

An Optimization Framework for Real-Time Systems with Sustainable Schedulability Analysis

Yecheng Zhao*, Runzhi Zhou⁺, and Haibo Zeng*

*Virginia Tech, USA. *Emails:* {zyecheng, hbzeng}@vt.edu

⁺Case Western Reserve University, USA. *Email:* rxz229@case.edu



Abstract—The design of modern real-time systems not only needs to guarantee their timing correctness, but also involves other critical metrics such as control quality and energy consumption. As real-time systems become increasingly complex, there is an urgent need for efficient optimization techniques that can handle large-scale systems. However, the complexity of schedulability analysis often makes it difficult to be directly incorporated in standard optimization frameworks, and inefficient to be checked against a large number of candidate solutions. In this paper, we propose a novel optimization framework for the design of real-time systems. It leverages the sustainability of schedulability analysis that is applicable for a large class of real-time systems. It builds a counterexample-guided iterative procedure to efficiently learn from an unschedulable solution and rule out many similar ones. Compared to the state-of-the-art, the proposed framework may be ten times faster while providing solutions with the same quality.

I. INTRODUCTION

Design optimization techniques are becoming vital and urgent for a number of application domains for real-time systems. For example, the automotive industry is extremely cost sensitive yet its products are highly safety critical [13]. Unmanned aerial vehicles powered by batteries must carefully plan and operate according to their tight energy budget [18].

There has been a rich set of research on the development of timing and schedulability analysis over the past years. However, many of the analysis techniques are either impossible, or too complex and inefficient, to use in well-established optimization frameworks (i.e., mathematical programming) [39]. As a result, the existing practice for design optimization often has to rely on ad-hoc approaches. This typically comes with the loss of solution quality as well as limited applicability. Hence, an efficient and general framework that bridges the gap between schedulability analysis and optimization is critical to the future success of real-time systems design.

In this paper, we seek to address this urgent need and propose a general framework for optimizing real-time systems. We leverage the concept of sustainable schedulability analysis that is recommended as a good engineering practice for real-time systems [5]. Specifically, if a task system is schedulable under a sustainable schedulability analysis, then it should remain to be schedulable with, for example, decreased worst case execution time (WCET), or increased period.

The design of our framework centers around the concept of Maximal Unschedulable Assignment (MUA) to variables that are sustainable in the schedulability analysis, including task WCET and period. It avoids the direct formulation of

schedulability region, but uses MUAs to provide an efficient abstraction of the schedulability constraints. It develops a counterexample (i.e., unschedulable solutions) guided paradigm to learn from an unschedulable solution and rule out many similar unschedulable ones. It also builds an MUA-driven branching algorithm that takes advantage of the special structure in the optimization problem.

Our framework is generally applicable to a broad range of optimization problems for real-time systems. We use two case studies to demonstrate the benefit of our framework. The first is the optimization of energy consumption on platforms with dynamic voltage and frequency scaling (DVFS), where task WCETs change based on the selected CPU frequency. The second is the selection of task periods to optimize control quality under schedulability constraints [24]. Compared to the state-of-the-art, the proposed framework may be ten times faster while providing solutions with the same quality.

Paper organization. The rest of the paper is organized as follows. Section II summarizes the related work. Section III presents the system model and the problem setting that fits our framework. Section IV describes the concept of Maximal Unschedulable Assignment (MUA). Section V develops the framework based on MUA. Section VI discusses the applicability and limitations. Section VII shows the experiments to demonstrate the advantage of our framework. Finally, Section VIII concludes the paper.

II. RELATED WORK

There is a rich literature for the optimization of real-time systems. Generally speaking, the current approaches can be classified into four categories: (i) meta heuristics such as simulated annealing (e.g., [7], [32]) and genetic algorithm (e.g., [17], [31]); (ii) problem specific heuristics (e.g., [29], [34], [38]); (iii) directly applying existing optimization frameworks such as branch-and-bound (BnB) (e.g., [35]), Mixed Integer Linear Programming (MILP) (e.g., [26], [37]), and convex programming (e.g., [20]); (iv) customized optimization frameworks that are tuned for specific design variables in real-time systems, such as the optimization of priority assignment [40], [42], [43] and period [39]. The first two categories either do not have any guarantee on solution quality, or suffer from scalability issues and may have difficulty to handle large industrial designs. For the third category, besides the possible scalability issues, it also requires that the schedulability analysis can be formulated in the chosen

framework, which may not always be possible. The current approaches in the fourth category are also limited in their applicability. For example, unlike the proposed framework in this paper, none of them [39], [40], [42], [43] considers task WCETs as design variables.

The problem of task period selection has been studied in the literature [8], [30]. In particular, Bini et al. [8] propose a problem specific branch-and-bound technique for optimizing task rate, but it describes the exact schedulability region in the domain of task rates, hence it may be applicable to problems with optimization objectives that depend only on task periods (not on other variables such as task response times). Mancuso et al. [24] develop a branch-and-bound algorithm, where a linear lower bound is adopted as an approximation to task response time. In [31], a genetic algorithm is used for the problem to minimize the sum of end-to-end delays in networked control systems. Davare et al. consider the period optimization to minimize end-to-end latencies for a set of paths, and formulate it in mixed integer geometric programming (MIGP) framework [10]. Differently, Zhao et al. [39] propose a customized procedure specialized for the minimization of end-to-end latencies, which is several orders of magnitude faster than [10]. Our approach is also to develop a customized framework. However, it not only is more generally applicable than [39], but also may run $100\times$ faster.

There are various approaches proposed to address the problem of optimizing energy consumption for systems with DVFS, see a related review in [4]. However, they are all focusing on one particular scheduling model and associated schedulability analysis. For example, Huang et al. [20] consider mixed-criticality systems scheduled with Earliest Deadline First with Virtual Deadline, which allows to formulate the problem as a convex program. Instead, our framework is generally applicable to any systems as long as the schedulability analysis is sustainable.

In summary, compared to the existing approaches, our framework is applicable to a larger class of optimization problems in real-time systems. It does not pertain to a particular scheduling model or schedulability analysis, but can be used for any systems with sustainable schedulability analysis. It applies to the optimization of various decision variables including task WCET, period, deadline, or priority assignment. Finally, it may still be much faster than the state-of-the-art, as demonstrated in the experimental results.

III. SYSTEM MODEL

In this paper, we consider a general setting of real-time systems for which the associated schedulability analysis is sustainable [5]. It contains m tasks indexed from 1 to m . The design optimization of a real-time system is to select the appropriate values for design variables such that (a) a given cost function is minimized, and (b) system schedulability is satisfied. Mathematically it can be expressed as follows

$$\begin{aligned} & \min F(\mathbf{X}) \\ \text{s.t.} & \text{system schedulability} \\ & x_j \in [x_j^l, x_j^u], \forall j = 1, \dots, n \end{aligned} \quad (1)$$

where $\mathbf{X} = (x_1, \dots, x_n)$ is the vector of variables. Here we first focus on the case that \mathbf{X} may include the response time R_i , WCET C_i , deadline D_i and/or period T_i for any task τ_i . In Section VI we will discuss how to handle the case where priority assignment is also part of the decision variables. Each variable x_j in \mathbf{X} takes integer values within a bounded range $[x_j^l, x_j^u]$ (i.e., the design variables have finite resolutions). We do not impose any particular form of schedulability analysis, as long as it is sustainable [5].

Sustainability is proposed as a guideline for the development of schedulability analysis techniques in real-time systems [5]. Specifically, a schedulability analysis is defined as sustainable if any schedulable task system remains schedulable with (i) decreased WCET C_i ; (ii) larger period T_i ; (iii) larger deadline D_i for any task τ_i , among others.¹

Here we discuss how to leverage the sustainability with respect to task deadlines to handle the case that response time R_i appears in the objective function. In this case, we not only need to make sure that $R_i \leq D_i$ (which can be satisfied by any R_i that is no larger than D_i), but also a precise value of R_i in order to compare different schedulable solutions. In this case, we replace R_i with a virtual deadline \hat{D}_i [42], which can be interpreted as a safe estimation on R_i (hence $R_i \leq \hat{D}_i \leq D_i$) when the system is schedulable. It is easy to see any schedulability analysis that is sustainable with respect to the deadline D_i is also sustainable with respect to \hat{D}_i .

For systems with sustainable schedulability analysis, without loss of generality, they satisfy the following property.

Property 1. The system schedulability constraints can be written as $G(\mathbf{X}) \leq 0$, where each function in $G(\mathbf{X})$ is **monotonically non-increasing** with respect to each variable x_j in \mathbf{X} . Hence, if a smaller assignment to x_j (e.g., the period, virtual deadline, or the additive inverse $-C_j$ of WCET C_j) makes the system schedulable, then a larger assignment to x_j also does (assuming all other variables remain unchanged).

