ILP-Based Approaches to Partitioning Recurrent Workloads Upon Heterogeneous Multiprocessors

ILP-based approaches to partitioning recurrent workloads upon heterogeneous multiprocessors

Renato Bruni, Alberto Marchetti-Spaccamela, Sanjoy K Baruah, Vincenzo Bonifaci

► To cite this version:

Renato Bruni, Alberto Marchetti-Spaccamela, Sanjoy K Baruah, Vincenzo Bonifaci. ILP-Based Approaches to Partitioning Recurrent Workloads Upon Heterogeneous Multiprocessors ILP-based approaches to partitioning recurrent workloads upon heterogeneous multiprocessors. 28th Euromicro Conference on Real-Time Systems, ECRTS 2016, Jul 2016, Toulouse, France. pp.215-225, 10.1109/ECRTS.2016.10 . hal-01397810

HAL Id: hal-01397810

<https://hal.inria.fr/hal-01397810>

Submitted on 16 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/307572271>

ILP-Based Approaches to Partitioning Recurrent Workloads Upon Heterogeneous Multiprocessors

Conference Paper · July 2016

DOI: 10.1109/ECRTS.2016.10

CITATIONS

0

READS

7

4 authors, including:

[Renato Bruni](#)

Sapienza University of Rome

57 PUBLICATIONS 301 CITATIONS

[SEE PROFILE](#)

[Alberto Marchetti-Spaccamela](#)

Sapienza University of Rome

173 PUBLICATIONS 1,973 CITATIONS

[SEE PROFILE](#)

ILP-based approaches to partitioning recurrent workloads upon heterogeneous multiprocessors

Sanjoy K. Baruah*, Vincenzo Bonifaci[†], Renato Bruni[‡], Alberto Marchetti-Spaccamela[‡]

*University of North Carolina at Chapel Hill, E-mail: baruah@cs.unc.edu

[†]Istituto di Analisi dei Sistemi ed Informatica, CNR, Roma, Italy, E-mail: vincenzo.bonifaci@iasi.cnr.it

[‡]Università di Roma “La Sapienza”, Roma, Italy, E-mail: {bruni,alberto}@diag.uniroma1.it

Abstract—The problem of partitioning systems of independent constrained-deadline sporadic tasks upon heterogeneous multiprocessor platforms is considered. Several different integer linear program (ILP) formulations of this problem, offering different tradeoffs between effectiveness (as quantified by speedup bound) and running time efficiency, are presented.

I. INTRODUCTION

Heterogeneous multicore CPUs – CPUs in which the processing elements differ from one another with respect to functionality or processing speed – are currently widely available and increasingly becoming the common case. The presence of such heterogeneity requires choices to be made when mapping software components onto processing elements. The need to resolve such choices adds considerable complexity to resource allocation, and inhibits the adoption of such platforms by the embedded computing industry despite significant potential benefits in terms of balancing performance and energy.

We consider here real-time systems that are modeled as collections of independent sporadic tasks (the model is described in detail in Section II). We seek to devise algorithms for implementing such systems upon heterogeneous multiprocessor platforms under the partitioned paradigm. To our knowledge, this topic has not been studied much previously:

- On the one hand, most prior real-time scheduling research that considers heterogeneous platforms (see, e.g., [10], [22], [21], [23], [24]; [20] has a nice survey) has restricted attention to *implicit-deadline* sporadic tasks.
- On the other hand, prior research that does address the partitioned scheduling of task systems represented using models that are more general than the implicit-deadline model considers *identical* multiprocessor platforms only (see, e.g., [6], [9]).

In this paper, we initiate a methodical study of the problem of partitioning, upon heterogeneous multiprocessor platforms, task systems that are represented using the more general *constrained-deadline* sporadic task model¹. We assume that once the partitioning has been performed and tasks assigned to the processors, run-time scheduling is done on each processor using the earliest deadline first (EDF) scheduling algorithm, which is known to be optimal for this purpose [18], [11].

¹Although we expect that most of our results will also extend to the arbitrary-deadline sporadic task model, for ease of presentation we do not explore this issue any further in this paper, but leave it for future work.

Our approach. We will derive various approaches to task partitioning. These algorithms share the commonality that they are all based upon formulating the task partitioning problem as an integer linear program (ILP). For implicit-deadline task systems, this is not particularly difficult to do; indeed most of the research on partitioning implicit-deadline sporadic task systems (including the works [22], [21], [23], [24] cited above) has been based upon first formulating such an ILP, and then seeking polynomial-time algorithms for obtaining approximate solutions to these ILPs (solving an ILP is known to be NP-hard [15], and hence unlikely to be solvable exactly in polynomial time).

Despite this inherent intractability of solving ILPs, however, the optimization community has recently been devoting immense effort to devise extremely efficient implementations of ILP solvers, and highly-optimized libraries with such efficient implementations are widely available today. Modern ILP solvers, particularly when running upon powerful computing clusters, are often capable of solving ILPs with tens of thousands of variables and constraints. We therefore believe that it is reasonable to attempt to solve ILPs exactly rather than only approximately, and seek to obtain ILP formulations that we will seek to solve exactly to solve the partitioning problem for constrained-deadline sporadic task systems. Since the running time of ILP solvers tends to increase with the number of variables and constraints in the ILP to be solved, we seek to develop ILPs for task partitioning in which the number of variables and constraints are restricted to be low-order polynomials of the representation of the task system. While the number of constraints may not always be a good indicator of the complexity of an ILP formulation, we use it as a first approximation: indeed, the best known complexity bounds for solving ILPs do increase with the number of linear constraints [12, Theorem 5.3]. Possibly more refined metrics, such as the *constrained induced-width* [13] or the *constraint density* [2], have also been suggested in other settings, but not in the context of the problem of partitioning tasks onto heterogeneous processors – not even implicit-deadline tasks [14]; these refinements fall outside the scope of this work.

Our results. In partitioning implicit-deadline sporadic task systems, an ILP represents an exact solution to the partitioning problem — solving an ILP exactly therefore constitutes an optimal algorithm for performing such partitioning. For

	Number of constraints	Speedup factor	Comments
1.	$n + m + m \log_2 d_{\max}$	4	In Section III-B
2.	$n + m + m \log_\rho d_{\max}$	2ρ	In Section III-C. ρ is a constant > 1 . A generalization of row 1 (which is obtained if $\rho \leftarrow 2$)
3.	$n + m + nmk$	$1 + \frac{1}{k}$	In Section IV. k is any integer ≥ 1 .
4.	$n + m + m \log_\rho d_{\max}$	$2 + \rho + \frac{\rho^2}{\rho-1}$	In Section V. A “poorer” version of row 2 — same number of constraints, larger speedup factor. But more amenable to polynomial-time approximation — see Section V for details

Table I
SUMMARY OF RESULTS

partitioning constrained-deadline systems, however, we do not know how to obtain such an exact ILP representation of this problem with only polynomially many constraints — this difficulty was previously identified even for partitioning upon identical multiprocessors in [4]. Instead, our goal here is to obtain polynomial-sized ILP representations of the problem of partitioning constrained-deadline sporadic task systems upon heterogeneous multiprocessor platforms with the property that exact solutions of the ILP constitute approximate solutions to the partitioning problem. Our metric of effectiveness of such an approximate solution is the *speedup factor* — an ILP formulation has speedup factor f , $f \geq 1$, if any constrained-deadline sporadic task system that can be partitioned upon a particular heterogeneous platform by a hypothetical optimal algorithm can be partitioned using this ILP formulation upon a platform in which each processor is at least f times as fast.

