

Generalized Elastic Scheduling *

Thidapat Chantem and Xiaobo Sharon Hu
Department of Computer Science & Engineering
University of Notre Dame
Notre Dame, IN 46556
{tchantem,shu}@cse.nd.edu

M.D. Lemmon
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556
lemmon@nd.edu

Abstract

The elastic task model proposed by Buttazzo, et. al. [9] is a powerful model for adapting real-time systems in the presence of uncertainty. This paper generalizes the existing elastic scheduling approach in several directions. It reveals that the original task compression algorithm in [9] in fact solves a quadratic programming problem that seeks to minimize the sum of the squared deviation of a task's utilization from initial desired utilization. This finding indicates that the task compression algorithm may be applied to efficiently solve other similar types of problems. In particular, an iterative approach is proposed to solve the task compression problem for real-time tasks with deadlines less than respective periods. Furthermore, a new objective for minimizing the average difference of task periods from desired values is introduced and a closed-form formula is derived for solving the problem without recursion.

1. Introduction

A desirable property of any real-time system is the guarantee that it will perform at least beyond some pre-specified thresholds defined by system designers. This is usually not a concern under normal situations where analysis has been done offline to ensure system performance based on the regular workload. However, in response to an event such as user's input or changing environment, the load of the system may dynamically change in such a way that a temporal overload condition occurs. The challenge, then, is to provide some mechanism to guarantee the minimum performance level under such circumstances.

Many real-time task models have been proposed to extend timing requirements beyond the hard and soft deadlines based on the observation that jobs can be dropped

without severely affecting performance ([4], [18]). For example, Ramanathan et al. proposed both online [16] and offline [25] scheduling algorithms that are based on the (m,k) model, which is analyzed in [17]. In this model, up to $k - m$ consecutive jobs are allowed to be dropped in any sliding window of k . Moreover, [29] presented the Dynamic Window-Constrained Scheduling (DWCS) algorithm, which is similar except that the window k is fixed. Further enhancement to these successful models can be found in [24] and [23]. Other frameworks such as the imprecise computation model [13] and reward-based model [2] can be used to capture the situations where the quality of service is proportional to the amount of workload completed.

Despite the success of the abovementioned models in alleviating overload situations, it is sometimes more suitable to execute jobs less often instead of dropping them or allocating fewer cycles. For example, limitations on the throughput capacity of ad hoc communication networks [1] make it highly desirable to reduce overall network traffic by having control tasks adaptively adjust their periods in response to the actual activity level of the control application.

The work in [20] was among the first to address this type of requirements. Seto, et al. considered the problem of finding a feasible set of task periods to optimize a specific form of control performance measure and formulated the problem as a non-linear programming problem [26]. In [27], finding all feasible periods of a given set of tasks was studied for the Rate Monotone (RM) scheduling algorithm. Cervin et al. used optimization theory to solve the period selection problem online by adaptively adjusting task periods [11] with the focus on optimizing specific control performance. Recently, [5] offered an optimal search algorithm that solves the period selection problem for fixed-priority scheduling schemes. The algorithm may be applicable only during the design phase due to its potentially high complexity. Another interesting framework was introduced in [19] where task periods are adjusted in response to varying computation times.

*This work is supported in part by NSF under grant number CNS-0410771.

Buttazzo and his colleagues proposed an elegant yet more flexible framework known as the elastic task model [8] where deadline misses are avoided by increasing tasks periods until some desirable utilization level is achieved. The work in [10] extends the basic elastic task model to handle cases where the computation time is unknown, [9] incorporated a mechanism to handle resource constraints within the elastic framework, and [7] provided a means to smoothly adjust task execution rate. In addition, [14] uses a control performance metric as a cost function to find an optimal sampling interval for each task.