Note that some variables such as WCET C_i may be opposite to the above property (smaller C_i corresponds to easier schedulability). In this case, we can simply perform a variable conversion (i.e., replacing C_i with $C'_i = -C_i$) to make it conform to the assumption.

A. Examples on Sustainable Schedulability Analysis

Sustainability is a general property that applies to many schedulability analysis techniques in real-time systems. For example, for the classical Liu-Layland task model scheduled with fixed priority [22], the response time based schedulability analysis [2] is sustainable

$$R_i = C_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i}{T_j} \right\rceil C_j \leq D_i \quad (2)$$

A more complicated example is mixed-criticality systems scheduled with Adaptive Mixed-Criticality (AMC) scheduling

¹For simplicity we call such a schedulability analysis sustainable, but it is termed as self-sustainable analysis in [3], see a detailed discussion therein.

policy [6]. The proposed AMC-rtb and AMC-max schedulability analyses [6] are both sustainable [16]. The two analyses mainly differ in the estimation of interferences from higher priority tasks during the criticality change.

AMC-rtb analysis calculates the response time of a HI-criticality task τ_i during the criticality change as

$$R_i(HI) = C_i(HI) + \sum_{\forall j \in hpH(i)} \left\lceil \frac{R_i(HI)}{T_j} \right\rceil C_j(HI) + \sum_{\forall j \in hpL(i)} \left\lceil \frac{R_i(LO)}{T_j} \right\rceil C_j(LO) \quad (3)$$

Here $C_i(HI)$ represents the WCET of τ_i in HI criticality mode. $R_i(LO)$ is the response time of τ_i in LO criticality mode given by Equation (2). $hpH(i)$ and $hpL(i)$ represent the set of HI- and LO-criticality tasks of higher priority than τ_i , respectively.

Intuitively, AMC-rtb assumes that a HI-criticality task always executes in HI-criticality mode and LO-criticality task may execute up to $R_i(LO)$. AMC-max improves upon AMC-rtb by considering different specific time instants of criticality change and dividing the workload of higher priority HI-criticality tasks into LO-mode and HI-mode. Specifically, given a time instant s of criticality change, AMC-max computes the WCRT of τ_i as follows

$$R_i(HI, s) = C_i(HI) + \sum_{\forall j \in hpL(i)} \left(\left\lfloor \frac{s}{T_j} \right\rfloor + 1 \right) C_j(LO) + \sum_{\forall j \in hpH(i)} M(j, s, R_i(HI, s)) C_j(HI) + \sum_{\forall j \in hpH(i)} \left(\left\lceil \frac{t}{T_j} \right\rceil - M(j, s, R_i(HI, s)) \right) C_j(LO) \quad (4)$$

where $M(j, s, t)$ represents the maximum number of instances of τ_j that are released as HI-criticality instances during the time interval $[s, t]$, which is expressed as

$$M(j, s, t) = \min \left\{ \left\lceil \frac{t - s - (T_j - D_j)}{T_j} \right\rceil + 1, \left\lceil \frac{t}{T_j} \right\rceil \right\} \quad (5)$$

The WCRT of τ_i during the criticality change can be computed by examining all possible time instants s of criticality change

$$R_i(HI) = \max_{\forall s} R_i(HI, s) \quad (6)$$

It is shown to be sufficient to only consider those s in the interval $[0, R_i(LO))$ [6].

B. Example Optimization Problems

We now provide two examples that fit our framework.

Optimizing control quality. The first problem is to optimize control performance for a set of periodic tasks scheduled on a uniprocessor [24]. The objective function, which represents the control cost, are approximated as a weighted sum of task period T_i and response time R_i for each task τ_i [24]

$$F(\mathbf{X}) = \sum_{i=1}^m \alpha_i T_i + \beta_i R_i \quad (7)$$

where α_i and β_i are given constant weights. Our framework works for any scheduling policy as long as its schedulability analysis is sustainable. In the experiments (Section VII-B) we assume the same as those in [24], i.e., the schedulability analysis in Equation (2).

Energy minimization with DVFS. Platforms with DVFS capabilities allow to adjust the CPU clock rate to save energy. Higher clock rate gives smaller WCET, which generally helps schedulability. However, this comes with an increased energy consumption. The goal is to determine the clock rate f_i for executing each task τ_i such that the system is schedulable while the total energy is minimized.

Specifically, suppose τ_i has an execution time C_i^b measured at a base clock rate f^b , then its execution time at another clock rate f_i can be estimated as $C_i = C_i^b \times \frac{f^b}{f_i}$. Thus $f_i = f^b \times \frac{C_i^b}{C_i}$. We normalize f^b to be 1 and consider that C_i^b is given, which makes C_i a decision variable in the optimization. We adopt the energy consumption objective formulated in [20]:

$$F(\mathbf{X}) = \sum_{\forall \tau_i} \frac{C_i^b f^b}{T_i} \cdot \beta \cdot (f_i)^{\alpha-1} = \sum_{\forall \tau_i} \frac{1}{T_i} \cdot \beta \cdot \frac{(C_i^b)^{\alpha}}{(C_i)^{\alpha-1}} \quad (8)$$

where β is a circuit-dependent constant. A common assumption for α is 3 [27], [28]. Like the previous case, we do not impose any constraint on the scheduling policy as long as the associated schedulability analysis is sustainable. In the experiments (Section VII-A), we use Equation (2) for a large number of random systems. For an industrial case study, we assume mixed-criticality systems scheduled with AMC, and adopt the AMC-rtb and AMC-max schedulability analyses [6], both of which are sustainable [16].

IV. MAXIMAL UNSCHEDULABLE ASSIGNMENT

We now introduce the concept of Maximal Unschedulable Assignment (MUA).

Definition 1. An assignment

$$\mathcal{X} = (v_1, \dots, v_n) \quad (9)$$

is a valuation of each variable $x_i = v_i$ in \mathbf{X} . An assignment $\mathcal{X} = (v_1, \dots, v_n)$ is said to **dominate** another assignment $\mathcal{X}' = (v'_1, \dots, v'_n)$, denoted as $\mathcal{X} \succeq \mathcal{X}'$, if \mathcal{X} is component-wise no smaller than \mathcal{X}' , i.e., $v_i \geq v'_i, \forall i$.

Definition 2. An assignment \mathcal{X} is said to be unschedulable if it violates the schedulability constraints. \mathcal{X} is a **maximal unschedulable assignment (MUA)** if (a) \mathcal{X} is unschedulable and (b) there is no other unschedulable assignment that dominates \mathcal{X} .

We remark that the concept of MUAs is well-defined, since by Property 1 of schedulability constraints with respect to the variables in \mathbf{X} , any assignment \mathcal{X}' that is dominated by an MUA \mathcal{X} must also be unschedulable. Also, any two MUAs cannot dominate each other, otherwise one of them is not an MUA. Note that we assume variables with integer values (or in general discrete variables), hence the MUAs always exist.

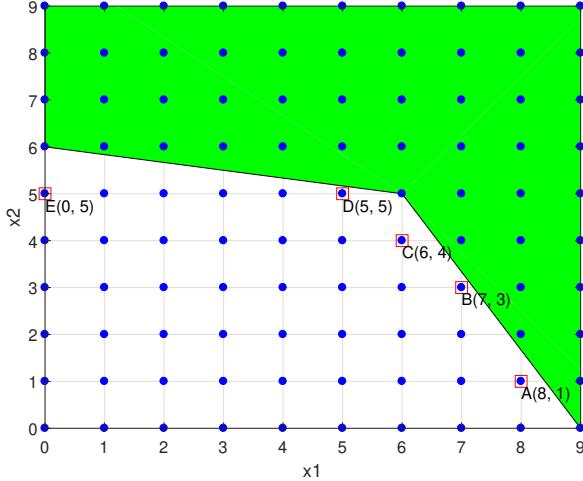


Fig. 1: The MUAs of the problem in Equation (10).

Example 1. Consider a hypothetical optimization problem formulated in (10) where \parallel denotes the logic OR operation. The decision variables are $\mathbf{X} = (x_1, x_2)$, which take values in the range $[0, 9]$. The feasibility region are formed by the disjunction of two linear constraints, which is shown in Figure 1 as the green shadowed area.

$$\begin{aligned} \min \quad & F(\mathbf{X}) = x_1 + x_2 \\ \text{s.t.} \quad & \begin{cases} x_1 + 6x_2 \geq 36 \\ 5x_1 + 3x_2 \geq 45 \\ 0 \leq x_1, x_2 \leq 9 \end{cases} \end{aligned} \quad (10)$$

In the figure, the five points $A = (8, 1)$, $B = (7, 3)$, $C = (6, 4)$, $D = (5, 5)$, $E = (0, 5)$ marked in the figure are all unschedulable assignments as they lie outside of the schedulability region. E is not an MUA, since D dominates E . Meanwhile, A , B , C and D are all MUAs since they are unschedulable assignments and there exists no other unschedulable assignment that dominates any of them.

