ERRATA: JULY 12TH, 2023

We found implementation errors for the schedulability tests RTA-EDF and UGB. Therefore, we re-simulate the task sets and revise the evaluation section as follows.

VI. EVALUATION

In this section, we compare our RTA framework designed for global gang scheduling, with existing schedulability tests for both global and non-global gang scheduling.

A. Evaluation Setting

The randomly generated task sets are based on [5]. For each number of processors m (i.e., 8, 16, 32, 64, 128 and 256), we consider four parameters: (S1) the type of the task set, i.e., implicit-deadline ($D_i = T_i$) and constrained-deadline ($D_i \leq T_i$), (S2) the distribution of task utilization $u_i \stackrel{\text{def.}}{=} C_i/T_i$, i.e., the binomial distribution with $p = 0.1, 0.3, 0.5, 0.7$ and 0.9,[3] (S3) the range of task parallelism m_i, i.e., $[1, \frac{1}{2}m]$ (S4) the range of task set utilization $U \stackrel{\text{def.}}{=} \sum_{\tau_i \in \tau} \frac{u_i \cdot m_i}{m}$, i.e., $[0.0, 0.1)$, $[0.1, 0.2)$, ... ,$[0.9, 1.0)$. For each task, the period T_i is uniformly selected in $[10ms, 1000ms]$; C_i is set to $u_i \cdot T_i$, where u_i is generated by S2; for implicit- and constrained-deadline tasks, D_i is set to T_i and uniformly selected in $[C_i, T_i]$, respectively; and m_i is uniformly distributed in the range assigned by S3. For every combination of S1, S2, S3 and S4, we generate 1000 task sets, yielding $2 \cdot 5 \cdot 2 \cdot 10 \cdot 1000 = 200,000$ task sets in total for each m.

Using the generated sets, we compare our schedulability tests designed for preemptive *global* gang scheduling, with all existing schedulability tests for preemptive *global/non-global* gang scheduling subject to our task model (explained in Section II), as follows.

- WaPe: global scheduling for FP in [10]
- DoLi: global scheduling for EDF in [5], [15]
- UGB: non-global scheduling (i.e., a generalization of partitioned scheduling) for FP in [11]
- RTA-FP and RTA-EDF: Theorem 2 for FP and EDF
- RTA^1-FP and RTA^1-EDF: Theorem 4 for FP and EDF
- RTA^2-FP and RTA^2-EDF: Theorem 6 for FP and EDF
- RTA*-FP and RTA*-EDF: Theorem 7 for FP and EDF

For fair comparison for the tests that employ FP as a prioritization policy, we apply DM (Deadline Monotonic) [24] to all the tests.

We count the number of task sets deemed schedulable by each of the above schedulability tests, and show the ratio of those task sets. We observe that the trend for the relative ratio among individual tests does not much vary with m. Therefore, we explain the representative results for $m = 64$ in the next subsections. For $m = 64$, we present the overall results for all generated task sets without any figure, *and* some interesting results with Fig. 3 for a subset of generated task sets subject to a pair of S1 and S3, denoted by (I/D, L/H), where I and D imply implicit-deadline and constrained-deadline task sets in S1, respectively, and L and H imply $m_i \in [1, \frac{1}{2}m]$ (i.e., low m_i) and $m_i \in [1, m)$ (i.e., high m_i), respectively. In each of Fig. 3, the X-axis represents the task set utilization U (i.e., S4), while the Y-axis represents the ratio of task sets deemed schedulable by each schedulability tests. Therefore, each point in Fig. 3 targets task sets subject to a given combination of S1, S3 and S4 while the target task sets include all parameters of S2.

[3]For given p, task utilization is uniformly distributed in $[0.5, 1.0]$ and $[0.0, 0.5]$ with probability of p and $1.0-p$, respectively. Therefore, the average number of tasks in each task set decreases as p increases.

B. Comparison of Global Gang Scheduling

As our RTA framework targets global gang scheduling, we now compare our RTA framework with a given prioritization policy, to an existing schedulability test for global gang scheduling with the same policy, i.e., RTA*-FP (and RTA-FP) versus WaPe, and RTA*-EDF (and RTA-EDF) versus DoLi.

For global gang FP scheduling, RTA*-FP and RTA-FP outperform WaPe under every combination of S1 and S3, and they respectively achieve 34.6% and 23.1% overall improvement over WaPe. This is because, while RTA-FP and WaPe share a similar schedulability analysis structure, RTA-FP tightly calculates a response time using the notion of k-interference slots/processors (and RTA*-FP more tightly does). The most favorable and unfavorable settings for RTA*-FP against WaPe, are (C, L) and (I, H), respectively, shown in Figs. 3(a) and (b); under the settings, RTA*-FP respectively finds 67.4% and 18.0% more schedulable task sets, compared to WaPe.

When it comes to global gang EDF scheduling, RTA*-EDF and RTA-EDF also outperforms DoLi in that RTA*-EDF and RTA-EDF find 41.0% and 34.6% more schedulable task sets than DoLi, respectively. Since DoLi uses a notion of the maximum idle parallelism when a task cannot be executed, it is more effective for task sets with low m_i. As a result, the performance of RTA*-EDF against DoLi varies with the setting for S3. For example, as shown in Figs. 3(c) and (d), RTA*-EDF finds 119.6% and 80.9% *more* schedulable task sets compared to DoLi, respectively under (C, L) and (C, H).

In summary, RTA*-FP and RTA*-EDF significantly outperforms the existing schedulability tests that target global gang FP and EDF scheduling, respectively.