We have derived several different ILP representations for the problem of partitioning constrained-deadline sporadic task systems upon heterogeneous multiprocessor platforms, all of which have number of variables and constraints polynomial in the representation of the task system. All these ILP formulations have nm integer variables, where n is the number of tasks and m the number of processors, but specify different numbers of constraints and offer different speedup guarantees — they are summarized in Table I.

II. SYSTEM MODEL, BACKGROUND, AND NOTATION

We seek to partition a sporadic task system τ comprising the n independent sporadic tasks $\tau_1, \tau_2, \dots, \tau_n$ upon an unrelated multiprocessor platform π comprising the m processors $\pi_1, \pi_2, \dots, \pi_m$. The i 'th sporadic task τ_i is characterized by a period p_i and a relative deadline d_i , and m worst-case execution time (WCET) parameters $c_{i,1}, c_{i,2}, \dots, c_{i,m}$, with $c_{i,j}$ denoting the WCET of τ_i if it executes upon processor π_j . In this paper, we restrict attention to task systems in which $d_i \leq p_i$ for each task $\tau_i \in \tau$ — such sporadic task systems are called *constrained-deadline* sporadic task systems. During run-time, τ_i releases a sequence of jobs at time-instants that are not known beforehand, but with the constraint that successive jobs are released at least p_i time units apart. Each job of τ_i has a deadline d_i time units after its release time; the amount of execution needed by this job depends upon the identity of

the processor on which it executes. More specifically, since we are studying partitioned scheduling in this paper, given task system τ and multiprocessor platform π , our objective is to obtain a partitioning of the tasks upon the processors. Let $f(i) \in \{1, 2, \dots, m\}$ denote the index of the processor to which each τ_i is assigned under such a partitioning; each job of τ_i needs to execute for up to $c_{i,f(i)}$ time units upon processor $\pi_{f(i)}$.

Since the preemptive Earliest Deadline First scheduling algorithm (EDF) is known to be optimal for scheduling a single preemptive processor, we will use EDF as the scheduling algorithm upon each individual processor during run-time. The demand bound function (dbf) [5] of a sporadic task is widely used to quantify the computational demand of such a task, where the $\text{dbf}(\tau_i, t)$ of sporadic task τ_i with period p_i , relative deadline d_i , and WCET c_i for an interval of duration t is defined as follows

$$\text{dbf}(\tau_i, t) := \left\lfloor \frac{t + p_i - d_i}{p_i} \right\rfloor c_i$$

It is known that a collection of sporadic tasks can be scheduled to always meet all deadlines upon a preemptive uniprocessor by EDF if and only if for all $t \geq 0$, the sum of the dbf's of all the tasks in the collection for an interval of duration t does not exceed t .

Some additional notation:

- Let $N := \{1, 2, \dots, n\}$ denote the task index set.
- Let $M := \{1, 2, \dots, m\}$ denote the processor index set.
- Let $u_{i,j} := c_{i,j}/p_i$ denote the *utilization* of task τ_i on processor π_j .
- Let $d_{\max} := \max_{1 \leq i \leq n} \{d_i\}$ denote the largest relative deadline parameter of any task in τ .

As above, let $f : N \rightarrow M$ denote a partitioning of the task system τ upon multiprocessor platform π . We use the notation $\text{dbf}_{f,j}(t)$ to denote the sum of the dbf's of all the tasks in τ that have been assigned to processor π_j under the partitioning f , for interval duration t :

$$\text{dbf}_{f,j}(t) := \sum_{i \in N: f(i)=j} \left(\left\lfloor \frac{t + p_i - d_i}{p_i} \right\rfloor c_{i,j} \right)$$

III. A SIMPLE ILP MODEL FOR TASK PARTITIONING

Marchetti-Spaccamela et al. have previously [19] derived a polynomial-time algorithm for assigning sporadic tasks to heterogeneous processors, and shown that this algorithm has a speedup bound of $(8+2\sqrt{6})$ or ≈ 12.9 . An intermediate step in [19] is the derivation of a 0/1 ILP representation of the partitioning problem with nm variables and $(n+m+m \log d_{\max})$ linear constraints, and a proof that this ILP representation has a speedup factor of 6. In this section, we present two results:

- 1) We derive, in Theorem III.3 below, an improved ILP with the same number of variables and constraints and show (Corollary III.4) that it has a superior (i.e., smaller) speedup factor of 4.
- 2) In Theorem III.5, we generalize the derivation of this improved ILP in the following manner. For any constant $\rho > 1$, we can derive an ILP with nm variables and $(n+m+m \log_{\rho} d_{\max})$ linear constraints and speedup bound 2ρ ; by choosing ρ to be smaller than two, a smaller speedup bound than 4 is thus obtained at the cost of needing to solve an ILP with a larger number of constraints.

A. A review of some results from [19]

First a preliminary definition. Let D denote the set of values

$$\{0, 1, 2, 4, \dots, 2^{\lceil \log_2 d_{\max} \rceil}\}$$

We call D the *deadline checkpoint set* for task system τ .

The starting point of the reasoning in [19] is the following lemma.

Lemma III.1 (from [19]). *Let $f : N \rightarrow M$ denote an assignment of the tasks in task system τ to the processors of unrelated multiprocessor platform π such that, for each $j \in M$*

$$\sum_{i: f(i)=j} u_{i,j} \leq \beta$$

and for each $j \in M$ and $k, 1 \leq k \leq \lceil \log_2 d_{\max} \rceil$,

$$\left(\sum_{i: (f(i)=j) \wedge (2^{k-1} < d_i \leq 2^k)} c_{i,j} \right) \leq \beta \cdot 2^k.$$

Then for each $j \in M$, $\text{dbf}_{f,j}(t) \leq 6\beta t$ for all $t \geq 0$.

Let us try and understand what this lemma means. The first set of inequalities requires that the cumulative utilization assigned to each processor not exceed β ; the second, that the sum of the WCETs of all tasks assigned to a processor that have relative deadlines between two successive powers of two not exceed β times the larger power of two. (For example, considering $k \leftarrow 6$, the second constraint mandates that the sum of the WCETs of all tasks with relative deadline in the range $(32, 64]$ not exceed 64β .) The lemma asserts that if these conditions are satisfied by an assignment, then this assignment constitutes a valid partitioning upon processors of speed $6 \times \beta$.

Now, suppose that τ is feasible upon π under partitioned scheduling, i.e., there is an assignment $f : N \rightarrow M$ of τ upon

π such that all jobs of all tasks will always complete by their deadlines if tasks are assigned according to this partitioning, and each processor scheduled during run-time by EDF. For this assignment f , it is evident that the utilization constraints of Lemma III.1 are satisfied for $\beta = 1$. The second set of constraints in Lemma III.1 will also be satisfied for $\beta = 1$, by the following reasoning:

- Since the partitioning f is feasible, the sum of the dbf's of all the tasks assigned to the j th processor for interval duration 2^k is no more than 2^k .
- Each task with relative deadline in $(2^{k-1}, 2^k]$ must have $\text{dbf}(\tau_i, 2^k) \geq c_{i,j}$.
- Hence the sum of the $c_{i,j}$'s for all tasks τ_i that have relative deadline in $(2^{k-1}, 2^k]$ and are assigned to the j th processor must be $\leq 2^k$.

Hence if τ is feasible upon π , there exists an $f : N \rightarrow M$ for which the conditions of Lemma III.1 are satisfied with $\beta = 1$; therefore $\text{dbf}_{f,j}(t) \leq 6t$ for all $t \geq 0$ for each j . Thus [19] derives the following speedup result:

Corollary III.2. *Let $f : N \rightarrow M$ denote an assignment of tasks to processors satisfying the conditions of Lemma III.1. Then f constitutes a feasible partitioning of τ upon π under a speedup factor of 6β . In particular, if the conditions of the lemma are satisfied with $\beta \leq 1/6$, then assignment f constitutes a feasible partitioning of τ upon the given platform π .*

In summary, the conditions of Lemma III.1 are necessary when $\beta = 1$ and sufficient when $\beta = 1/6$, i.e., they yield a speedup factor of 6 for partitioning constrained-deadline sporadic tasks upon heterogeneous multiprocessor platforms.