By Property 1, any assignment dominated by an unschedulable assignment is also unschedulable. Therefore, an MUA $\mathcal{X} = (v_1, \dots, v_n)$ implies the following constraint that must be satisfied by any schedulable solution

$$\neg \left\{ \begin{array}{l} x_1 \leq v_1 \\ \dots \\ x_n \leq v_n \end{array} \right\} \Leftrightarrow \left\{ \begin{array}{l} x_1 \geq v_1 + 1 \\ \dots \\ x_n \geq v_n + 1 \end{array} \right\} \quad (11)$$

where $\{$ denotes the logic AND operation. We call (11) the **implied constraint** by \mathcal{X} . Note that no matter how complicated the schedulability analysis is, the MUA-implied constraints will always take the form in (11), hence they are an abstract interpretation of the schedulability constraints. The higher the values in \mathcal{X} , the stronger the implied constraint. In this sense, constraints implied by MUAs are the “strongest” type of constraints for schedulability. In Example 1, the implied constraints by MUAs $A = (8, 1)$, $B = (7, 3)$, $C = (6, 4)$, $D = (5, 5)$ together represent the exact schedulability region.

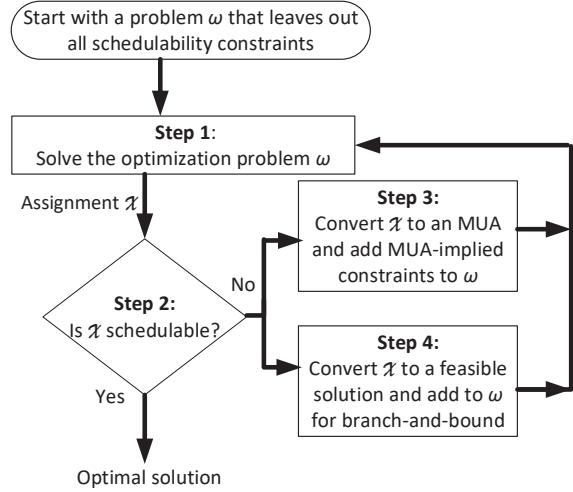


Fig. 2: MUA-guided optimization framework.

V. OPTIMIZATION FRAMEWORK

We now present the optimization framework that builds upon the concept of MUA. By the sustainability of schedulability analysis (hence Property 1), once we find an unschedulable solution, we can generalize it to MUAs to simultaneously rule out many similar unschedulable solutions. This leads to the following key idea for the design of our framework: we use MUA-implied constraints as an efficient abstraction (as opposed to a direct formulation) of the schedulability region, and employ an iterative procedure to gradually learn those MUAs that are critical for determining the optimal solution (Section V-A). Furthermore, we maintain an MUA-driven branching structure to allow incremental update of the branching tree and efficient solution to each leaf problem (Section V-B).

A. MUA-guided Iterative Procedure

In real-time systems, the complexity of the schedulability analysis may prevent us from leveraging existing optimization frameworks. For example, the most accurate schedulability analysis for AMC, AMC-max [6], hinders a possible formulation in MILP [41]. As in Equation (6), it requires to check, for each possible time instant s of criticality change, whether the corresponding response time is within the deadline. However, the range of s is unknown a priori as it depends on the task response time in LO mode.

Our key observation is that sustainability of schedulability analysis (hence Property 1) allows to generalize from one unschedulable solution to MUAs, which can simultaneously rule out many similar unschedulable solutions. This is leveraged to develop the iterative optimization framework guided by counterexamples (i.e., unschedulable solutions), as illustrated in Figure 2. Initially, it starts with an optimization problem ω that leaves out all the schedulability constraints. It then enters an iterative procedure that contains four steps. The first

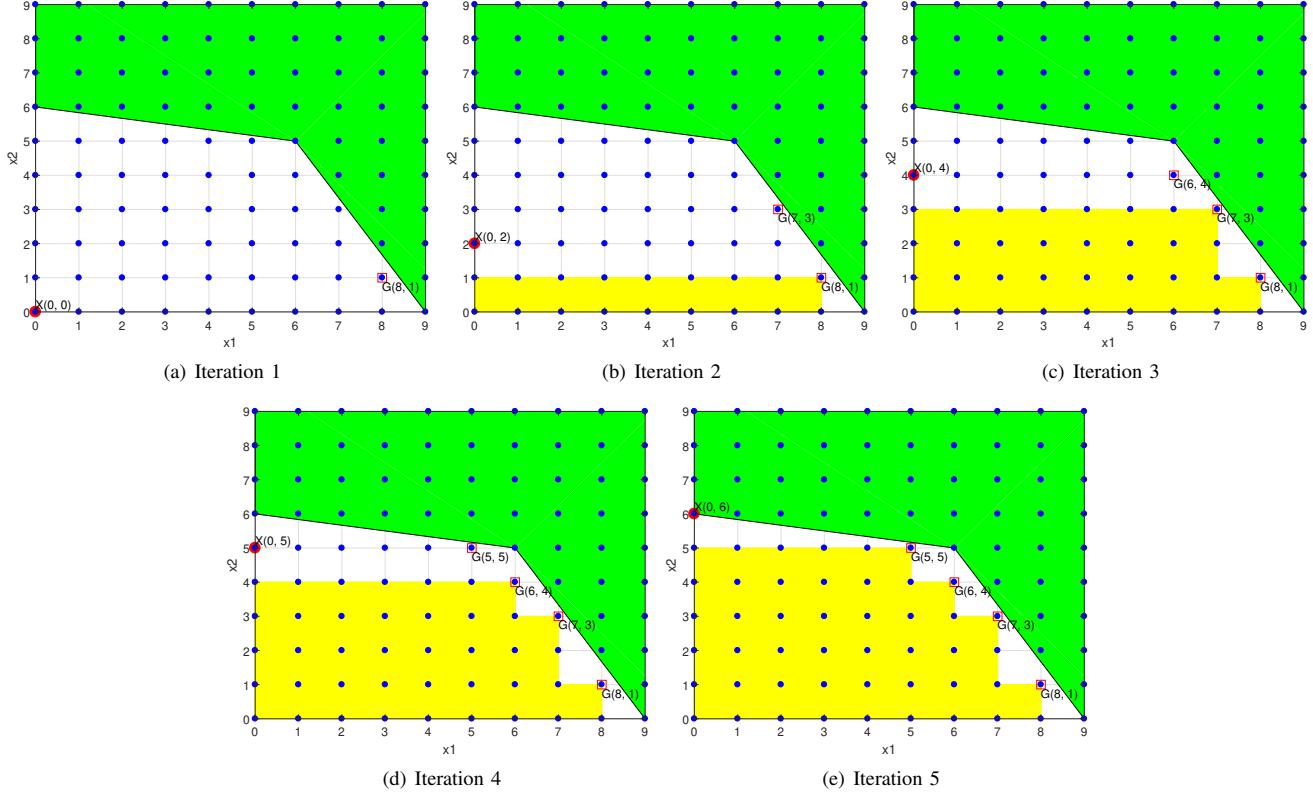


Fig. 3: The iterative procedure applied to the problem in Equation (10). The yellow shadowed area is the pruned unschedulability at the end of the iteration.

step is to solve ω using an MUA-driven branching algorithm. Here ω maintains to be a relaxed version of the original problem in Equation (1), since it only includes the implied constraints from a subset of all MUAs (and hence only part of the schedulability constraints). The second step is to use the associated (sustainable) schedulability analysis to check if the returned solution \mathcal{X} from Step 1 is schedulable or not. If yes, then \mathcal{X} must also be an optimal solution to the original problem (Theorem 1). Otherwise, it performs two operations, Steps 3 and 4, that can be executed *in parallel*. In Step 3, it converts \mathcal{X} to an MUA and adds the implied constraints (11) to ω . In this way, the counterexample (i.e., the unschedulable solution \mathcal{X} returned in Step 1) is generalized as much as possible, such that many similar unschedulable solutions can be ruled out together. In Step 4, it converts \mathcal{X} to a feasible solution and adds it to ω , which allows both efficient branch-and-bound and early termination of the algorithm. These steps will be explained from Section V-C to Section V-E, after we present the MUA-driven tree structure (Section V-B).

We now discuss two important properties of the framework, as stated in the following theorem.

Theorem 1. The algorithm in Figure 2 guarantees to terminate. If Step 1 in each iteration is solved optimally w.r.t. the added constraints, the algorithm guarantees to return an

optimal solution upon termination.

