAIR to check for schedulability and only solve the ILP to find the mapping if the task set is, in fact, schedulable.

Tardiness Calculation: On solving the mixed linear program, we get n_i for each high utilization task and $\hat{u}_i$ for each low utilization task. Therefore, we can use Inequalities (5) and (7) to calculate the tardiness of these tasks, respectively.

Note that the mixed linear program criterion is a little imprecise; maximizing ζ does not directly optimize the overall tardiness bound. Instead, it only tries to balance parts of the tardiness. After applying Inequalities (7) and (5) for calculating tardiness, the resulting tardiness of each high-utilization task is actually less than the optimized bound ζ^{-1}, while the tardiness of low-utilization tasks is actually higher than ζ^{-1}.

To further balance the overall tardiness, instead of using the strict upper bound of $\delta^2_{t_{i,j}}$ (from Inequality (9)) in the calculation of ζ, we can approximate it. The reason we cannot directly use Inequality (4) to calculate $\delta^2_{t_{i,j}}$ is because we do not know n_i before we solve the integer linear program. However, we can approximate $\delta^2_{t_{i,j}}$ by using $\hat{n}_i(p = 0.5)$ instead of n_i. Then,

we have $\delta^2_{t_{i,j}} = \frac{\delta^2_{\mathcal{L}_i}(\hat{n}_i - 1)^2/\hat{n}_i + \delta^2_{\mathcal{C}_i}/\hat{n}_i + 2\sigma(\mathcal{L}_i, \mathcal{C}_i)(\hat{n}_i - 1)/\hat{n}_i}{n_i} = \frac{\delta^2_i}{n_i}$. This may provide a better tardiness bound for all tasks.

However, when the worst-case execution time of a low-utilization task is large, the achieved mapping may still result in a larger maximum tardiness (for that task) than optimal.

VI. STOCHASTIC CAPACITY AUGMENTATION BOUND OF 2 FOR STOCHASTIC FEDERATED SCHEDULING

A. Stochastic Capacity Augmentation Bound for BASIC

Theorem 2. *The BASIC federated scheduling algorithm has a stochastic capacity augmentation bound of $b = 2$.*

In order to prove Theorem 2, we first prove that the BASIC mapping strategy always admits all *eligible* task sets — task sets that satisfy Conditions (1) and (2) in Definition 1 for $b = 2$.

BASIC admits a task set if, $\mathrm{E}[\mathcal{U}_{\text{low}}] \leq n_{\text{low}}/b$ for $b = 2$. Therefore, we must prove that for all task sets that satisfy Conditions (1) and (2), n_{low} is large enough for BASIC to admit the task set.

First, we prove that the number of cores assigned to high-utilization tasks n_{high} is bounded by $b \sum_{\text{high}} \mathrm{E}[\mathcal{U}_i]$.

Lemma 6. *For a high-utilization task τ_i ($1 \leq \mathrm{E}[\mathcal{U}_i]$), if $D_i > b\mathrm{E}[\mathcal{L}]_i$ (Condition (2)), then the number of assigned cores $n_i \leq b\mathrm{E}[\mathcal{U}_i]$ with $b = 2$.*

Proof: For $\mathrm{E}[\mathcal{U}_i] > 1$, since $b(\mathrm{E}[\mathcal{L}_i] + \alpha_i) = D_i$, so

$$\mathrm{E}[\mathcal{C}_i] = b(\mathrm{E}[\mathcal{L}_i] + \alpha_i)\mathrm{E}[\mathcal{U}_i]$$
$$\Rightarrow \quad D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i = (b-1)(\mathrm{E}[\mathcal{L}_i] + \alpha_i)$$

Therefore, we can bound n_i by

$$
\begin{aligned}
n_i &= \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} \right\rceil \\
&< \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} + 1 \\
&= \frac{2(\mathrm{E}[\mathcal{L}_i] + \alpha_i)\mathrm{E}[\mathcal{U}_i] - (\mathrm{E}[\mathcal{L}_i] + \alpha_i)}{\mathrm{E}[\mathcal{L}_i] + \alpha_i} + 1 \\
&= 2\mathrm{E}[\mathcal{U}_i]
\end{aligned}
$$