B. An ILP with speedup factor 4

We now prove an improved version of Lemma III.1 that yields a superior speedup bound (of four, rather than six). The conditions specified in our theorem below differs from those in Lemma III.1 only in that the summation of $c_{i,j}$'s in the second inequality is over all tasks with relative deadline $\leq 2^k$ (rather than only those with relative deadline $> 2^{k-1}$ and $\leq 2^k$).

Theorem III.3. *Let $f : N \rightarrow M$ denote an assignment of the tasks in task system τ to the processors of unrelated multiprocessor platform π such that, for each $j \in M$*

$$\sum_{i: f(i)=j} u_{i,j} \leq \beta \tag{1}$$

for each $k, 1 \leq k \leq \lceil \log_2 d_{\max} \rceil$,

$$\left(\sum_{i: (f(i)=j) \wedge (d_i \leq 2^k)} c_{i,j} \right) \leq \beta \cdot 2^k. \tag{2}$$

Then for each $j, 1 \leq j \leq m$, $\text{dbf}_{f,j}(t) \leq 4\beta t$ for all $t \geq 0$.

Proof. Consider any $t \geq 0$ and let $s := 2^k$ denote the smallest integer power of 2 that is not smaller than t (i.e. $s = 2^k \geq t > s/2$).

Consider an assignment $f : N \rightarrow M$ of the tasks to the processors. For any $j \in M$, we have

$$\begin{aligned}
\text{dbf}_{f,j}(t) &\leq \text{dbf}_{f,j}(s) \\
&= \sum_{i:f(i)=j \wedge d_i \leq s} \left\lfloor \frac{s + p_i - d_i}{p_i} \right\rfloor c_{ij} \\
&\leq \sum_{i:f(i)=j \wedge d_i \leq s} \left(s \frac{c_{ij}}{p_i} + c_{ij} \right) \\
&\leq s \sum_{i:f(i)=j} \frac{c_{ij}}{p_j} + \sum_{i:f(i)=j \wedge d_i \leq s} c_{ij} \\
&= s \sum_{i:f(i)=j} u_{ij} + \beta \cdot 2^k \\
&\leq \beta s + \beta s \\
&\leq 4\beta t
\end{aligned}$$

and the theorem is proved. $\square$

This implies that any f satisfying Inequalities 1 and 2 of Theorem III.3 constitutes a feasible partitioning of the tasks in τ upon the set of heterogeneous processors π , when the processors receive a speedup factor of 4β :

Corollary III.4. *Let $f : N \rightarrow M$ denote an assignment of tasks to processors satisfying Inequalities 1 and 2 of Theorem III.3. Then f constitutes a feasible partitioning of τ upon π under a speedup factor of 4β . In particular, if Inequalities 1 and 2 are satisfied with $\beta \leq 1/4$, then assignment f constitutes a feasible partitioning of τ upon the given platform π . $\square$*

We now apply the result of Theorem III.3 above to construct a 0/1 integer linear program (ILP) for specifying a feasible partitioning of sporadic task system τ upon the platform π . That is, we will construct a 0/1 ILP, a solution to which will yield a partitioning $f : N \rightarrow M$ that satisfies Inequalities 1 and 2. This ILP is constructed as follows:

- For each $i \in N, j \in M$, we have a 0/1 integer variable (i.e., an integer variable that takes on either the value zero or the value one) $x_{i,j}$, denoting whether τ_i is to be assigned to processor π_j .
- We specify that each task gets assigned to exactly one processor; this is done by the following n constraints:

$$\forall i \in N \quad \left(\sum_{j \in M} x_{i,j} \right) = 1 \quad (3)$$

- We next specify that Inequality 1 of Theorem III.3 should be satisfied; this is achieved by the following m constraints:

$$\forall j \in M \quad \left(\sum_{i \in N} x_{i,j} u_{ij} \right) \leq \beta \quad (4)$$

- Finally, we specify Inequality 2 of Theorem III.3 by the following $(\log d_{\max} \times m)$ constraints:

$$\forall k \in D, \forall j \in M \quad \left(\sum_{(i \in N) \wedge (d_i \leq 2^k)} c_{i,j} x_{i,j} \right) \leq \beta \cdot 2^k \quad (5)$$

By Theorem III.3, any assignment of integer values to the $\{x_{i,j}\}$ variables satisfying the Constraints 3–5 above bears witness to the schedulability of τ , with a speedup of 4β , upon the platform π . Moreover, for a τ that is feasible upon π the model always admits a solution with $\beta = 1$, since all inequalities are clearly valid; thus, this ILP guarantees a speedup factor of at most 4. When the ILP model is feasible with $\beta \leq 1/4$, Theorem III.3 guarantees schedulability on the original platform; hence, a reasonable objective function for the ILP with Constraints 3–5 would be to **minimize** β .

C. A generalization

Above, we derived an ILP model for the problem of partitioned scheduling of constrained-deadline sporadic task systems upon unrelated multiprocessors, such that any solution to the ILP immediately yields a partitioning algorithm at a speedup bound of 4. We also saw that this ILP has $(n + m + m \times \lceil \log_2 d_{\max} \rceil)$ constraints; we now briefly describe how to reduce the speedup bound by increasing the number of constraints.

Recall that we had defined the *deadline checkpoint set* D as powers of two: $D = \{0, 2^0, 2^1, 2^2, \dots, 2^{\lceil \log_2 d_{\max} \rceil}\}$. For any given constant $\rho > 1$, we could instead have chosen to define it as

$$D_\rho = \{0, 1, \rho, \rho^2, \rho^3, \dots, \rho^{\lceil \log_\rho d_{\max} \rceil}\},$$

The following generalization of Theorem III.3 is easily proved via a proof analogous to the proof of Theorem III.3:

Theorem III.5. *Let $f : N \rightarrow M$ denote an assignment of the tasks in task system τ to the processors of unrelated multiprocessor platform π such that, for each $j \in M$*

$$\sum_{i:f(i)=j} u_{i,j} \leq \beta,$$

and for each $j \in M$ and $k, 1 \leq k \leq \lceil \log_\rho d_{\max} \rceil$,

$$\left(\sum_{i \in N: (f(i)=j) \wedge (d_i \leq \rho^k)} c_{i,j} \right) \leq \beta \cdot \rho^k$$

Then for each $j \in M$, $\text{dbf}_{f,j}(t) \leq 2\rho\beta t$ for all $t \geq 0$. $\square$

It follows, from arguments virtually identical to those of Corollary III.4, that the speedup bound for the ILP constructed based on Theorem III.5 above is 2ρ ; hence by choosing ρ to be smaller than 2 a speedup bound smaller than 4 is obtained. The tradeoff is that the number of constraints increases to $(n + m + m \times \lceil \log_\rho d_{\max} \rceil)$; this is $> \log_2 d_{\max}$ for $\rho < 2$, becoming larger as $\rho \rightarrow 1$.

Theorem III.5 thus suggests one approach for obtaining speedup bounds arbitrarily close to 2, by simply selecting a denser deadline checkpoint set. In Section IV below, we explore another approach, that allows for speedup bounds arbitrarily close to one (once again at the cost of having additional constraints).

IV. A STRENGTHENED ILP FORMULATION

We now explore a different idea that also trades off an increase in the number of constraints in the ILP for a superior speedup bound. Specifically, for any positive integer constant k we will derive an ILP model with $(n+m+mnk)$ constraints, finding a feasible solution to which corresponds a partitioning at a speedup bound of $(1 + \frac{1}{k})$.