Proof. During each iteration, the procedure has to find a solution \mathcal{X} that satisfies the constraints implied by all the previously added MUAs. This means that if \mathcal{X} is still unschedulable, then it must correspond to MUAs that are different from the known ones. Since the total number of MUAs is clearly finite and bounded by $\Omega(\prod_{i=1}^n (x_i^u - x_i^l + 1))$, the algorithm guarantees to terminate.

The implied constraint (11) only cuts away unschedulable decision space. Since the problem starts with no MUAs in ω , at any point during the optimization, the feasibility region defined by the added MUA-implied constraints maintains to be an over-approximation of the exact feasibility region. Hence the optimization problem ω in Step 1 has the same objective function but a larger feasibility region than the original problem (1). This implies that upon termination, where the algorithm finds a schedulable solution, the solution must also be optimal to (1). \square

Example 2. As an example, we apply the framework to the problem (10). Here we focus on how the framework prunes the unschedulable solutions, and ignore the step of finding feasible solutions.

Iteration 1. The algorithm initially ignores all schedulability constraints. Optimizing $F(\mathbf{X})$ gives the assignment

$\mathcal{X} = (x_1, x_2) = (0, 0)$. Since \mathcal{X} is clearly unschedulable, the algorithm proceeds to convert \mathcal{X} into an MUA. Depending on the strategy for MUA computation (which will be discussed in Section V-C), the algorithm may obtain any one of points $A = (8, 1)$, $B = (7, 3)$, $C = (6, 4)$, or $D = (5, 5)$ in Figure 1. Suppose A is returned as the converted MUA.

Iteration 2. The feasibility region is updated to Figure 3(b), where the area colored in yellow represents the region cut away by the constraints implied by the MUA A . Optimizing $F(\mathbf{X})$ gives the assignment $\mathcal{X} = (0, 2)$. Since \mathcal{X} is not schedulable, the algorithm proceeds to convert \mathcal{X} into an MUA. Suppose $B = (7, 3)$ is returned as the converted MUA.

Iteration 3. The feasibility region is updated to Figure 3(c). Optimizing $F(\mathbf{X})$ returns the assignment $\mathcal{X} = (0, 4)$. Suppose $C = (6, 4)$ is returned as the computed MUA.

Iteration 4. The feasibility region is updated to Figure 3(d). Optimizing $F(\mathbf{X})$ returns the assignment $\mathcal{X} = (0, 5)$. Now the only possible MUA that can be computed is $D = (5, 5)$.

Iteration 5. The feasibility region is updated to Figure 3(e). Optimizing $F(\mathbf{X})$ gives the assignment $\mathcal{X} = (0, 6)$. This assignment is now schedulable and the algorithm terminates with an optimal solution $(x_1, x_2) = (0, 6)$.

B. MUA-driven Branching Tree

As in Figure 2, the above procedure requires to solve an instance of ω during each iteration, which may only be slightly different from the previous iterations since the procedure only adds a handful of new MUAs to ω . Directly calling mathematical programming solvers to solve ω is not necessarily efficient since many solvers do not support incremental solving, i.e., they have to solve each instance of ω from scratch.²

Hence we build a branching tree structure to represent the MUA-implied constraints. It allows incremental updating of the tree, as well as efficiently solving each subproblem at the leaf nodes. Initially (i.e., before entering the iteration in Figure 2) the tree only contains one (root) node. Each time a new MUA \mathcal{X} is added to the problem ω , we add a new layer in the tree to represent the disjunction (i.e., logic OR) of the constraints implied by \mathcal{X} , where each branch (edge) in the new layer is a constraint in the disjunction. Each node \mathcal{N} in the tree represents the set of conjunctive (i.e., logic AND) constraints along the path from the root node to this node \mathcal{N} . Since the constraint on each branch takes the particular form $x_j \geq u_j$ for some value u_j , the constraint represented by a node \mathcal{N} can be simplified as

$$\begin{cases} x_1 \geq u_1 \\ \dots \\ x_n \geq u_n \end{cases} \quad (12)$$

For simplicity, we denote the node as $\mathcal{N} = [u_1, \dots, u_n]$.

²The only known exception is the MILP solver CPLEX [21], which provides an interface for building customized branching tree. But the price is that it can no longer run in parallel on multiple cores, which actually makes the whole procedure slower.

Some nodes in the tree are redundant, in the sense that the corresponding constraints may have already been satisfied by constraints along the path from the root. These nodes can be pruned to make the tree structure more compact.

Example 3. Consider the branching tree after Iteration 2 in Example 2, i.e., after adding the constraints implied by MUAs $A = (8, 1)$ and $B = (7, 3)$. The left hand side of Figure 4 shows the resulted tree. The constraint corresponding to A is $x_1 \geq 9 \vee x_2 \geq 2$, and the constraint corresponding to B is $x_1 \geq 8 \vee x_2 \geq 4$, where \vee denotes the logic OR operation. Hence, the root node (node 1) is first branched to two children, nodes 2 and 3, where the branch to node 2 represents the constraint $x_1 \geq 9$, and the branch to node 3 represents the constraint $x_2 \geq 2$. When adding the constraints for MUA $B = (7, 3)$, each of nodes 2 and 3 is branched to two children. Node 6, for example, represents the constraints along the path from the root node, i.e., $x_2 \geq 2 \wedge x_1 \geq 8$, where \wedge denotes the logic AND operation.

When adding the constraint $x_1 \geq 8 \vee x_2 \geq 4$ implied by MUA B to node 2, this constraint is already satisfied by the constraint $x_1 \geq 9$ represented by node 2, i.e., $(x_1 \geq 9) \wedge (x_1 \geq 8 \vee x_2 \geq 4) = x_1 \geq 9$. The right hand side of Figure 4 shows the tree structure after pruning the redundant nodes.

The special form of MUA-implied constraints make it easy to check the redundancy of nodes.

Theorem 2. Assume $\mathcal{N} = [u_1, \dots, u_n]$ is an existing node in the tree. If an MUA $\mathcal{X} = (v_1, \dots, v_n)$ to be added satisfies that $\exists i : u_i \geq v_i + 1$, then the constraints implied by \mathcal{X} are redundant.

Proof. This is easy to see by checking the represented constraints of \mathcal{N} in (12) and those of MUA \mathcal{X} in (11).

$$\begin{array}{l} \left\{ \begin{array}{l} x_1 \geq u_1 \\ \dots \\ x_n \geq u_n \end{array} \right. \Rightarrow x_i \geq u_i \Rightarrow x_i \geq v_i + 1 \Rightarrow \left\{ \begin{array}{l} x_1 \geq v_1 + 1 \\ \dots \\ x_n \geq v_n + 1 \end{array} \right. \end{array} \quad \square$$

In the following we explain how the steps in Figure 2 are handled. For Step 2, it can use any schedulability analysis as long as it is sustainable.

C. Step 1: Solving ω

For Step 1, it requires to solve the following optimization problem at each leaf node $\mathcal{N} = [u_1, \dots, u_n]$ in the tree:

$$\begin{array}{ll} \min_{x_j \in \mathbf{X}} & F(\mathbf{X}) \\ \text{s.t.} & \text{constraints of form (12)} \\ & x_j \in [x_j^l, x_j^u], \forall j = 1, \dots, n \end{array} \quad (13)$$

The constraints in (13) are of a particularly simple form: there are no coupled constraints for any pair of variables x_i and x_j . This means that the overall optimization can be done

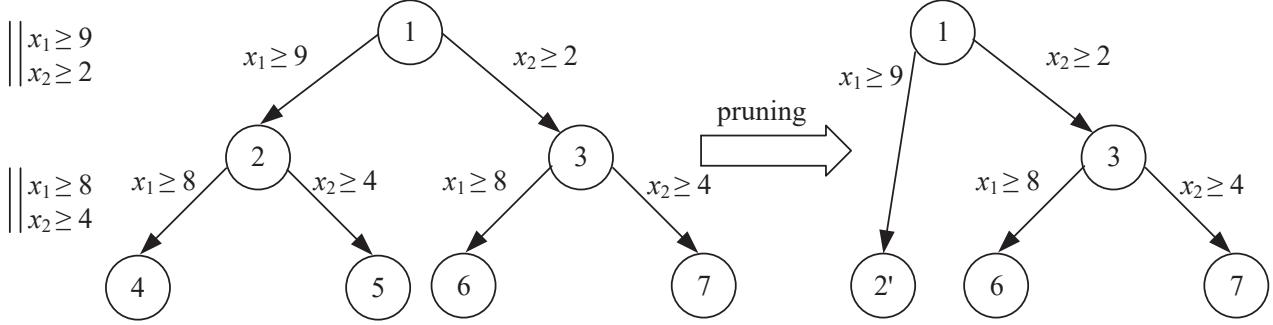


Fig. 4: Branching tree at the end of Iteration 2 of Example 2: original tree (left); after pruning (right).

by optimizing each variable in sequence [9]. In other words, (13) can be transformed to

$$\begin{aligned} & \min_{x_1} \min_{x_2} \dots \min_{x_n} F(\mathbf{X}) \\ \text{s.t.} & \text{constraints of form (12)} \\ & x_j \in [x_j^l, x_j^u], \forall j = 1, \dots, n \end{aligned} \quad (14)$$

For example, if $F(\mathbf{X})$ is convex with respect to the variables, then solving (14) amounts to solving a series of single-variable convex programs, which is a lot easier than those with coupled constraints.