For $\mathrm{E}[\mathcal{U}_i] = 1$, $n_i = 2 = 2\mathrm{E}[\mathcal{U}_i]$. Therefore, $n_{\text{high}} = \sum_{\text{high}} n_i \leq b \sum_{\text{high}} \mathrm{E}[\mathcal{U}_i]$ for $b = 2$. $\square$

Since task set τ satisfies Condition (1), the total utilization $\sum \mathrm{E}[\mathcal{U}_i] \leq m/b$ for $b = 2$. So we have

$$n_{\text{low}} = m - n_{\text{high}} \geq b \sum_i \mathrm{E}[\mathcal{U}_i] - b \sum_{\text{high}} \mathrm{E}[\mathcal{U}_i] = b \sum_{\text{low}} \mathrm{E}[\mathcal{U}_i]$$

Hence, BASIC's admission criterion is satisfied and it admits any task set satisfying Conditions (1) and (2). Since BASIC always provides bounded tardiness to task sets it admits (Section IV-B), by Definition 1 this establishes Theorem 2.

B. Stochastic Capacity Augmentation Bound for FAIR

Theorem 3. *The FAIR federated scheduling algorithm has a stochastic capacity augmentation bound of $b = 2$.*

To prove Theorem 3, we simply prove if the BASIC admits a task set, then FAIR does as well; since BASIC admits any task

set that satisfies Conditions (1) and (2) of Definition 1 for $b = 2$, FAIR also admits them. Since FAIR always provides bounded tardiness to task sets it admits, this establishes Theorem 3.

First, we show that the minimum core assignment $\hat{n}_i(p = 0.5)$ to each high-utilization task by the FAIR algorithm is at most the number of cores n_i that the BASIC algorithm assigns.

Lemma 7. *If* $n_i = \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} \right\rceil$ *for* $\mathrm{E}[\mathcal{U}_i] > 1$ *and* $n_1 = 2$ *for* $\mathrm{E}[\mathcal{U}_i] = 1$; *and* $\hat{n}_i(p = 0.5) = \left\lfloor \frac{\mathcal{C}_i(p) - \mathcal{L}_i(p)}{D_i - \mathcal{L}_i(p)} + 1 \right\rfloor = \left\lfloor \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} + 1 \right\rfloor$; *then* $\hat{n}_i \leq n_i$ *for all* $\mathrm{E}[\mathcal{U}_i] \geq 1$.

Proof: To make the proof straightforward, we use the two cases from our definition of $\hat{n}_i$ in Section V.

For $\mathrm{E}[\mathcal{U}_i] > 1$, obviously $\frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} > \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} > 1$, since $D_i - \mathrm{E}[\mathcal{L}_i] > \alpha_i > 0$. So we denote $\epsilon > 0$ and

$$\frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} = \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} + \epsilon$$

When $\frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$ is not integer,

$$\begin{aligned} \hat{n}_i(p = 0.5) &= \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} \right\rceil \\ &\leq \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} \right\rceil = n_i \end{aligned}$$

When $\frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]}$ is integer, since $\epsilon > 0$,

$$\begin{aligned} n_i &= \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i] - \alpha_i}{D_i - \mathrm{E}[\mathcal{L}_i] - \alpha_i} \right\rceil = \left\lceil \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} + \epsilon \right\rceil \\ &\geq \frac{\mathrm{E}[\mathcal{C}_i] - \mathrm{E}[\mathcal{L}_i]}{D_i - \mathrm{E}[\mathcal{L}_i]} + 1 = \hat{n}_i(p = 0.5) \end{aligned}$$

For $\mathrm{E}[\mathcal{U}_i] = 1$, $\hat{n}_i = 2 = n_i$. Therefore, for all cases, $\hat{n}_{\mathrm{high}} = \sum_{\mathrm{high}} \hat{n}_i \leq \sum_{\mathrm{high}} n_i = n_{\mathrm{high}}$. $\square$

FAIR has more cores available for low utilization tasks than BASIC does, since $\hat{n}_{\mathrm{low}}(p = 0.5) = m - \hat{n}_{\mathrm{high}}(p = 0.5) \geq m - n_{\mathrm{high}} = n_{\mathrm{low}}$. It also allows the total utilization of low-utilization tasks to be as high as $\hat{n}_{\mathrm{low}}(p = 0.5)$, while basic only allows it to be n_{low}/b. Therefore, FAIR admits any task set that BASIC admits.