Approximation schemes have been defined for computing the value of dbf to any desired degree of accuracy (see, e.g. [3]). Equation 6 below gives such an approximation scheme; for any fixed positive integer value of k , $\text{dbf}^{(k)}(\tau_i, t)$ defines an approximation of $\text{dbf}(\tau_i, t)$ that is exact for the first k “steps” of $\text{dbf}(\tau_i, t)$, and an upper bound for larger values of t :

$$\text{dbf}^{(k)}(\tau_i, t) = \begin{cases} c_{i,j} \times \left\lfloor \frac{t+p_i-d_i}{p_i} \right\rfloor & \text{if } t \leq (k-1)p_i + d_i \\ c_i + (t-d_i)u_i & \text{otherwise} \end{cases} \quad (6)$$

The following lemma (see, e.g., [4]) provides a quantitative bound on the degree by which $\text{dbf}^{(k)}$ may deviate from dbf:

Lemma IV.1. *For all $t \geq 0$*

$$\text{dbf}(\tau_i, t) \leq \text{dbf}^{(k)}(\tau_i, t) < (1 + \frac{1}{k}) \text{dbf}(\tau_i, t).$$

That is, $\text{dbf}^{(k)}(\tau_i, t)$ provides an upper bound on $\text{dbf}(\tau_i, t)$ that is no more than a fraction $1/k$ greater than the actual value of $\text{dbf}(\tau_i, t)$.

As previously stated, it is known that a collection of sporadic tasks can be scheduled to always meet all deadlines upon a preemptive uniprocessor by EDF if and only if for all $t \geq 0$, the sum of the dbf’s of all the tasks in the collection over an interval of duration t does not exceed t . For schedulability, it is clearly necessary that the utilizations of all the tasks in the collection sum to no more than 1. Since $\text{dbf}^{(k)}$ is an upper bound on dbf, it follows that a *sufficient* uniprocessor EDF-schedulability test for a collection of tasks is that the sum of the $\text{dbf}^{(k)}$ functions of all the tasks in the collection over an interval of duration t not exceed t . Albers and Slomka showed [3, Lemma 4] that it suffices to validate this fact only for those values of t at which one or more of the $\text{dbf}^{(k)}$ functions has a step discontinuity. More precisely, let

$$S_{i,k} = \{t : t = d_i + hp_i, h = 0, 1, \dots, k\}$$

and $S_k = \bigcup_{\text{all } i} S_{i,k}$

It suffices to test that the sum of the $\text{dbf}^{(k)}$ functions of all the tasks in the collection over an interval of duration t not exceed t , only for values of $t \in S_k$.

We can use this result to define a revised ILP formulation for modeling the partitioned scheduling of sporadic task systems upon heterogeneous multiprocessors. The first part of this revised ILP is identical to the one constructed in Section III:

- As in Section III, we will have a zero-one integer variable $x_{i,j}$, denoting whether τ_i is to be assigned to processor π_j , for each $i \in N, j \in M$.
- Again as in Section III, the following n constraints specify that each task gets assigned to exactly one processor:

$$\forall i \in N \quad \left(\sum_{j \in M} x_{i,j} \right) = 1 \quad (7)$$

- The following m constraints bound the total utilization on each processor:

$$\forall j \in M \quad \left(\sum_{i \in N} x_{i,j} u_{i,j} \right) \leq \beta \quad (8)$$

The final set of constraints replace the Inequalities 5 of the ILP in Section III with constraints based upon the $\text{dbf}^{(k)}$ approximation, that express the requirement that the sum of the $\text{dbf}^{(k)}$ functions of all tasks assigned each processor over an interval duration not exceed the duration. As we had stated above, this condition only needs to be validated for interval durations in S_k ; this motivates the following set of constraints:

$$\forall t \in S_k, \forall j \in M \quad \left(\sum_{i \in N} (x_{i,j} \times \text{dbf}_j^{(k)}(\tau_i, t)) \right) \leq \beta \cdot t \quad (9)$$

(where for each $j \in M$, $\text{dbf}_j^{(k)}(\tau_i, t)$ denotes the function $\text{dbf}^{(k)}(\tau_i, t)$ when the WCET of τ_i is set equal to $c_{i,j}$.) Since each τ_i contributes at most k distinct points to S_k , it follows that $|S_k| \leq nk$; hence there are at most mnk such constraints.

Note that the Inequalities 9 are constructed for specified values of t ; i.e., for each $t \in S_k$. For each such specified t , it is straightforward to observe that the inequality is indeed a linear one, since inspection of Expression 6 reveals that for a given value of t for each τ_i the expression $\text{dbf}_j^{(k)}(\tau_i, t)$ is a *constant*.

Theorem IV.2. *A feasible solution to the ILP on the 0/1 variables $\{x_{i,j}\}, i \in N, j \in M$, with Constraints 7–9 (defined above) yields a feasible partitioning of the tasks in τ to a platform in which each processor in π is speeded up by a multiplicative factor of $(1 + 1/k)\beta$. In particular, if the inequalities are satisfied with $\beta \leq (1 + 1/k)^{-1}$, then a feasible solution to the ILP yields a feasible partitioning on the given platform.*

Proof. It is evident from the result of Albers and Slomka that satisfying Constraints 7–9 is sufficient for feasibility upon a speed- $\beta(1 + 1/k)$ platform. Additionally we conclude from the lower bound in Lemma IV.1 that failure to satisfy these conditions implies infeasibility upon a speed- β platform. $\square$

The quality of the solution that is obtained by solving the ILP specified by Constraints 7–9 depends on the value of k : the larger this value, the better is the quality (i.e., the lower the speedup factor) of the obtained solution. However we observe that the number of constraints increases with k . It follows that large values of k lead to an ILP that is not solvable with state of the art packages. We explore this tradeoff via schedulability experiments In Section VI.

V. AN ILP THAT IS MORE AMENABLE TO APPROXIMATING IN POLYNOMIAL TIME

In the previous sections, we discussed assignment algorithms based on solving reasonably-sized ILPs. However, in some scenarios the solution of an ILP may be a computational bottleneck – solving an ILP is, after all, a prototypical NP-hard problem. Therefore, the design of efficient (polynomial time) assignment and schedulability algorithms retains interest.

A standard technique that has been developed within the optimization community for efficiently obtaining an approximate solution to an ILP is to first consider the linear program (LP) obtained by “relaxing” (i.e., ignoring) the integrality requirement, solve this LP (this can be done in polynomial time), and then “rounding” the solution so obtained to obtain an integral solution as desired. The main challenge in designing the rounding procedure is to ensure that such rounding does not degrade the feasibility or the quality of the solution (i.e., the value of the objective function that was optimized) too much.

Recently, a new approach to rounding, known as *iterative rounding* [17], has been shown to provide improved rounding guarantees. Analogously to prior rounding approaches, the first step requires that an LP-relaxation be solved and a non-integer solution (say, X) be obtained. However, instead of rounding all non integral values of X at the same time, only one variable is rounded; assume, for example, that the value of variable x_1 is set to $\hat{x}_1$. Then the method iterates and solves a new LP-relaxation that is obtained from the original LP by fixing the value of x_1 to $\hat{x}_1$. In this way, a new solution X' is obtained; as in the previous case, the method now rounds one variable of X' ; the method is iterated until all variables satisfy the given integrality constraints.