In addition, even if $F(\mathbf{X})$ is not convex (hence in general difficult to solve [9]), we may still be able to provide a simple solution to it. In real-time systems, often times there exists some tradeoff between schedulability and the metrics in the objective function. For example, when the CPU frequency is increased, the task WCET is lower, but the energy consumption will be higher. Similarly, increasing the task periods will make the task system easier to schedule, but the control quality and stability will be worsened. In component-based design, modularity (i.e., the number of exposed interfaces) is a critical metric, but better modularity may lead to larger code size and longer WCET [33]. We leverage this observation to derive the following corollary, which is a direct consequence from Equation (14). In fact, Corollary 3 applies to both problems in the experiments.

Corollary 3. If $F(\mathbf{X})$ is monotonically non-decreasing with respect to each of the variables in \mathbf{X} , then the optimal solution to (13) is $x_1 = u_1, \dots, x_n = u_n$.

This corollary is intuitive since the objective function will push each variable x_i to its lowest possible value u_i .

D. Step 3: Converting an Unschedulable Assignment to MUA

In this section, we discuss the algorithm for converting an unschedulable assignment \mathcal{X} into an MUA \mathcal{U} . It utilizes Property 1, i.e., the schedulability analysis is sustainable with respect to the variables, to maximally increase each entry in \mathcal{X} while maintaining its unschedulability. The procedure is summarized in Algorithm 1. It uses binary search to sequentially find the maximal value that each entry in \mathcal{U} can be increased to while maintaining unschedulability. The

algorithm requires $O(n \log x^{\max})$ number of schedulability analysis where $x^{\max} = \max\{x_1^u, \dots, x_n^u\}$.

Algorithm 1 Conversion to MUA

```

1: function CONVERTTOMUA(Unschedulable Assignment  $\mathcal{X}$ )
2:    $\mathcal{U} = (u_1, \dots, u_n) = \mathcal{X}$ 
3:   for each entry  $u_i$  in  $\mathcal{U}$  do
4:     Use binary search to maximally increase  $u_i$  while
       keeping the system unschedulable
5:   end for
6:   return  $\mathcal{U}$ 
7: end function

```

This step is critical in the efficiency of the overall algorithm. Unlike typical counterexample guided algorithms, once we find an unschedulable solution, in the next iteration we rule out not only this solution but also many similar ones. Here the concept of MUA is critical: it is essentially a generalization from one unschedulable solution to many, which is a key in allowing a fast convergence rate of the framework.

Example 4. Using Algorithm 1 on Example 1 yields the exact trace of iterations shown in Figure 3 in Example 2. For instance, in the first iteration, when a solution $\mathcal{X} = (0, 0)$ is returned, Algorithm 1 first increases u_1 to 8 since that is the largest value of x_1 that still makes the system unschedulable. It then increases u_2 from 0 to 1 to get the MUA $\mathcal{U} = (8, 1)$.

E. Step 4: Finding Feasible Solutions

The branching tree structure (as in Figure 4) allows to implement a typical branch-and-bound algorithm, i.e., to use the known best solution \mathcal{X} to cut the branches that are certainly no better than \mathcal{X} and hence are surely suboptimal. Getting a feasible solution also allows returning useful results for designers even if they are not necessarily the global optimal solution. This facilitates the possible early termination of the overall procedure, otherwise it will not be able to get any feasible solution until the optimal solution is found.

Our main idea is to make use of the assignment returned in Step 1 and convert it heuristically into a good-quality schedulable assignment. The conversion can be performed concurrently

with the rest of the optimization process (Steps 1 and 3). Specifically, in Step 2, if the assignment $\mathcal{X} = (v_1, \dots, v_n)$ is unschedulable, we simply scale each variable x_i by a factor $a \in [0, 1]$, until \mathcal{X} becomes schedulable

$$x_i = v_i + a(x_i^u - v_i) \quad (15)$$

When $a = 1$, x_i takes the upper-bound value x_i^u . We use binary search to find the minimum a such that the assignment \mathcal{X} becomes schedulable after scaling.

VI. APPLICABILITY AND EFFICIENCY

In this section, we discuss the applicability and expected efficiency of the proposed techniques.

Applicability to Optimizing Priority Assignment. As mentioned earlier, the framework is applicable to optimizing task response time, period, WCETs or deadline, as long as the schedulability analysis is sustainable with respect to these variables. Regarding task priority assignment as part of the decision variables, this comes with two cases: (a) the objective function is independent from the task priorities; (b) it is sensitive to task priority assignments, such as the memory consumption in the software implementation that preserves the semantics of the synchronous reactive models [36].

For case (a), we can leverage Audsley's algorithm [1] that is applicable to many task models and scheduling schemes [12]: if there exists a schedulable priority assignment, Audsley's algorithm will be able to find it. Hence, we can leave out the variables of priority assignment in the problem ω , but instead incorporate the optimization of priority assignment in the procedure for checking schedulability (e.g., Step 2 in Figure 2, and Line 4 in Algorithm 1): After fixing the values of all other variables, the existence of a schedulable priority assignment now can be efficiently checked by Audsley's algorithm.

For case (b), it is necessary to explicitly include those binary variables for task priority orders in the optimization problem ω in the framework of Figure 2. In this case, we can leverage the concept of unschedulability core [40]. Intuitively, it is an irreducible representation of the reason why a given total priority order is unschedulable, in the sense that relaxing any order will make the system schedulable. We leave the details of this discussion to future work.

Efficiency. The proposed technique fits the best for problems that have the following characteristics:

- 1) The schedulability analysis is complex.
- 2) The rest of the problem is relatively simple.

For the first characteristic, the exact schedulability analysis is often NP-complete, even for the basic settings, e.g., periodic task with fixed priority scheduling [15], or EDF scheduling with arbitrary deadlines [14]. In this case, even if the problem may be formulated in some standard mathematical programming framework, our approach might still be better, since it uses MUAs to abstract away the details of schedulability analysis, and only performs such an analysis on a small number of design choices guided by the objective function. Of course, there are cases that the schedulability analysis is particularly simple, such as the condition for tasks with

implicit deadlines scheduled by EDF [22], or the utilization based bound for EDF-VD [20]. In this case, although our framework is still applicable, it is faster to directly handle the schedulability condition.

For the second characteristic, this is satisfied when the objective function has some friendly properties (such as monotonicity or convexity with respect to the variables), and the constraints only include the system schedulability. If the problem includes some additional constraints other than schedulability, then the problem at each leaf node of the branching tree (see Figure 4) is not necessarily easy to solve, since there might be coupled constraints among the variables. In this case, it can be more efficient to use other appropriate solvers to directly solve ω at Step 1 of the framework.

VII. EXPERIMENT RESULT

We now use two problems to demonstrate the advantage of our framework. The first is the minimization of energy consumption, the second is the optimization of control quality.

A. Optimizing Energy Consumption with DVFS

In this experiment, we consider the energy consumption model in Equation (8). We use both random systems as well as an industrial case study to compare different approaches. To demonstrate that our framework is applicable to various scheduling models and schedulability analysis techniques, we assume the periodic task model and Equation (2) in the experiments on random systems, and use AMC-rtb and AMC-max (Equations (3)–(6)) in the experiments on the industrial case study. The relative simplicity of Equation (2) compared to AMC-rtb and AMC-max also allows us to perform experiments on a large number of random systems.

Random Systems. For random systems, we compare the following methods:

- **MIGP:** A mixed-integer geometric programming formulation, solved by the geometric programming solver gpposy [25] with the BnB (bmi) solver in YALMIP [23].
- **MUA-MILP:** The proposed iterative procedure in Figure 2, but the subproblem ω with MUA-implied constraints is formulated as an MILP and solved using CPLEX [21].
- **MUA-incremental:** The proposed technique depicted in Figure 2 with MUA-driven branching tree for incremental update, where each problem at the leaf node is solved using Corollary 3.
- **Minimum-single-speed:** A simple heuristic that uses binary search to find the minimum single speed at which all tasks become schedulable.