We will focus our attention on the elastic task model where the selection of task periods is central. The existing elastic scheduling algorithm determines task periods based on an elegant analogy between spring systems and task scheduling in which a task's resistance to changing its period is viewed as a spring's resistance to being compressed. In accordance with the *principle of least action* found in classical mechanics, this suggests that the elastic task model is really attempting to minimize some overall measure of the task set's *energy*, whose precise nature was not made clear in the original work.

This paper generalizes the existing elastic scheduling approach in several directions. First, we re-examine the problem of period determination in the elastic task model and show that the original task compression algorithm in [9] in fact solves a quadratic programming (QP) problem. The QP problem seeks to minimize the sum of the squared deviation of every task's utilization from its initial utilization. Identifying the nature of the optimization problem underlying the task compression algorithm is important in several aspects. For instance, it may suggest other relevant optimization objectives and shed light on determining task periods in the presence of uncertainty for other task models.

Second, we select an alternative objective function, i.e., minimizing the average difference of task periods from their desired minimum periods, as the sampling period is directly related to the performance of some controllers. We show that the solution to this optimization problem can be evaluated with a closed-form formula.

Lastly, the proposed framework is extended to cases where task deadlines are less than task periods. Such requirements often arise in control systems to reduce jitters, for example. Moreover, in some situations, it is desirable that tasks finish executing sooner, even if their periods are not up. We formulate the problem of period determination as a constrained optimization problem and propose a heuristic approach based on the task compression algorithm in [9] to solve the problem. The heuristic is guaranteed to find a feasible solution, if one exists. It is quite efficient and is hence suitable for online period adjustment. Experimental results show that the heuristic is able to find solutions that are close to the global optimal solution.

The remainder of the paper is organized as follows. We begin by reviewing key background materials in Section 2. Section 3 presents the solutions to the period determination problem for tasks with deadlines equal to their periods. The optimization approach is then extended in Section 4 to treat the case where task deadlines are less than task periods. Experimental results are presented and discussed in Section 5. Finally, the paper concludes with Section 6.

2. Preliminaries

This section describes the system under consideration as well as important assumptions made throughout the paper. We also briefly review the task compression algorithm used for period selection [9].

2.1. System model

We consider the following real-time task model. Each task τ_i is periodic and is characterized by the following 6-tuple: $(C_i, D_i, T_i, T_{i_{\min}}, T_{i_{\max}}, e_i)$, for $i = 1, \dots, N$, where N is the number of tasks in the system, C_i is the worst case execution time of τ_i , D_i is its deadline, and T_i is τ_i 's initial period. Furthermore, $T_{i_{\min}}$ denotes the most desirable period of τ_i , as specified by control system designers, whereas $T_{i_{\max}}$ represents the maximum time interval between two instances of a task that will prevent system performance from falling below some thresholds (e.g., a control system becomes unstable). The elastic coefficient, e_i , represents the resistance of task τ_i to increasing its period in face of changes. The smaller the elastic coefficient of a task, the harder it is to increase that task's period.

Task deadlines are first assumed to be equal to task periods. This requirement will be relaxed in Section 4. All task attributes are real values and are assumed to be known *a priori*. The current utilization of τ_i is $U_i = \frac{C_i}{T_i}$. Similarly, the minimum and maximum utilizations of τ_i is $U_{i_{\min}} = \frac{C_i}{T_{i_{\max}}}$ and $U_{i_{\max}} = \frac{C_i}{T_{i_{\min}}}$, respectively.

2.2. Elastic task model

In [9], Buttazzo, et. al. modeled a task as a spring system, where increasing or decreasing a task period is analogous to compressing or decompressing a spring. The elastic coefficient, e_i , introduced above hence has its intuitive meaning of the hardness of the spring. The purpose of increasing task periods is to drive the total utilization of the system down to some desired utilization level, U_d , analogous to a spring system trying to minimize its energy under an external force.