In this section, we seek to construct an ILP formulation of the problem of partitioning sporadic tasks upon heterogeneous multiprocessors that is more amenable to iterative rounding than the ILP formulations we have seen above. It will turn out (Theorem V.1 below) that this ILP has the same number of variables and constraints, but a poorer (larger) speedup factor than the one described in Section III-C (Theorem III.5). Hence from the perspective of just developing an ILP, this is not a particularly useful result. However, we will see that this ILP can in fact be rounded iteratively in a manner that we were unable to pull off with the earlier ILP formulations, resulting in a polynomial-time algorithm for partitioning sporadic tasks upon heterogeneous multiprocessors that has speedup bound of ≈ 7.83 , thereby improving the ≈ 12.9 speedup bound of Marchetti-Spaccamela et al. [19].

As before, we use variables x_{ij} for each pair $(i, j) \in N \times M$, modeling the assignment of τ_i to π_j . Apart from the usual assignment constraints, the first constraints we consider are the utilization bounds on the tasks assigned to the same processor. That is, we require that

$$\sum_{i \in N} u_{ij} x_{ij} \leq 1 \quad \forall j \in M. \quad (10)$$

Now, let ρ denote any constant, $\rho > 1$. We de-

fine the function $r(x) := \rho^{\lceil \log_\rho x \rceil}$, and the set $D_\rho := \{\rho^0, \rho^1, \dots, r(d_{\max})\}$. We want to express the requirement that for all tasks assigned to the same processor with deadline at most ρ^k , the sum of their WCETs is at most ρ^k . Note that this is the set of tasks with $r(d_i) \leq \rho^k$. For technical reasons that will become apparent later (in Lemma V.2), we adopt the slightly weaker constraint

$$\sum_{i \in N: r(d_i) \leq d} c_{ij} \left(1 - \frac{d_i}{p_i}\right) x_{ij} \leq d \quad \forall d \in D_\rho, \forall j \in M. \quad (11)$$

We call these constraints (11) the *relaxed dbf constraints*. It is clear that these constraints have to be fulfilled by any feasible task assignment. (In particular, if $d_i \leq \rho^k$ and $x_{ij} = 1$, then $\text{dbf}_{f,j}(\rho^k) \geq c_{ij} > c_{ij}(1 - d_i/p_i)$, where f is the assignment represented by $\mathbf{x}$). We therefore arrive at the following ILP, denoted poly-ILP.

$$\sum_{j \in M} x_{ij} = 1 \quad \forall i \in N \quad (12a)$$

$$\sum_{i \in N} u_{ij} x_{ij} \leq 1 \quad \forall j \in M \quad (12b)$$

$$\sum_{i \in N: r(d_i) \leq d} c_{ij} \left(1 - \frac{d_i}{p_i}\right) x_{ij} \leq d \quad \forall d \in D_\rho, \forall j \in M \quad (12c)$$

$$x_{ij} \in \{0, 1\} \quad \forall i \in N, \forall j \in M$$

s.t. $c_{ij} \leq d_i. \quad (12d)$

If poly-ILP is infeasible, then there can be no feasible task assignment. Now assume that it is feasible and consider its relaxation, which is obtained by replacing each constraint (12d) by

$$x_{ij} \geq 0 \quad \forall i \in N, \forall j \in M \text{ s.t. } c_{ij} \leq d_i. \quad (13)$$

Since it is an LP and not an ILP, the relaxation can be solved in polynomial time. Let $\mathbf{x}^*$ denote its solution. For each $j \in M$ and deadline $d \in D_\rho$, we compute the value

$$b_{j,d} := \sum_{i \in N: r(d_i) = d} c_{ij} (1 - d_i/p_i) x_{ij}^*.$$

Note that, by (12c),

$$\sum_{d' \leq d} b_{j,d'} \leq d \quad \forall d \in D_\rho, \forall j \in M. \quad (14)$$

Based on these computed values, we define a variation of poly-ILP, denoted by sparse-ILP in the sequel. We obtain the latter by replacing the constraints (12c) with the following set of constraints:

$$\sum_{i \in N: r(d_i) = d} c_{ij} \left(1 - \frac{d_i}{p_i}\right) x_{ij} \leq b_{j,d} \quad \forall d \in D_\rho, \forall j \in M.$$

By dividing both sides of the inequality by d , these constraints can also be written as

$$\sum_{i \in N: r(d_i) = d} \bar{c}_{ij} \left(1 - \frac{d_i}{p_i}\right) x_{ij} \leq \bar{b}_{j,d} \quad \forall d \in D_\rho, \forall j \in M. \quad (15)$$

where $\bar{b}_{j,d} = b_{j,d}/d \leq 1$ (since $\mathbf{x}^*$ is feasible for the relaxation of poly-ILP) and $\bar{c}_{ij,d} = c_{ij}/d \leq 1$ (since if $r(d_i) = d$ then $c_{ij} \leq d_i \leq d$). Let $\mathcal{A}$ be the set of vectors $\mathbf{x}$ satisfying (12a). We can now express sparse-ILP in matrix notation as

$$\{\mathbf{x} \in \mathcal{A} \cap \{0, 1\}^{N \times M} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\},$$

where $\mathbf{A}$ and $\mathbf{b}$ are, respectively, the matrix of coefficients and the vector of right hand sides of the constraints (12b) and (15). Note that all entries of $\mathbf{A}$ and $\mathbf{b}$ take values between 0 and 1.

By construction, if $\mathbf{x}^*$ is a feasible solution for the LP relaxation of poly-ILP it is also feasible for the LP relaxation of sparse-ILP, and if the LP relaxation of sparse-ILP is infeasible, then no feasible task assignment exists. Our goal will be to round $\mathbf{x}^*$ to an *integral* vector which *approximately* satisfies the constraints of sparse-ILP.

A reason for preferring sparse-ILP to poly-ILP is that the former is an ILP formulation in which the constraint matrix is sparse: each variable appears in only a small number of constraints. This *sparsity* gives the potential to derive efficient rounding schemes which result in integral solutions, violating the relaxed dbf-constraints only by constant factors. We present such a rounding scheme below; to this end, the following theorem shows that—even when violated up to constant factors—the relaxed dbf constraints (15), together with the utilization constraints (12b), are approximately sufficient.

Theorem V.1. *Let $\beta \geq 1$ and let $f : N \rightarrow M$ be an assignment encoded by a vector $\hat{\mathbf{x}} \in \mathcal{A} \cap \{0, 1\}^{N \times M}$ such that, for each $j \in M$ and $d \in D_\rho$,*

$$\sum_{i \in N} u_{ij} \hat{x}_{ij} \leq \beta \quad (16)$$

and

$$\sum_{i \in N : r(d_i) = d} c_{ij} \left(1 - \frac{d_i}{p_i}\right) \hat{x}_{ij} \leq b_{j,d} + (\beta - 1) \cdot d. \quad (17)$$

Then $\text{dbf}_{f,j}(s) \leq (\beta + \rho + (\beta - 1)\rho^2/(\rho - 1))s$ for all $s \geq 0$. In particular, if $\beta \leq 2$, f is a feasible assignment under a speedup factor of $(2 + \rho + \rho^2/(\rho - 1))$.