C. Comparison of Any Gang Scheduling

We now present the performance of our schedulability tests with all other existing ones, regardless of prioritization policies (i.e., EDF or FP) and scheduling categories (i.e., global or non-global). Overall, the schedulability ratio of RTA*-EDF, DoLi, WaPe, and UGB, normalized by that of RTA*-FP, is 86.4%, 61.3%, 74.3% and 114.5%, respectively. Between the two highest schedulability-performance tests RTA*-FP and UGB, we observe that UGB finds 19.1%, 7.2%, 25.1% and 6.7% more schedulable task sets than RTA*-FP under (I, L), (I, H), (C, L) and (C, H), respectively (the best and the worst performance of RTA*-FP against UGB is shown in Fig. 3(e) and (f), respectively). However, among task sets schedulable by RTA*-FP, 1.1%, 0.5%, 1.6% and 0.6% task sets are not deemed schedulable by UGB, respectively, under (I, L), (I,

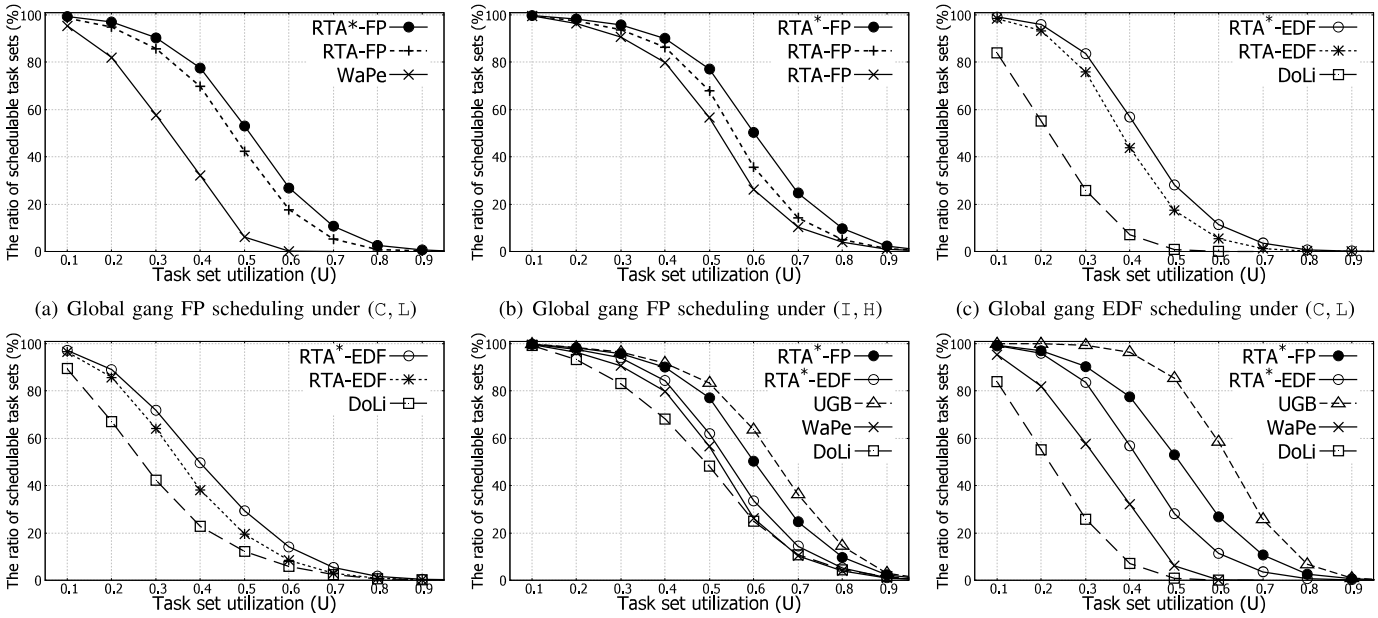

(a) Global gang FP scheduling under (C, L)

(b) Global gang FP scheduling under (I, H)

(c) Global gang EDF scheduling under (C, L)

(d) Global gang EDF scheduling under (C, H)

(e) Any gang scheduling under (I, H)

(f) Any gang scheduling under (C, L)

Fig. 3. Schedulability performance comparison of our schedulability tests with existing ones

H), (C, L) and (C, H). Considering there has been discussion of superiority between global and partitioned scheduling for the sequential task model, e.g., [25], [26], [27], it is interesting to observe the existence of task sets, which are schedulable by RTA*-FP but not schedulable by UGB, both of which share the same prioritization policy.

We also observe that the schedulability performance of RTA*-EDF is less than that of RTA*-FP. This accords with the corresponding results for the RTA framework for the sequential task model [23].

D. Comparison of Our RTA Frameworks

Finally, we present how our novel techniques in Sections IV and V and its composition (corresponding to $\texttt{RTA}^1$, $\texttt{RTA}^2$, $\texttt{RTA}^*$) improve our basic response time analysis RTA. Compared to RTA for EDF, $\texttt{RTA}^1$, $\texttt{RTA}^2$ and $\texttt{RTA}^*$ for EDF yield 4.6%, 8.7% and 11.2% overall schedulability improvement, respectively. A similar trend is observed for FP, yielding 4.3%, 7.3% and 9.3% overall schedulability improvement, respectively. In particular, if we focus on task sets with (C, H), the improvement for EDF and FP is increased to 6.7%, 8.8% and 13.3%, and 6.6%, 7.4% and 11.2%, respectively. The results demonstrate the effectiveness of the proposed two techniques and its composition in reducing pessimistic interference calculation.