The attractiveness of the elastic task model is its accompanying task compression algorithm, which is quite efficient ($O(N^2)$) and can readily be used online. (In fact, the

elastic task model and the task compression algorithm have already been implemented in the S.Ha.R.K. kernel [15].) The task compression algorithm works as follows. If it is possible to drive the system utilization down to U_d without violating any period bounds, the algorithm will return a set of feasible periods (T_1, T_2, \dots, T_N) that can be used by the system. Tasks whose periods are fixed (if $e_i = 0$ or $T_{i_{\min}} = T_{i_{\max}}$) are considered inelastic and are treated as special cases. The amount of utilization that each remaining, non-inelastic task should receive is computed based on its elastic coefficient, initial period, and the amount of utilization that must be reduced to achieve U_d . The resultant period of a task τ_i is guaranteed to fall somewhere between $T_{i_{\min}}$ and $T_{i_{\max}}$.

Throughout this paper, we will assume that the EDF (Earliest Deadline First) scheduling algorithm [21] is used. Furthermore, we will focus our attention on cases where tasks need to either increase or decrease its utilization in response to either internal (e.g., change in sampling rate of one or more tasks in the system) or external (e.g., network traffic) factors.

3. General period selection

Given a particular set of real-time tasks, there may exist numerous sets of feasible periods. It is not difficult to see that different sets of periods would lead to different performance of the resultant system. In general, the period selection problem can be expressed as an optimization problem. That is,

```
optimize: performance metric
s.t.:      tasks are schedulable
          period bounds are satisfied
```

Below we introduce two specific performance metrics and discuss their implications. We assume that tasks deadlines equal task periods.

3.1. Minimize utilization perturbation

Processor utilization by each task is an important measure for any real-time system. It not only reveals the amount of system resource dedicated to the task but also impacts schedulability. In the elastic task model, one consequence of changing task periods is changing the utilization of tasks. From the stand point of performance preservation, it is desirable to minimize the changes in task utilization. This objective can be captured by the following constrained optimization problem.

$$\min: \quad E(U_1, \dots, U_N) = \sum_{i=1}^N w_i (U_{i0} - U_i)^2 \quad (1)$$

$$\text{s.t.}: \quad \sum_{i=1}^N U_i \leq U_d \quad (2)$$

$$U_i \geq U_{i_{\min}} \quad \text{for } i = 1, 2, \dots, N \quad (3)$$

$$U_i \leq U_{i0} \quad \text{for } i = 1, 2, \dots, N \quad (4)$$

In the formulation, N is the number of tasks in the system, U_{i0} is the initial utilization of task τ_i and $U_{i0} \geq U_{i_{\min}}$, U_i is the utilization of τ_i to be determined, and U_d is the desired total utilization. (U_d is usually set to 1 for EDF scheduling.) Constant $w_i (\geq 0)$ is a weighting factor and reflects the criticality of a task. More critical tasks would have larger w_i 's. The first constraint simply states the schedulability condition under EDF. The rest of the constraints bounds the utilization, equivalently bounds the task period by $T_{i_{\min}}$ and $T_{i_{\max}}$ where $T_{i_{\min}} = C_i/U_{i_{\max}}$ and $T_{i_{\max}} = C_i/U_{i_{\min}}$.

The above constrained optimization problem belongs to the category of quadratic programs and can be solved in polynomial time. However, solving such a problem during runtime can be too costly. What makes the above formulation attractive is that its solution is exactly the same as that found by the task compression algorithm in [9]. We introduce a number of lemmas and a theorem to support this argument.

Lemma 1 *Given the constrained optimization problem as specified in (1)–(4) and $\sum_{i=1}^N U_{i0} > U_d$, any solution, U_i^* , to the problem must satisfy $\sum_{i=1}^N U_i^* = U_d$.*