Proof. For any $s \geq 0$ and $j \in M$, we bound $\text{dbf}_{f,j}(s)$ as follows:

$$\begin{aligned} \text{dbf}_{f,j}(s) &= \sum_{i \in N : d_i \leq s} \left\lfloor \frac{s + p_i - d_i}{p_i} \right\rfloor c_{ij} \hat{x}_{ij} \\ &\leq \sum_{i \in N : d_i \leq s} \left(s \frac{c_{ij}}{p_i} + c_{ij} \left(1 - \frac{d_i}{p_i}\right) \right) \hat{x}_{ij} \\ &\leq \sum_{i \in N : r(d_i) \leq r(s)} \left(s \frac{c_{ij}}{p_j} + c_{ij} \left(1 - \frac{d_i}{p_i}\right) \right) \hat{x}_{ij} \\ &\leq s \sum_{i \in N} \frac{c_{ij}}{p_i} \hat{x}_{ij} + \sum_{i \in N : r(d_i) \leq r(s)} c_{ij} \left(1 - \frac{d_i}{p_i}\right) \hat{x}_{ij} \\ &= s \sum_{i \in N} u_{ij} \hat{x}_{ij} + \sum_{k=0}^{\log_\rho(r(s))} \sum_{i \in N : r(d_i) = \rho^k} c_{ij} \left(1 - \frac{d_i}{p_i}\right) \hat{x}_{ij} \end{aligned}$$

$$\begin{aligned} &\stackrel{(16)}{\leq} \beta s + \sum_{k=0}^{\log_\rho(r(s))} \sum_{i \in N : r(d_i) = \rho^k} c_{ij} \left(1 - \frac{d_i}{p_i}\right) \hat{x}_{ij} \\ &\stackrel{(17)}{\leq} \beta s + \sum_{k=0}^{\log_\rho(r(s))} (b_{j,\rho^k} + (\beta - 1)\rho^k) \\ &\stackrel{(14)}{\leq} \beta s + \rho^{\log_\rho(r(s))} + \sum_{k=0}^{\log_\rho(r(s))} (\beta - 1)\rho^k \\ &= \beta s + r(s) + (\beta - 1) \sum_{k=0}^{\log_\rho(r(s))} \rho^k \\ &= \beta s + r(s) + (\beta - 1) \cdot \frac{\rho^{\log_\rho(r(s))+1} - 1}{\rho - 1} \\ &= \beta s + r(s) + (\beta - 1) \frac{\rho \cdot r(s) - 1}{\rho - 1} \\ &\leq (\beta + \rho + (\beta - 1) \frac{\rho^2}{\rho - 1}) s. \end{aligned}$$

The last inequality follows from $r(s) \leq \rho \cdot s$. $\square$

To construct $\hat{\mathbf{x}}$, we adopt an iterative rounding procedure that is similar to the procedure presented in [16], [19]. The idea of the iterative rounding procedure is the following. In each iteration k , we first compute an extreme point solution $\mathbf{x}^k$ of a linear program LP^k , where LP^0 is the relaxation of sparse-ILP, and each LP^k is obtained by fixing the value for some variables or removing some constraints of LP^{k-1} .

Given a feasible fractional solution $\mathbf{x}^k$, to define LP^{k+1} we first fix all variables which are integral in $\mathbf{x}^k$, i.e., those variables are not allowed to be changed anymore in the remainder of the procedure. Let s be the number of variables in LP^k and let r_a , r_b and r_c be the number of constraints of types (12a), (12b) and (15), respectively. Let $r = r_a + r_b + r_c$. To obtain LP^{k+1} we either delete one or more variables, in case $s > r$, or delete a constraint while ensuring that in the final solution that constraint will not be violated too much. Along the way we ensure that the constraints of type (12a) are always satisfied exactly, so that $\mathbf{x}^k \in \mathcal{A}$ at all times.

Note that if there is some variable x_{ij} that is fixed at value 1 and removed from the program, then for all $j' \in M \setminus \{j\}$, $x_{ij'}$ will be set to 0 and also be removed from the program. The constraint of type (12a) corresponding to this i is then superfluous and will also be removed.

To derive the bounds (16)–(17), we need to study the coefficient matrix $\mathbf{A}$ in more detail. Let γ be the maximum, over all x_{ij} , of the sum of the values of the coefficients of variable x_{ij} in constraints (12b) and (15). We first derive a bound on γ .

Lemma V.2. *For any task set τ , $\gamma \leq 1$.*

Proof. Observe that γ is just the maximum value of $u_{ij} + \bar{c}_{ij,d}(1 - d_i/p_i)$ across all variables x_{ij} in the program. Recall that for all such pairs (i, j) , $c_{ij} \leq d_i$, i.e., $\bar{c}_{ij,d} \leq 1$, otherwise the variable x_{ij} is forced to 0 and removed from the LP. We

can now bound

$$u_{ij} + \bar{c}_{ijd} \left(1 - \frac{d_i}{p_i}\right) \leq \frac{c_{ij}}{p_i} + 1 - \frac{d_i}{p_i} \leq \frac{c_{ij}}{p_i} + 1 - \frac{c_{ij}}{p_i} = 1.$$

□

The following technical lemma is instrumental to our rounding procedure. It is a specialization to our setting of a more general rounding result for assignment LPs, from another paper currently under review [8].

Lemma V.3. *Let LP^k be the linear program that is solved in iteration k of the rounding procedure, with s variables and r constraints. Let $\mathbf{x}^k$ be an extreme point solution to this LP. Then either,*

- (i) $\mathbf{x}^k$ has at least one integer component; or
- (ii) there is $j \in M$ and a corresponding constraint of type (12b) such that $\sum_{i \in N} u_{ij} z_{ij} - \sum_{i \in N} u_{ij} x_{ij}^k \leq \gamma$ for any integer solution $\mathbf{z}$; or
- (iii) there are $j \in M$, $d \in D_\rho$ and a corresponding constraint of type (15) such that

$$\sum_{i \in N: r(d_i)=d} \bar{c}_{ijd} \left(1 - \frac{d_i}{p_i}\right) (z_{ij} - x_{ij}^k) \leq \gamma$$

for any integer solution $\mathbf{z}$.

Proof. Let $\mathbf{A}$ be the coefficient matrix of LP^k . If $s > r$, the null space of $\mathbf{A}$ is nontrivial, so let $\mathbf{x}_0$ be a nonzero vector in the null space. Since $\mathbf{x}^k$ is an extreme-point solution to LP^k , it cannot be expressed as the convex combination of two (or more) solutions to LP^k . If $\mathbf{x}^k$ does not have any integral entry, then we can find a value $\delta > 0$ such that $\mathbf{x}^k + \delta \mathbf{x}_0$ and $\mathbf{x}^k - \delta \mathbf{x}_0$ are both solutions to LP^k (since $\mathbf{A}(\mathbf{x} \pm \delta \mathbf{x}_0) = \mathbf{A}\mathbf{x}$) and, in particular, $\mathbf{x}^k$ is a convex combination of these two solutions. Therefore $\mathbf{x}^k$ must have at least one integral entry.

If $s \leq r$, we show that there always exists a constraint of type (12b) or type (15) such that $\max_{\mathbf{z} \in S} \{(\mathbf{A}\mathbf{z})_l - (\mathbf{A}\mathbf{x})_l\} \leq \gamma$, where γ is the maximum sum of coefficients in a column of constraints (12b) and (15) and where S is the integer solution space for all remaining variables, i.e., $S = \{0, 1\}^s$.