Note that FAIR will only increase $\hat{n}_i$ to $\hat{n}_i(p > 0.5)$ if it admits the task set. Thus, as far as the schedulability and capacity augmentation bound are concerned, this will not affect the proof above. In the most loaded case, $\hat{n}_i(\hat{p}) = \hat{n}_i(p = 0.5)$.

VII. Numerical Evaluation

To compare the different performances of these schedulability tests for *stochastic task sets*, here, we present a numerical evaluation on randomly generated task sets with probability distributions on execution time and critical-path length.

A. Task Sets Generation and Experimental Setup

We evaluate the schedulability results on a varying number of cores m: $4, 8, 16, 32,$ and 64. For various total task set utilizations U starting from 10% to 80%, we generate task sets, add tasks and load the system to be exactly mU — fully loading a unit speed machine. Results for 4 and 64 cores are similar to the rest, so we omit them for brevity.

For each task, we assume normal distributions of execution time and critical-path length. We uniformly generate the expected execution time $\mathrm{E}[\mathcal{C}_i]$ between 1 and 100. For tasks with small variance, we uniformly generate variance to be from 5% to 10% of $\mathrm{E}[\mathcal{C}_i]$; for tasks with large variance, we let it be from 5% to 500%. We generate the critical-path length following the same rules and ensure the average parallelism $\mathrm{E}[\mathcal{C}_i]/\mathrm{E}[\mathcal{L}_i]$ is 32. To ensure a reasonable amount of high-utilization tasks in a task set on m cores, we uniformly generate the task utilization u_i between 0.4 to $\sqrt{m}$. Since we assume a normal distribution for execution time and critical-path length, with the expected mean and standard deviation, we can calculate the worst-case execution time by calculating the value $\hat{c}_i$ of the distribution when the probability of a longer execution time is less than 0.01. Deadline is calculated by $u_i \mathrm{E}[\mathcal{C}_i]$

Using the task set setups above, we run each setting for 100 task sets. We conduct two sets of experiments:

1) We want to evaluate the performance of the two schedulability tests: BASIC and FAIR. In addition, we use the simple schedulability test from the stochastic capacity augmentation bound as a baseline comparison.
2) We want to evaluate the different tardiness bounds of each individual task using different federated mapping algorithms. For task sets that are schedulable according the BASIC test, we record the maximum, mean and minimum tardiness of each task set.

B. Experiment Results

1) Schedulability Performance: We evaluate the performances of different schedulability tests: BOUND (as a baseline), BASIC and FAIR. Note that, as we have proved, schedulability with the ILP mapping algorithm is exactly the same as with the FAIR mapping algorithm (denoted as FAIR/ILP in the figure). Also, since the exact variance of a task is not needed to run these schedulability tests, the calculated schedulability of task sets with small variance and large variance is the same. Therefore, we do not include these curves in the figures.

From Figure 2, we can see that across different numbers of cores, the FAIR/ILP algorithm performs the best, while BOUND performs the worst. Even though the bound indicates that task sets with total utilization larger than $50\% m$ may not be schedulable in terms of bounded tardiness, the two other linear time schedulability tests can still admit task sets up to around 60% for BASIC and 80% for FAIR.

Also note that some task sets with 10% utilization are deemed unschedulable by BOUND. This is due to the critical-path length requirement for parallel tasks by BOUND. For a few tasks with 100% utilization, the FAIR algorithm still guarantees bounded tardiness, because all tasks in the set are low-utilization tasks, and the GEDF scheduler can ensure bounded tardiness for sequential tasks with no utilization loss.

2) Tardiness of Tasks with Small and Large Variance: For task sets for which bounded tardiness is guaranteed, we would like to compare the guaranteed expected tardiness. Note that both the LP and ILP optimization in the FAIR and ILP mapping algorithms only try to optimize the maximum tardiness of the