To avoid excessive waiting, in MUA-incremental we set a limit for the number of nodes in each iteration (i.e., the size of each layer in the branching tree) to be $K = 10000$. Also, the time limits of MUA-MILP and MUA-incremental are set to 600 seconds for each task system. The BnB (bmi) solver in YALMIP does not have a time limit setting and only allows to set a limit on the number of iterations. Therefore, we set a

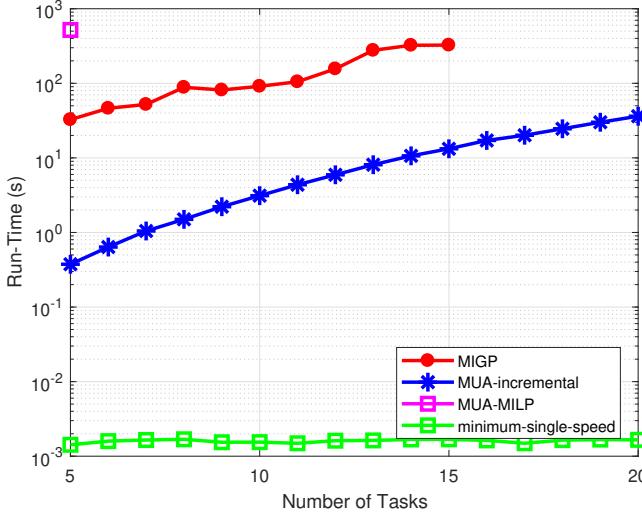


Fig. 5: Runtime of minimizing energy consumption on random systems.

maximum iteration limit of 2000. This gives roughly a similar time limit for MIGP as those of the other methods.

The task sets are generated synthetically as follows. For each task set, we first randomly select a system utilization in the range $[0.5, 0.9]$. We then generate a period T_i for each task according to log-uniform distribution in the range $[100, 100000]$, and a utilization U_i for each task using UUni-fast algorithm [11]. The corresponding WCET $C_i^b = T_i \cdot U_i$ is treated as the execution time at the base clock rate. The range of decision variable C_i is taken as $[C_i^b, 2C_i^b]$. This means that the clock rate can be decreased as low as half the base clock rate. The deadline D_i of each task τ_i is generated randomly in the range $[C_i^b, T_i]$. Priorities are assigned according to the deadline monotonic policy.

Figure 5 illustrates the average runtime over 1000 random systems for each m , the number of tasks in the system. The runtime of Minimum-single-speed is very short (a few milliseconds), as it is simply a binary search on a single period value. MUA-incremental is about one to two orders of magnitude faster than MIGP. The capping of MIGP that occurs for systems with 14 or 15 tasks is mainly due to a large number of cases reaching the iteration limit. Meanwhile, MUA-incremental is able to finish all the instances in the time limit. As for MUA-MILP, it is very slow such that even for systems with 5 tasks, most of the cases are timed out. This is because the MILP solver CPLEX is unable to perform efficient incremental solving of the problems. Again, it demonstrates the benefit of designing a branching algorithm that takes advantage of the MUA-guided framework, as in the case of MUA-incremental.

Since MUA-MILP is unable to finish for most of the cases, we only compare the quality of solutions from MUA-incremental, MIGP and Minimum-single-speed in terms of relative gap. For each random system, we define the relative gap of MIGP (or Minimum-single-speed) with respect to

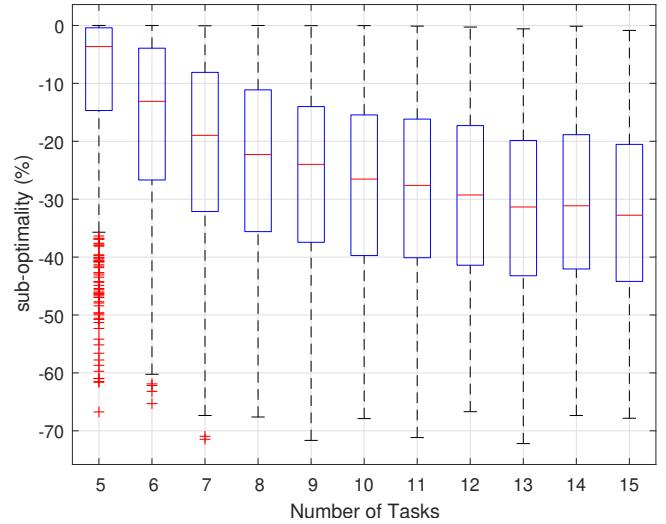


Fig. 6: Relative gap of MIGP compared to MUA-incremental.

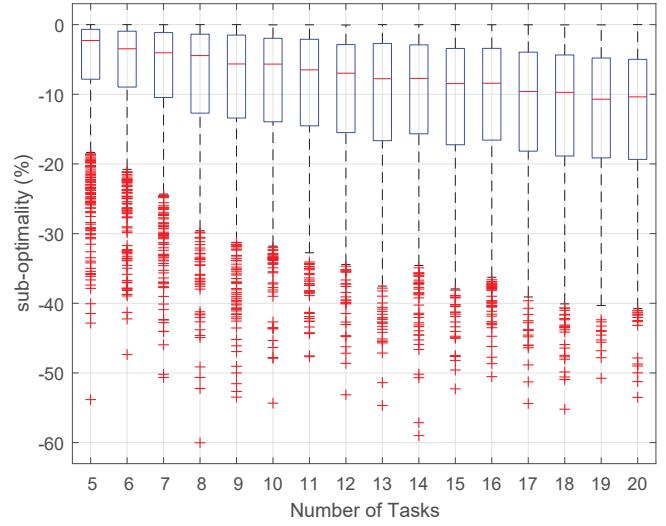


Fig. 7: Relative gap of Minimum-single-speed compared to MUA-incremental.

MUA-incremental as the relative difference in the objective values. Since MUA-incremental always provides a better objective value than the other two, we define the relative gap of MIGP (resp. Minimum-single-speed) as

$$s = \frac{p_A - p_B}{p_B} \times 100\% \quad (16)$$

where p_A is the objective value of MUA-incremental, and p_B represents that of MIGP (resp. Minimum-single-speed).

Figure 6 shows the whisker box plot of the distribution of the relative gap for MIGP compared to MUA-incremental. On average MUA-incremental finds 3% to 30% better solutions (i.e., with less energy consumption) than MIGP within the time limit. Likewise, Figure 7 shows the distribution for

TABLE I: Flight Management System case study [19]

τ	τ_1	τ_2	τ_3	τ_4	τ_5	τ_6
T	5000	200	1000	1600	100	1000
$C(LO)$	[5, 40]	[5, 40]	[5, 40]	[5, 40]	[5, 40]	[5, 40]
HI/LO	HI	HI	HI	HI	HI	HI
τ	τ_7	τ_8	τ_9	τ_{10}	τ_{11}	
T	1000	1000	1000	1000	1000	
$C(LO)$	[5, 40]	[50, 400]	[50, 400]	[50, 400]	[50, 400]	
HI/LO	HI	LO	LO	LO	LO	

Minimum-single-speed. As in the figure, MUA-incremental is on average 3% to 10% better than Minimum-single-speed.

Flight Management System. We next evaluate the techniques on an avionics case study consisting of a subset of Flight Management System application [19]. The system contains 11 tasks of different criticality, which implement functions such as localization and flight planning. Each task is abstracted into an implicit deadline sporadic task characterized by a minimum inter-arrival time, a range of execution time that is typical in practice, and a criticality level. 7 tasks are of HI-criticality and the other 4 are of LO-criticality. The parameter configuration of the case study is summarized in Table I.

We consider fixed-priority uniprocessor scheduling according to Adaptive-Mixed-Criticality (AMC) for these tasks. For schedulability analysis, AMC-max has much higher computational complexity comparing to AMC-rtb, but it is more accurate and may help find better quality solutions when used in optimization. In the following, we consider the problem of minimizing LO-criticality energy consumption [19] given by Equation (8). The range of the LO-criticality WCET $C_i(LO)$ of each task τ_i is determined as follows. We first take the upper-bound C_i^u of the execution time range given in the case study. Then we consider $C_i(LO)$ to be freely adjustable in the interval $[\frac{C_i^u}{4}, 2C_i^u]$ by CPU clock rate adjustment. For HI-criticality task, $C_i(HI)$ is obtained by scaling $C_i(LO)$ by a fixed criticality factor γ , i.e. $C_i(HI) = \gamma C_i(LO)$. In the experiments, we vary γ to take three possible values 3, 4, 5. The optimization problem is to find a $C_i(LO)$ for each task that minimizes Equation (8).