We show the statement by contradiction. Assume that the statement is not true, that is, for each constraint l of type (12b) or (15) it holds that there exists a vector $\mathbf{z} \in S$ such that

$$(\mathbf{A}\mathbf{z})_l - (\mathbf{A}\mathbf{x})_l > \gamma. \quad (18)$$

Note that all variables still present in the linear program correspond to a processor $j \in M$ and a task $i \in N$ that is not yet assigned fully to one processor, but fractionally to multiple processors. Hence, the constraint of type (12a) corresponding to each τ_i is still present in the linear program. It follows that

$$\sum_{j \in M} \sum_{i \in N: x_{ij} \in (0,1)} x_{ij} = r_a, \quad (19)$$

where r_a is the number of constraints of type (12b) remaining in LP^k . Define L as the set of constraints of types (12b) and (15) present in the current linear program, and let r_b and r_c be their number, respectively (so $r_a + r_b + r_c = r$). For any

$q = (i, j)$, let L^q denote the set of these constraints containing variable x_q ; by definition of γ , we have

$$\max_q \sum_{l \in L^q} a_{lq} \leq \gamma. \quad (20)$$

Then,

$$\begin{aligned} \gamma(r - r_a) &= \gamma(r_b + r_c) \\ &\stackrel{(18)}{<} \sum_{l \in L} \max_{\mathbf{z} \in S} ((\mathbf{A}\mathbf{z})_l - (\mathbf{A}\mathbf{x})_l) \\ &\stackrel{\text{as all } a_{lq} \geq 0}{=} \sum_{l \in L} ((\mathbf{A}\mathbf{1})_l - (\mathbf{A}\mathbf{x})_l) \\ &= \sum_{l \in L} \sum_q a_{lq} (1 - x_q) \\ &= \sum_q \sum_{l \in L^q} a_{lq} (1 - x_q) \\ &\leq \sum_q \gamma (1 - x_q) \\ &= \gamma s - \sum_q \gamma x_q \\ &= \gamma(s - r_a). \end{aligned} \quad (21)$$

The second inequality follows from (20).

The chain of inequalities implies that $\gamma(r - r_a) < \gamma(s - r_a) \Rightarrow r < s$ which is a contradiction to being in the case that $s \leq r$. Hence we conclude that if $s \leq r$, there must be a constraint l of type (12b) or (15) for which $\max_{\mathbf{z} \in S} \{(\mathbf{A}\mathbf{z})_l - (\mathbf{A}\mathbf{x})_l\} \leq \gamma$. □

Lemma V.3 is used to guide the rounding process. If Case (i) applies, the variables that have an integer value are fixed at that value and removed from the LP. If a variable x_{ij} is fixed at value 1, then for all $j' \in M \setminus \{j\}$, the variables $x_{ij'}$ are fixed at value 0 and the constraint of type (12a) corresponding to i is removed. If we are in Case (ii) or (iii), the constraint for which the claim holds can be found in polynomial time by checking, for each constraint $l \in L$ of type (12b) or (15), whether $\sum_q (1 - x_q) \leq \gamma$. This is sufficient since all $a_{lq} \geq 0$ and the maximum value any variable x_q can take is 1. If such a constraint is of type (12b) (Case (ii)) or (15) (Case (iii)), the final task assignment will satisfy (16) or (17) for that constraint, respectively, *even if the constraint is dropped*; thus, we drop the constraint, obtaining the next (smaller) LP.

After either all constraints have been removed or the values of all variables have been fixed at an integer value, we obtain an integral vector $\hat{\mathbf{x}}$ which satisfies $\sum_{i \in N} u_{ij} \hat{x}_{ij} \leq 1 + \gamma$ for each $j \in M$ and $\sum_{i \in N: r(d_i)=d} \bar{c}_{ijd} \left(1 - \frac{d_i}{p_i}\right) \hat{x}_{ij} \leq \bar{b}_{j,d} + \gamma$ for all $j \in M$ and all deadlines $d \in D_\rho$. Hence, the vector $\hat{\mathbf{x}}$ satisfies constraints (16), (17) with $\beta := 1 + \gamma$. We are now in the position to invoke Theorem V.1 to obtain our final guarantee.

Theorem V.4. *There is a polynomial-time partitioning algorithm with a speedup bound of $(5 + 2\sqrt{2}) \approx 7.83$ for the*

problem of assigning constrained-deadline tasks to heterogeneous processors.

Proof. All steps required to construct $\hat{\mathbf{x}}$ can be carried out in polynomial time. The assignment induced by $\hat{\mathbf{x}}$ satisfies (16)–(17) with $\beta = 1 + \gamma \leq 2$ (Lemma V.2). Thus, by Theorem V.1 with $\rho = 1 + 1/\sqrt{2}$, the assignment induced by $\hat{\mathbf{x}}$ is feasible with speedup

$$2 + \rho + \frac{\rho^2}{\rho - 1} = 5 + 2\sqrt{2} \approx 7.83. \quad \square$$

VI. SCHEDULABILITY EXPERIMENTS

In the sections above, we saw how the problem of partitioned scheduling of sporadic task systems upon unrelated multiprocessors could be modeled by ILPs. Our motivation for doing so is that the optimization community has devoted immense effort to coming up with extremely efficient (although still exponential-time, since solving ILPs is NP-hard) algorithms for solving ILPs, and highly-optimized libraries implementing these efficient algorithms are widely available. This is particularly true for ILPs like the ones we have constructed above, in which each variable is further constrained to take in only the values zero or one. In this section, we validate the performance of our ILP-based schedulability tests against synthetic workloads in terms of the percentage of schedulable task sets. Our results strongly suggest that the ILP model discussed in Section IV is superior in practice.

A. Generation of the task sets and solutions

We developed a parametric framework with several parameters (m , κ , p , $\bar{U}$, α – they are detailed below) to randomly generate our workloads; this framework is general enough to support our entire range of experiments.

We consider m -processor platforms and $n = \kappa m$ tasks, with $\kappa \geq 1$ an integer-valued parameter. We randomly generate an affinity mask $R_{i,j}$ for each $i \in N$ and $j \in M$: $R_{i,j} \leftarrow 1$ with probability p and 0 with probability $(1 - p)$. These affinity masks help define the $C_{i,j}$ values: $C_{i,j}$ has a value $< \infty$ if and only if $R_{i,j} = 1$. If the generated mask does not allow a particular task to be processed on any processor, we then allow that task to be processed upon a randomly chosen processor.

We then generate utilization values for every allowed pair (i, j) for which $R_{i,j} = 1$. Tasks are grouped into m groups of size κ each: tasks $\tau_{1+(k-1)\kappa}$ to $\tau_{k\kappa}$ form the k th group. For each group we use the UNISORT algorithm [7] to generate randomly distributed utilizations with total value $\bar{U}$ for the allowed task-processor pairs in the group. Note that since there are m groups, each of total utilization $\bar{U}$, the value $\bar{U}$ represents the average load that a processor can expect if tasks are randomly assigned. Note that $\bar{U} \leq 1$ is *not* a necessary condition for schedulability (indeed some of our ILP formulations are able to schedule task sets with $\bar{U} > 1$).

We generate the periods by setting $p_i = 2^{\Delta_i}$, with each Δ_i uniformly distributed in the range 3...10. The worst-case execution times are computed directly from the periods and utilizations. Finally, the relative deadline of each task is

Figure 1. Experimental results I. The diagram shows the percentage of task sets found to be schedulable by Corollary III.4 (Eqs. (3)–(5)) and by Theorem IV.2 (Eqs. (7)–(9)) with $k = 3$. (a) Variation of $\bar{U}$ ($m = 10$, $\kappa = 10$, $p = 0.5$, $\alpha = 0.2$).

sampled uniformly in the range $[(1 - \alpha) \cdot (\max_j c_{ij}) + \alpha p_i, p_i]$, where $\alpha \in [0, 1]$.

In the experiments we consider the two ILP models discussed in Section III-B (Eqs. (3)–(5)) and Section IV (Eqs. (7)–(9)), the latter with $k = 3$. All optimization models have been solved by using a branch-and-cut approach implemented in the mathematical programming solver Gurobi 6.50 [1] on a PC with Intel i7-4770 CPU at 3.4 GHz and 16Gb RAM. Instances of the ILP model of Section III-B are solved within 3 seconds, and within less than 1 second on average; instances of the ILP model of Section IV are solved within 109 seconds, and within less than 15 seconds on average.