Fig. 2: Task Set Utilization vs. Schedulability Ratio (both in percentages) for different number of cores.
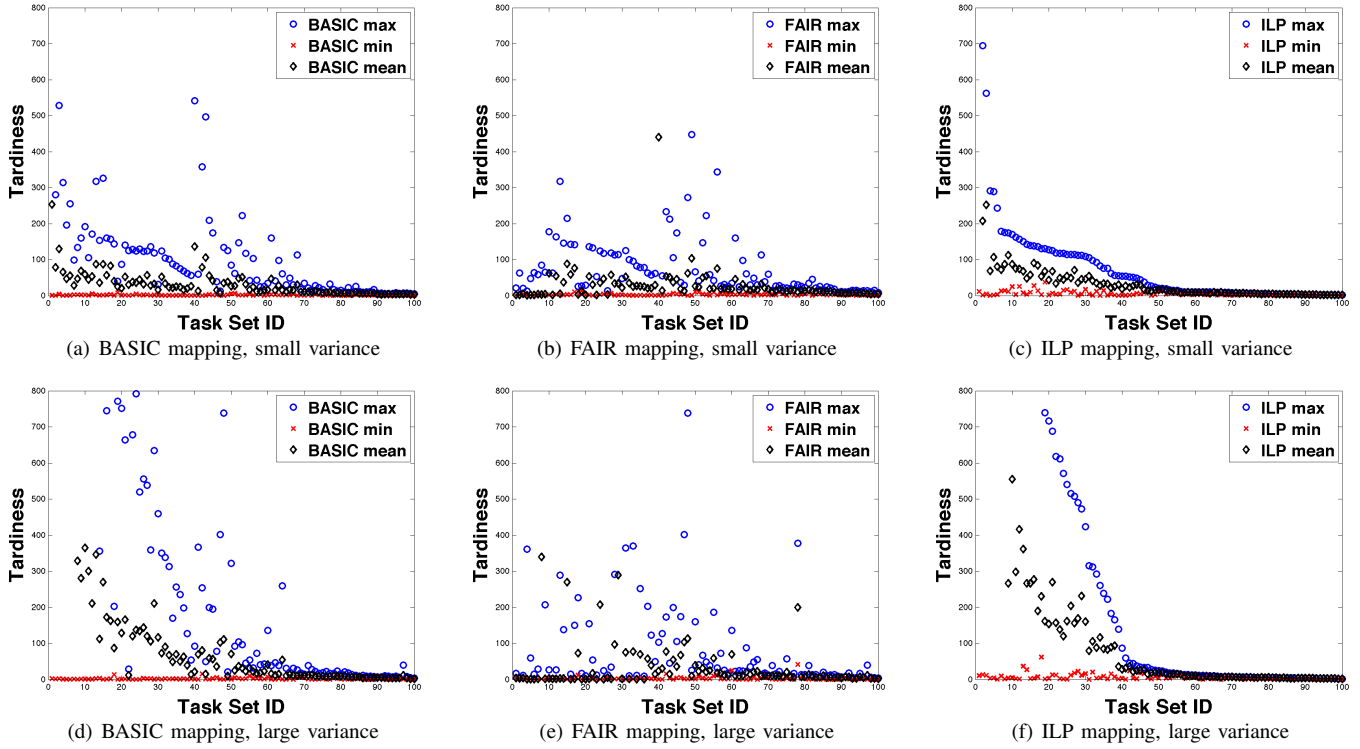

Fig. 3: Maximum, mean and minimum tardiness for parameters with small and large variances.

entire task sets. Therefore, it would be more interesting to see the different expected tardiness bound for the individual tasks.

Figure 3 shows the maximum, mean and minimum expected tardiness calculated from the BASIC, FAIR and ILP mappings for task sets with small and large execution time variations respectively. To make it easy to compare them, we sort all the figures according to the maximum tardiness of the ILP mapping for that corresponding setting (low and high variances).

Not surprisingly, BASIC performs the worst among all three mappings, if we count the number of task sets for which BASIC generates the largest maximum tardiness. In fact, out of all randomly generated task sets, 92% and 85% of the task sets have smaller maximum tardiness with ILP than with BASIC, given small and large variance respectively. Compare FAIR and BASIC, 58% and 76% respectively have lower maximum tardiness under FAIR.

However, we can also see that the maximum tardiness from

the BASIC mapping is comparable to (only slightly worse than) that from the FAIR mapping, when variances of execution time and critical-path length are small. It is also comparable to ILP when the variances are large. This is probably because all compared task sets satisfy the requirement of the bound. Therefore, there are enough cores for BASIC mapping to approximate the better core assignment. Hence, when the variations are small, one could use the BASIC mapping to bound the tardiness.