We compare the following four methods:

- **MIGP-AMCrtb:** Formulating the problem with AMC-rtb analysis as a mixed integer geometric program, and use the gposy and YALMIP solvers to solve it.
- **MUA-MILP-AMCrtb:** The iterative framework in Figure 2, where the problem ω is directly solved as an MILP program, and the schedulability analysis uses AMC-rtb.
- **MUA-incremental-AMCrtb:** The framework in Figure 2, where problems ω are solved with the MUA-driven branching tree, and the schedulability test is AMC-rtb.
- **MUA-incremental-AMCmax:** The same as MUA-incremental-AMCrtb, except that the schedulability analysis uses AMC-max.

We let each approach run for 48 hours. The size limits for MUA-incremental-AMCrtb and MUA-incremental-AMCmax are both set to $K = 50000$. Note that we do not consider AMC-max for MIGP as the analysis is too complicated to formulate in the MIGP framework. Likewise,

we do not apply AMC-max to MUA-MILP, since it is already very slow under AMC-rtb: it cannot find any feasible solutions in 48 hours.

We first assume priority assignment is given by rate monotonic (ties are broken by criticality level). The results are summarized in Table II. For $\gamma = 3$, MIGP-AMCrtb gets stuck in a suboptimal solution (about 51% worse than the optimal) early and cannot find any better solution even after 48 hours. MUA-MILP-AMCrtb is even worse, as it cannot find any feasible solution in 48 hours. Comparably, MUA-incremental performs much better: regardless of whether AMC-max or AMC-rtb is used, it finds the best solution in about 15 minutes, or about 200 times faster than MUA-MILP-AMCrtb and MUA-MILP-AMCrtb.

The cases of $\gamma = 4$ and $\gamma = 5$ are similar. Note that the use of the more accurate AMC-max analysis did not improve the quality of the solution. This is mainly due to the rate-monotonic priority assignment. For this case, LO-criticality tasks are mostly lower in priorities than HI-criticality tasks (except τ_1, τ_4 which have quite loose deadlines). As a result, most HI-criticality tasks only suffer interference from other HI-criticality tasks. Therefore, the worst-case scenario for the response time calculation occurs when the system is entirely in the HI-criticality mode, where AMC-max and AMC-rtb give the same response time. On the other hand, the small difference in the runtimes of MUA-incremental-AMCmax and MUA-incremental-AMCrtb demonstrates that the efficiency of our approach is relatively insensitive to the complexity of the schedulability analysis.

We next consider the setting where priority assignment is not given and need to be co-optimized with WCETs. We omit MIGP-AMCrtb from this experiment as it is no longer applicable. Note that for methods based on the proposed optimization framework (MUA-incremental-AMCrtb and MUA-incremental-AMCmax), including priority assignment in the design space is rather simple: since both AMC-rtb and AMC-max are both compatible with Audsley's algorithm [16], in the framework we just use a schedulability analysis procedure that incorporates both the response time calculation (either AMC-rtb or AMC-max) and Audsley's algorithm for finding a feasible priority assignment.

The results are summarized in Table III. A number of observations can be made. First, much better solutions are found when priority assignment is treated as decision variables. The overall objective values reduce by as much as more than half comparing to the results in Table II. This shows the benefit of the proposed optimization framework: it is able to handle the co-optimization of various variables. Second, the difference between AMC-rtb and AMC-max analyses is now more noticeable. For example, when the criticality factor $\gamma = 5$, the use of AMC-max provides a solution that is 23% better than that of AMC-rtb. This demonstrates the benefit of using a more accurate analysis for optimization. However, such a benefit can only be achieved when the optimization algorithm is capable of accommodating the analysis. This is difficult for MIGP as AMC-max is too complicated to formulate

TABLE II: Results on Flight Management System by Optimizing Task WCETs

Method	$\gamma = 3$		$\gamma = 4$		$\gamma = 5$	
	Time	Objective	Time	Objective	Time	Objective
MUA-MILP-AMCrtb	$\geq 48h$	N/A	$\geq 48h$	N/A	$\geq 48h$	N/A
MIGP-AMCrtb	$\geq 48h$	6.216e+007	$\geq 48h$	8.379e+007	$\geq 48h$	1.076e+008
MUA-incremental-AMCrtb	873.91s	4.108e+007	582.06s	5.803e+007	409.76s	7.723e+007
MUA-incremental-AMCmax	903.10s	4.108e+007	619.86s	5.803e+007	406.59s	7.723e+007

TABLE III: Results on Flight Management System by Co-optimizing Task priority assignments and WCETs

Method	$\gamma = 3$		$\gamma = 4$		$\gamma = 5$	
	Time	Objective	Time	Objective	Time	Objective
MUA-MILP-AMCrtb	$\geq 48h$	N/A	$\geq 48h$	N/A	$\geq 48h$	N/A
MUA-incremental-AMCrtb	2356.61s	2.626e+007	1204.74s	2.832e+007	517.68s	3.415e+007
MUA-incremental-AMCmax	2734.41s	2.626e+007	2260.71s	2.626e+007	2569.84s	2.626e+007

in geometric programming framework. Our framework, on the other hand, can incorporate any schedulability analysis as long as it is sustainable. Third, the runtime noticeably increases comparing to Table II. This is mainly because the schedulability analysis now incorporates Audsley's algorithm, which has a substantially higher computation complexity than just a plain response time calculation.

B. Control Performance

In this experiment, we consider the problem of optimizing control performance for a set of periodic tasks scheduled on a uniprocessor. The problem was originally introduced in [24], where the objective is formulated in Equation (7).

Mancuso et al. [24] proposed to relax the response time analysis (2) by removing the ceiling operator (hence it uses $\frac{R_i}{T_j}$ to estimate the number of interferences from task τ_j on task τ_i). Although the relaxation was claimed to be a close approximation to the exact response time in practice, it is possible that the relaxed analysis may give unsafe solutions. In fact, in the randomly generated systems below, the solution returned by [24] is always unschedulable with the exact analysis in Equation (2). Thus, in the comparison of other methods below, we use the exact analysis (2).

We evaluate on randomly generated synthetic task sets. The parameters are generated using a similar setting as [24]. Specifically, in Equation (7), α_i is randomly generated in the range $[1, 1000]$, β_i is randomly generated in the range $[1, 10000]$, and the WCET of each task C_i is randomly selected from $[1, 100]$. The upper-bound for the task period T_i is set to be 5 times the sum of all task WCETs, i.e., $T_i^u = 5 \sum_{j=1}^n C_j$. This sets 20% as the lower bound on system utilization. Task priority are assigned with the rate monotonic policy.

We compare the following three methods:

- **MIGP:** MIGP formulation proposed in [10] for optimizing task periods. We use geometric programming solver gposy [25] with the BnB (bmi) solver in YALMIP [23] for solving MIGP problems.
- **MUPDA-MILP:** The proposed optimization framework in [39], which is an iterative procedure that leverages CPLEX solver [21] to solve a series of MILP problems.

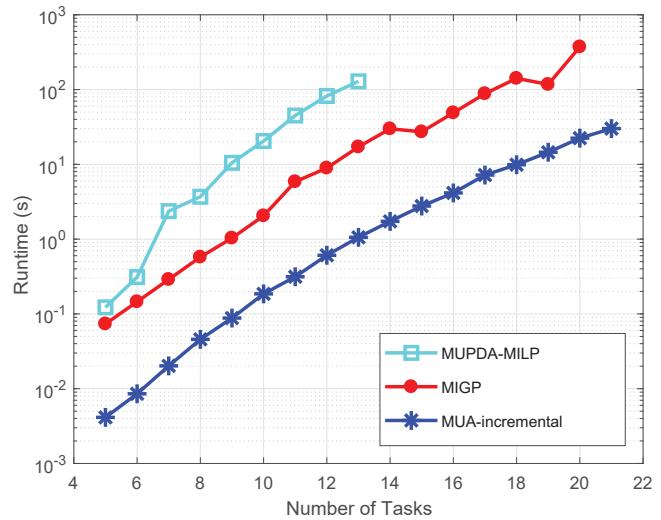


Fig. 8: Runtime of optimizing control performance.

- **MUA-incremental:** The proposed optimization framework in Figure 2, where the problem ω is solved using the MUA-driven branching tree.

Like in Section VII-A, we set a maximum iteration limit of 2000 for MIGP, and the time limits of the other two methods are set to 600 seconds for each problem instance. The size limit for MUA-incremental is set to $K = 10000$.

Figure 8 plots the runtime of all the optimization methods MIGP, MUPDA-MILP and MUA-incremental. While all methods give the same optimal solution when finishing within the time limit, MUPDA-MILP has the worst scalability. As discussed in [39], this confirms that it is often inefficient in solving problems with objective that are sensitive to many decision variables, since it takes many iterations to terminate, and each MILP problem has to be solved from scratch (instead of incrementally as in our approach). Meanwhile, MUA-incremental runs about 10 times faster than MIGP. This again demonstrates the advantage of our proposed framework that judiciously combines the MUA-guided iterative procedure with the incremental update of the branching tree.