B. Discussion of the results

a) Variation of $\bar{U}$ (Figure 1): In the first type of experiment, we use the average load $\bar{U}$ as the independent variable. We vary $\bar{U}$ from 0.2 to 1.5. We fix $m = 10$, $\kappa = 10$, $p = 0.5$, $\alpha = 0.2$. Again, for each experiment, we generated 10 task sets for each value on the x -axis. When the percentage of schedulable task sets was strictly between 0 and 1, we generated 20 additional task sets to achieve a higher precision. Indeed, as could be expected, schedulability decreases when the load factor $\bar{U}$ increases. It is interesting to note that for the strengthened ILP model of Theorem IV.2, more than 90% of the task sets are schedulable as long as $\bar{U} \leq 1$.

b) Variation of m (Figure 2(a)): In the second type of experiment, we look at the impact of the number of processors on schedulability, so we vary m from 2 to 10. We fix $\kappa = m$, so $n = m^2$. In this set of experiments, we fix $\bar{U} = 1$, $p = 1$, $\alpha = 0.2$. For each experiment, we generated 10 task sets for each value on the x -axis. Figure 1 shows the dependency on m of the percentage of task sets are guaranteed to be schedulable by Corollary III.4 (for the first ILP) and Theorem IV.2 (for the second ILP), the latter using a formulation in which $k = 3$. In both cases, schedulability increases with the number of processors even though $\bar{U}$ is fixed. This is due to the fact that, for a fixed total utilization value of each group, a higher

Figure 2. Experimental results II. The diagrams show the percentage of task sets found to be schedulable by Corollary III.4 (Eqs. (3)-(5)) and by Theorem IV.2 (Eqs. (7)-(9)) with $k = 3$. (a) Variation of m ($\kappa = m$, $\bar{U} = 1$, $p = 1$, $\alpha = 0.2$). (b) Variation of p ($m = 10$, $\kappa = 10$, $\bar{U} = 1$, $\alpha = 0.2$).

number of tasks in a group implies lower utilization values for the individual tasks, and therefore a more efficient partitioning.

c) *Variation of p (Figure 2(b))*: In the third type of experiment, we control the sparsity of the processor affinity matrix R . We vary p from 0.2 to 0.9. We fix $m = 10$, $\kappa = 10$, $\bar{U} = 1$, $\alpha = 0.2$. For each experiment, we generated 10 task sets for each value on the x -axis. We find out that the sparsity has a high impact on schedulability: there are clear schedulability thresholds around $p = 0.7$ (for the first ILP) and $p = 0.45$ (for the second ILP). This is not entirely unexpected, as when the affinity matrix is sparser, it may happen that several tasks of large combined utilization can only be assigned to a small set of processors.

VII. SUMMARY AND CONCLUSIONS

In this work, we proposed a partitioning approach for constrained-deadline tasks on heterogeneous (unrelated) processors. The approach is based on integer linear programming formulations and allows the derivation of guaranteed speedup bounds and consequently, sufficient schedulability tests.

Experiments among randomly generated task workloads clearly show that one of the proposed approaches is viable in terms of computation time and not as pessimistic in terms of schedulability as could be expected, especially when the task-processor affinity relation is dense.

ACKNOWLEDGMENTS

Work supported by NSF grants CNS 1115284, CNS 1218693, CNS 1409175, and CPS 1446631, AFOSR grant FA9550-14-1-0161, ARO grant W911NF-14-1-0499, and a grant from General Motors Corp. This work has also been partially supported by the research project “Designing Human-Agent Collectives for Sustainable Future Societies” (C26A15TXCF) of Sapienza University of Rome.

REFERENCES

- [1] Gurobi Optimizer. <http://www.gurobi.com>.
- [2] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. *Random Structures & Algorithms*, 38(3):251–268, 2011.
- [3] K. Albers and F. Slomka. An event stream driven approximation for the analysis of real-time systems. In *Proc. 16th Euromicro Conf. on Real-Time Systems*, pages 187–195, Los Alamitos, CA, 2004. IEEE.
- [4] S. Baruah and E. Bini. Partitioned scheduling of sporadic task systems: an ILP-based approach. In *Proceedings of the 2008 Conference on Design and Architectures for Signal and Image Processing*, 2008.
- [5] S. Baruah, A. Mok, and L. Rosier. Preemptively scheduling hard-real-time sporadic tasks on one processor. In *Proceedings of the 11th Real-Time Systems Symposium*, pages 182–190, Orlando, Florida, 1990. IEEE Computer Society Press.
- [6] S. K. Baruah and N. Fisher. The partitioned multiprocessor scheduling of deadline-constrained sporadic task systems. *IEEE Transactions on Computers*, 55(7):918–923, 2006.
- [7] E. Bini and G. C. Buttazzo. Measuring the performance of schedulability tests. *Real-Time Systems*, 30(1-2):129–154, 2005.
- [8] V. Bonifaci, G. D’Angelo, and A. Marchetti-Spaccamela. Hierarchical and semi-partitioned machine scheduling. TR-15-03, IASI-CNR, 2015.
- [9] J. Chen and S. Chakraborty. Resource augmentation for uniprocessor and multiprocessor partitioned scheduling of sporadic real-time tasks. *Real-Time Systems*, 49(4):475–516, 2013.
- [10] H. S. Chwa, J. Seo, J. Lee, and I. Shin. Optimal real-time scheduling on two-type heterogeneous multicore platforms. In *Real-Time Systems Symposium (RTSS), 2015 IEEE*, Dec 2015.
- [11] M. Dertouzos. Control robotics : the procedural control of physical processors. In *Proceedings of the IFIP Congress*, pages 807–813, 1974.
- [12] A. Frank and É. Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987.
- [13] E. C. Freuder. A sufficient condition for backtrack-bounded search. *J. ACM*, 32(4):755–761, 1985.
- [14] S. Kamath. Unrelated parallel machine scheduling—perspectives and progress. *OPSEARCH*, 48(4):318–334, 2011.
- [15] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.
- [16] R. M. Karp, F. T. Leighton, R. L. Rivest, C. D. Thompson, U. V. Vazirani, and V. V. Vazirani. Global wire routing in two-dimensional arrays. *Algorithmica*, 2:113–129, 1987.
- [17] L. C. Lau, R. Ravi, and M. Singh. *Iterative Methods in Combinatorial Optimization*. Cambridge University Press, 2011.
- [18] C. Liu and J. Layland. Scheduling algorithms for multiprogramming in a hard real-time environment. *Journal of the ACM*, 20(1):46–61, 1973.
- [19] A. Marchetti-Spaccamela, C. Ruten, S. van der Ster, and A. Wiese. Assigning sporadic tasks to unrelated machines. *Mathematical Programming*, pages 1–28, 2014.
- [20] G. Raravi. *Real-Time Scheduling on Heterogeneous Multiprocessors*. PhD thesis, Technical Institute of Porto (Portugal), 2014.
- [21] G. Raravi, B. Andersson, and K. Bletsas. Assigning real-time tasks on heterogeneous multiprocessors with two unrelated types of processors. *Real-Time Systems*, 49(1):29–72, 2013.

- [22] G. Raravi, B. Andersson, V. Nélis, and K. Bletsas. Task assignment algorithms for two-type heterogeneous multiprocessors. *Real-Time Systems*, 50(1):87–141, 2013.
- [23] G. Raravi and V. Nélis. A PTAS for assigning sporadic tasks on two-type heterogeneous multiprocessors. In *Proceedings of the 33rd IEEE Real-Time Systems Symposium, RTSS 2012, San Juan, PR, USA, December 4-7, 2012*, pages 117–126, 2012.
- [24] A. Wiese, V. Bonifaci, and S. Baruah. Partitioned EDF scheduling on a few types of unrelated multiprocessors. *Real-Time Systems*, 49(2):219–238, 2013.