Proof: We prove the lemma by utilizing the Kuhn-Tucker necessary conditions for constrained optimization problems. The Kuhn-Tucker necessary conditions for the solution to the given problem can be written in terms of the Lagrangian function for the problem as

$$J_a(\mathbf{U}, \mu) = \sum_{i=1}^N w_i (U_{i0} - U_i)^2 + \mu_0 \left(\sum_{i=1}^N U_i - U_d \right) + \sum_{i=1}^N \mu_i (U_{i_{\min}} - U_i) + \sum_{i=1}^N \lambda_i (U_i - U_{i0}) \quad (5)$$

where μ_i 's are Lagrange multipliers and $\mu_i \geq 0$ for $i = 1, \dots, N$. The necessary conditions for the existence of a relative minimum at U_i^* are

$$0 = \frac{\partial J_a}{\partial U_i^*} = -2w_i (U_{i0} - U_i^*) + \mu_0 - \mu_i - \lambda_i \quad (6)$$

$$0 = \mu_0 \left(\sum_{i=1}^N U_i^* - U_d \right) \quad (7)$$

$$0 = \mu_i (U_{i_{\min}} - U_i^*) \quad (8)$$

$$0 = \lambda_i (U_i^* - U_{i0}) \quad (9)$$

where $i = 1, \dots, N$.

Assume that (7) is inactive, i.e., $\mu_0 = 0$ and $\sum_{i=1}^N U_i^* < U_d$. Then at least one constraint in (8) or (9) must be active. Suppose the k -th constraint in (8) is active. That is, $U_k^* = U_{kmin}$ and $\mu_k \geq 0$. Then, the k -th constraint in (9) must be inactive, i.e., $\lambda_k = 0$. From Equation (6), we obtain

$$\mu_k = -2w_k(U_{k0} - U_{kmin}) \quad (10)$$

Since $U_{k0} > U_{kmin}$, then $\mu_k < 0$, which contradicts the assumption that $\mu_k \geq 0$ (i.e., the k -th constraint in (8) is active). Now consider some constraints in (9) are active, while others are not, say $\lambda_k = 0$. Then, (6) becomes

$$-2w_k(U_{k0} - U_k^*) = 0 \quad (11)$$

which again contradicts the assumption. Thus, for constraint (7) to be inactive, all constraints in (9) must be active, i.e., $U_k^* = U_{k0}$. However, since $\sum_{i=1}^N U_{k0} > U_d$, this cannot be a feasible solution to the optimization problem. Therefore, for any solution to the optimization problem, constraint (7) must be active, i.e., $\sum_{i=1}^N U_i^* = U_d$.

□

Lemma 2 Given the constrained optimization problem as specified in (1)–(4) and $\sum_{i=1}^N U_{i0} > U_d$, a solution to the problem, U_i^* , is optimal if and only if

$$U_i^* = U_{i0} - \frac{\frac{1}{w_i} \left(\sum_{U_j^* \neq U_{jmin}} U_{j0} - U_d + \sum_{U_j^* = U_{jmin}} U_j \right)}{\sum_{U_j^* \neq U_{jmin}} (1/w_j)} \quad (12)$$

if $U_i^* > U_{imin}$, and $U_i^* = U_{imin}$ otherwise.

Proof: Consider the Kuhn-Tucker conditions given in (6)–(9). From Lemma 1, we know that any solution to the given optimization problem must satisfy (7), i.e.,

$$U_d = \sum_{i=1}^N U_i^*. \quad (13)$$

Consider the case where all constraints in (9) are inactive, i.e. $\lambda_i = 0$, for all $i = 1, \dots, N$, and supposed that the k -th constraint in (8) is active, we have $U_k^* = U_{kmin}$, and

$$\mu_k = \mu_0 - 2w_k(U_{k0} - U_{kmin}), \quad (14)$$

Otherwise, we have $\mu_k = 0$. By summing up Equation (6) for all i and using the conclusions above, μ_0 can be expressed as

$$\frac{2 \left(\sum_{U_i^* \neq U_{imin}} U_{i0} - U_d + \sum_{U_i^* = U_{imin}} U_{imin} \right)}{\sum_{U_i^* \neq U_{imin}} (1/w_i)} \quad (15)$$