We also find that with large variances, the increase of maximum tardiness with FAIR is not significant, compared to BASIC and ILP. This is not surprising for BASIC result, because it confirms our hypothesis that the BASIC mapping does not take variation into account when allocating cores. However, ILP does try to balance the tardiness of all tasks, considering variance similarly to FAIR.

In fact, comparing FAIR and ILP, we notice that 67% and

58% task sets respectively have smaller maximum tardiness using ILP. ILP results seem much worse with large variances, only because for some task sets, the maximum tardiness comes from low-utilization tasks. Even through ILP can minimize the tardiness for high-utilization tasks, the LP calculation for low-utilization tasks only minimize the part of tardiness (the maximum tardiness of the PS scheduler in Lemma 4) but cannot directly minimize the overall tardiness in Lemma 5. As FAIR inflates the parameters for low-utilization tasks, the LP calculation may result in a better PS rate allocation and hence smaller tardiness.

VIII. Conclusions

This paper evaluates the soft real-time performance of federated scheduling for parallel real-time tasks under a stochastic task models. It provides a stochastic capacity augmentation bound of 2 for stochastic tasks with a soft real-time constraint of bounded expected tardiness. This is the first such result on stochastic parallel tasks.

The federated scheduling strategy is promising due to its simplicity since it separately schedules high-utilization tasks on dedicated cores and low-utilization cores on shared cores; therefore, one can potentially use out-of-the-box schedulers in a prototype implementation. It would be promising to implement the federated scheduling strategy on a real parallel platform to explore and quantify its performance.

In a more theoretical direction, while it doesn't make sense to put hard real-time constraints on stochastic tasks, we should definitely consider bounded tardiness scheduling of tasks with worst-case task parameters. Federated scheduling cannot be used for this objective, since the tardiness of high-utilization tasks with worst-case task parameters is either 0 or unbounded by the design of the scheduler. However, other scheduling policies, such as global EDF could be explored.

References

[1] J. Kim, H. Kim, K. Lakashmanan, and R. Rajkumar, "Parallel scheduling for cyber-physical systems: Analysis and case study on a self-driving car," in *ICCPS*, 2013.

[2] J. Li, J.-J. Chen, K. Agrawal, C.Lu, C. Gill, and A. Saifullah, "Analysis of federated and global scheduling for parallel real-time tasks," in *ECRTS*, 2014.

[3] R. I. Davis and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," *ACM Comp. Surv.*, vol. 43, pp. 35:1–44, 2011.

[4] M. Bertogna and S. Baruah, "Tests for global edf schedulability analysis," *J. Syst. Archit.*, vol. 57, no. 5, pp. 487–497, 2011.

[5] B. Andersson and J. Jonsson, "The utilization bounds of partitioned and pfair static-priority scheduling on multiprocessors are 50%," in *ECRTS*, 2003.

[6] J. M. López, J. L. Díaz, and D. F. García, "Utilization bounds for edf scheduling on real-time multiprocessor systems," *Real-Time Syst.*, vol. 28, no. 1, pp. 39–68, 2004.

[7] S. Baruah, V. Bonifaci, A. Marchetti-Spaccamela, and S. Stiller, "Improved multiprocessor global schedulability analysis," *Real-Time Syst.*, vol. 46, no. 1, pp. 3–24, 2010.

[8] V. Bonifaci, A. Marchetti-Spaccamela, and S. Stiller, "A constant-approximate feasibility test for multiprocessor real-time scheduling," *Algorithmica*, vol. 62, no. 3-4, pp. 1034–1049, 2012.

[9] W. Y. Lee and H. Lee, "Optimal scheduling for real-time parallel tasks," *IEICE Trans. Inf. Syst.*, vol. E89-D, no. 6, pp. 1962–1966, 2006.

[10] S. Collette, L. Cucu, and J. Goossens, "Integrating job parallelism in real-time scheduling theory," *Inf. Process. Lett.*, vol. 106, no. 5, 2008.

[11] G. Manimaran, C. S. R. Murthy, and K. Ramamritham, "A new approach for scheduling of parallelizable tasks in real-time multiprocessor systems," *Real-Time Syst.*, vol. 15, no. 1, pp. 39–60, 1998.

[12] S. Kato and Y. Ishikawa, "Gang EDF scheduling of parallel task systems," in *RTSS*, 2009.

[13] K. Lakshmanan, S. Kato, and R. R. Rajkumar, "Scheduling parallel real-time tasks on multi-core processors," in *RTSS*, 2010.