VIII. CONCLUSION

In this paper, we propose a framework for optimizing the design of real-time system with sustainable schedulability analysis. We propose the concept of Maximal-Unschedulable Assignment (MUA) and show how it can be used to abstractly interpret the schedulability constraints. Based on the concept, we develop an iterative optimization procedure that uses MUAs to iteratively refine the schedulability region. It contains three key steps: i) an algorithm for solving the optimization problem consisting of MUA-implied constraints, ii) an algorithm for computing MUAs that leads to faster convergence, and iii) use of MUAs to exploring good-quality schedulable solutions. We perform experiments on two optimization problems to demonstrate the advantage of the proposed approach in applicability and scalability.

ACKNOWLEDGMENT

This work is partially funded by NSF Grants No. 1812963 and No. 1837519.

REFERENCES

- [1] N. Audsley. On priority assignment in fixed priority scheduling. *Information Processing Letters*, 79(1):39 – 44, 2001.
- [2] N. C. Audsley, A. Burns, M. F. Richardson, and A. J. Wellings. in proc. ieee workshop on real-time operating systems and software. In *real-time scheduling: the deadline-monotonic approach*, pages 133–137, 1991.
- [3] T. Baker and S. Baruah. Sustainable multiprocessor scheduling of sporadic task systems. In *Euromicro Conference on Real-Time Systems*, 2009.
- [4] M. Bambagini, M. Marinoni, H. Aydin, and G. Buttazzo. Energy-aware scheduling for real-time systems: A survey. *ACM Transactions on Embedded Computing Systems (TECS)*, 15(1):1–34, 2016.
- [5] S. Baruah and A. Burns. Sustainable scheduling analysis. In *27th IEEE Real-Time Systems Symposium*, 2006.
- [6] S. Baruah, A. Burns, and R. Davis. Response-time analysis for mixed criticality systems. In *IEEE Real-Time Systems Symposium*, 2011.
- [7] I. Bate and P. Emberson. Incorporating scenarios and heuristics to improve flexibility in real-time embedded systems. In *IEEE Real-Time and Embedded Technology and Applications Symposium*, 2006.
- [8] E. Bini and M. Di Natale. Optimal task rate selection in fixed priority systems. In *26th IEEE Real-Time Systems Symposium*, 2005.
- [9] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [10] A. Davare, Q. Zhu, M. Di Natale, C. Pinello, S. Kanajan, and A. Sangiovanni-Vincentelli. Period optimization for hard real-time distributed automotive systems. In *ACM/IEEE Design Automation Conference*, 2007.
- [11] R. Davis and A. Burns. Priority assignment for global fixed priority preemptive scheduling in multiprocessor real-time systems. In *30th IEEE Real-Time Systems Symposium*, 2009.
- [12] R. I. Davis, L. Cucu-Grosjean, M. Bertogna, and A. Burns. A review of priority assignment in real-time systems. *Journal of systems architecture*, 65:64–82, 2016.
- [13] C. Ebert and J. Favaro. Automotive software. *IEEE Software*, 34(3):33–39, 2017.
- [14] P. Ekberg and W. Yi. Uniprocessor feasibility of sporadic tasks with constrained deadlines is strongly comp-complete. In *27th Euromicro Conference on Real-Time Systems*, pages 281–286, 2015.
- [15] P. Ekberg and W. Yi. Fixed-priority schedulability of sporadic tasks on uniprocessors is np-hard. In *IEEE Real-Time Systems Symposium*, 2017.
- [16] Z. Guo, S. Sruti, B. C. Ward, and S. Baruah. Sustainability in mixed-criticality scheduling. In *IEEE Real-Time Systems Symposium*, 2017.
- [17] A. Hamann, M. Jersak, K. Richter, and R. Ernst. Design space exploration and system optimization with symta/s - symbolic timing analysis for systems. In *IEEE Real-Time Systems Symposium*, 2004.
- [18] M. Hassanalian and A. Abdelkefi. Classifications, applications, and design challenges of drones: A review. *Progress in Aerospace Sciences*, 91:99–131, 2017.
- [19] P. Huang, G. Giannopoulou, N. Stoimenov, and L. Thiele. Service adaptions for mixed-criticality systems. In *IEEE Asia and South Pacific Design Automation Conference*, 2014.
- [20] P. Huang, P. Kumar, G. Giannopoulou, and L. Thiele. Energy efficient dvfs scheduling for mixed-criticality systems. In *ACM Conference on Embedded Software*, 2014.
- [21] International Business Machines Corporation. *CPLEX Optimizer*.
- [22] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1):46–61, Jan. 1973.
- [23] J. Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *IEEE Symposium on Computer Aided Control Systems Design*, 2004.
- [24] G. Mancuso, E. Bini, and G. Pannocchia. Optimal priority assignment to control tasks. *ACM Trans. Embedded Computing Systems*, 13(5s):161, 2014.
- [25] A. Mutapcic, K. Koh, S. Kim, and S. Boyd. Ggplab version 1.00: a matlab toolbox for geometric programming, January 2006.
- [26] M. D. Natale, L. Guo, H. Zeng, and A. Sangiovanni-Vincentelli. Synthesis of multi-task implementations of simulink models with minimum delays. *IEEE Trans. Industrial Informatics*, 6(4):637–651, 2010.
- [27] A. Nelson, O. Moreira, A. Molnos, S. Stuijk, B. T. Nguyen, and K. Goossens. Power minimisation for real-time dataflow applications. In *Euromicro Conference on Digital System Design*, 2011.
- [28] S. Pagani and J.-J. Chen. Energy efficiency analysis for the single frequency approximation (sfa) scheme. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(5s):158, 2014.
- [29] M. Saksena and Y. Wang. Scalable real-time system design using preemption thresholds. In *IEEE Real-Time Systems Symposium*, 2000.
- [30] D. Seto, J. P. Lehoczky, and L. Sha. Task period selection and schedulability in real-time systems. In *19th IEEE Real-Time Systems Symposium*, 1998.
- [31] M. Shin and M. Sunwoo. Optimal period and priority assignment for a networked control system scheduled by a fixed priority scheduling system. *International Journal of Automotive Technology*, 8:39–48, 2007.
- [32] K. Tindell, A. Burns, and A. Wellings. Allocating hard real-time tasks: An np-hard problem made easy. *Real-Time Syst.*, 4(2):145–165, 1992.
- [33] S. Tripakis and R. Lublinerman. Modular code generation from synchronous block diagrams: Interfaces, abstraction, compositionality. In *Principles of Modeling*, pages 449–477. Springer, 2018.
- [34] C. Wang, Z. Gu, and H. Zeng. Global fixed priority scheduling with preemption threshold: Schedulability analysis and stack size minimization. *IEEE Trans. Parallel and Distributed Systems*, 27(11):3242–3255, Nov. 2016.
- [35] Y. Wang and M. Saksena. Scheduling fixed-priority tasks with preemption threshold. In *International Conference on Real-Time Computing Systems and Applications*, 1999.
- [36] H. Zeng and M. Di Natale. Mechanisms for guaranteeing data consistency and flow preservation in autosar software on multi-core platforms. In *IEEE Symposium on Industrial and Embedded Systems*, 2011.
- [37] H. Zeng and M. Di Natale. Efficient implementation of autosar components with minimal memory usage. In *7th IEEE Symposium on Industrial Embedded Systems*, 2012.
- [38] Q. Zhao, Z. Gu, and H. Zeng. Design optimization for autosar models with preemption thresholds and mixed-criticality scheduling. *Journal of Systems Architecture*, 72:61–68, 2017.
- [39] Y. Zhao, V. Gala, and H. Zeng. A unified framework for period and priority optimization in distributed hard real-time systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37(11):2188–2199, 2018.
- [40] Y. Zhao and H. Zeng. The concept of unschedulability core for optimizing priority assignment in real-time systems. In *Conference on Design, Automation and Test in Europe*, 2017.
- [41] Y. Zhao and H. Zeng. An efficient schedulability analysis for optimizing systems with adaptive mixed-criticality scheduling. *Real-Time Systems*, 53(4):467–525, 2017.
- [42] Y. Zhao and H. Zeng. The virtual deadline based optimization algorithm for priority assignment in fixed-priority scheduling. In *IEEE Real-Time Systems Symposium*, 2017.
- [43] Y. Zhao and H. Zeng. The concept of response time estimation range for optimizing systems scheduled with fixed priority. In *IEEE Real-Time and Embedded Technology and Applications Symposium*, 2018.