[14] A. Saifullah, K. Agrawal, C. Lu, and C. Gill, "Multi-core real-time scheduling for generalized parallel task models," in *RTSS*, 2011.

[15] G. Nelissen, V. Berten, J. Goossens, and D. Milojevic, "Techniques optimizing the number of processors to schedule multi-threaded tasks," in *ECRTS*, 2012.

[16] H. S. Chwa, J. Lee, K.-M. Phan, A. Easwaran, and I. Shin, "Global edf schedulability analysis for synchronous parallel tasks on multicore platforms," in *ECRTS*, 2013.

[17] B. Andersson and D. de Niz, "Analyzing global-edf for multiprocessor scheduling of parallel tasks," *Principles of Distributed Systems*, 2012.

[18] C. Liu and J. Anderson, "Supporting soft real-time parallel applications on multicore processors," in *RTCSA*, 2012.

[19] L. Nogueira and L. M. Pinho, "Server-based scheduling of parallel real-time tasks," in *International Conference on Embedded Software*, 2012.

[20] S. Baruah, V. Bonifaciy, A. Marchetti-Spaccamelaz, L. Stougiex, and A. Wiese, "A generalized parallel task model for recurrent real-time processes," in *RTSS*, 2012.

[21] V. Bonifaci, A. Marchetti-Spaccamela, S. Stiller, and A. Wiese, "Feasibility analysis in the sporadic dag task model," in *ECRTS*, 2013.

[22] J. Li, K. Agrawal, C.Lu, and C. Gill, "Analysis of global edf for parallel tasks," in *ECRTS*, 2013.

[23] U. C. Devi, *Soft real-time scheduling on multiprocessors*. PhD thesis, University of North Carolina, 2006.

[24] A. Srinivasan and J. H. Anderson, "Efficient scheduling of soft real-time applications on multiprocessors.," in *ECRTS*, 2003.

[25] U. C. Devi and J. H. Anderson, "Tardiness bounds under global edf scheduling on a multiprocessor," *Real-Time Systems*, vol. 38, no. 2, pp. 133–189, 2008.

[26] J. Erickson, U. Devi, and S. Baruah, "Improved tardiness bounds for global edf," in *ECRTS*, 2010.

[27] H. Leontyev and J. H. Anderson, "Generalized tardiness bounds for global multiprocessor scheduling," *Real-Time Systems*, vol. 44, no. 1-3, pp. 26–71, 2010.

[28] J. P. Erickson and J. H. Anderson, "Fair lateness scheduling: Reducing maximum lateness in g-edf-like scheduling," in *ECRTS*, 2012.

[29] L. Palopoli, D. Fontanelli, N. Manica, and L. Abeni, "An analytical bound for probabilistic deadlines," in *ECRTS*, 2012.

[30] A. F. Mills and J. H. Anderson, "A stochastic framework for multiprocessor soft real-time scheduling," in *RTAS*, 2010.

[31] J. M. López, J. L. Díaz, J. Entrialgo, and D. García, "Stochastic analysis of real-time systems under preemptive priority-driven scheduling," *Real-Time Systems*, vol. 40, no. 2, pp. 180–207, 2008.

[32] J. L. Díaz, D. F. García, K. Kim, C.-G. Lee, L. Lo Bello, J. M. López, S. L. Min, and O. Mirabella, "Stochastic analysis of periodic real-time systems," in *RTSS*, 2002.

[33] "Intel CilkPlus." http://software.intel.com/en-us/articles/intel-cilk-plus.

[34] R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C. E. Leiserson, K. H. Randall, and Y. Zhou, "Cilk: An efficient multithreaded runtime system," *Journal of parallel and distributed computing*, pp. 207–216, July 1995.

[35] "OpenMP Application Program Interface v3.1," July 2011. http://www.openmp.org/mp-documents/OpenMP3.1.pdf.

[36] J. Reinders, *Intel threading building blocks: outfitting C++ for multi-core processor parallelism*. O'Reilly Media, 2010.

[37] K. Agrawal, C. E. Leiserson, Y. He, and W. J. Hsu, "Adaptive work-stealing with parallelism feedback," *ACM Trans. Comput. Syst.*, vol. 26, pp. 112–120, September 2008.

[38] J. Kingman, "Inequalities in the theory of queues," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 102–110, 1970.