as long as $\sum_{U_i^* \neq U_{imin}} U_{i0} + \sum_{U_i^* = U_{imin}} U_{imin} > U_d$, $\mu_0 > 0$, $\mu_i \geq 0$, and constraints in (4) are satisfied. Therefore, a solution, U_i^* , to the optimization problem either satisfies $U_i^* = U_{imin}$ or can be obtained by combining (15) with (6) for $U_i^* > U_{imin}$. (Note that $\mu_i = 0$ when $U_i^* > U_{imin}$.) That is, $U_i^* =$

$$U_{i0} - \frac{\frac{1}{w_i} \left(\sum_{U_j^* \neq U_{jmin}} U_{j0} - U_d + \sum_{U_j^* = U_{jmin}} U_{jmin} \right)}{\sum_{U_j^* \neq U_{jmin}} (1/w_j)} \quad (16)$$

Additionally, since both the objective function and the inequality constraints in (1)–(4) are convex, the necessary conditions for optimality provided by the Kuhn-Tucker conditions also become the sufficient conditions for optimality [22]. Hence, the solution found in Lemma 2 is a global minimum.

□

The following theorem can readily be proved based on the above two lemmas. The proof can be found in [12].

Theorem 1 Consider a set of N tasks where U_i is the utilization of the i th task. Let U_{i0} denote the initial desired utilization of task τ_i and let $e_i > 0$ be a set of elastic coefficients for $i = 1, \dots, N$. Let U_d be the desired utilization level and $\sum_{i=1}^N U_{i0} > U_d$. The task utilizations U_i , for $i = 1, \dots, N$ obtained from the task compression algorithm in [9] minimizes

$$E(U_1, \dots, U_N) = \sum_{i=1}^N \frac{1}{e_i} (U_{i0} - U_i)^2$$

subject to the inequality constraints $\sum_{i=1}^N U_i \leq U_d$, $U_i \geq U_{imin}$, and $U_i \leq U_{i0}$, for $(i = 1, \dots, N)$.

The above theorem has several significant consequences. Firstly, it reveals the optimization criterion inherent in the task compression algorithm. Secondly, it illustrates that the task compression algorithm can be used to solve certain convex programming problems. Finally, it may suggest other relevant optimization objectives and shed light on determining task periods in the presence of uncertainty for other task models.

3.2. Minimize task periods

For some controllers, instead of focusing on utilization, it may be more useful to examine task periods directly, as the sampling period is monotonically decreasing faster than the controller performance.

If the range of possible period values for each task is unbounded, the following constrained optimization problem

can be used to minimize the changes in task periods.

$$\min: \quad J(T_1, \dots, T_N) = \sum_{i=1}^N w_i(T_i - T_{i0}) \quad (17)$$

$$\text{s.t.}: \quad \sum_{i=1}^N \frac{C_i}{T_i} \leq 1 \quad (18)$$

The constraint in equation (18) is simply the condition assuring that the task set is schedulable under EDF. Let T_i^* denote a locally optimal set of task periods for the above optimization problem. The following theorem provides a closed-form formula to compute the solution to this minimization problem. The theorem can be proved through a straightforward application of the Kuhn-Tucker necessary conditions and the proof is thus omitted.

Theorem 2 *Given the constrained optimization problem specified in equations (17)–(18), a locally optimal solution is*

$$T_i^* = \sqrt{\frac{C_i}{w_i} \sum_{k=1}^N \sqrt{w_k C_k}} \quad (19)$$

Remark: If we require that $T_i \geq T_{i_{\min}}$ (in other words, there is a lower bound on the desired period), the constraint will be satisfied if the choice of w_i satisfies the following inequality:

$$\sqrt{w_i} T_{i_{\min}} \leq \sqrt{C_i} \sum_{k=1}^N \sqrt{w_k} \sqrt{C_k} \quad (20)$$

The above expression is obtained by simply letting (19) be greater than or equal to $T_{i_{\min}}$. The significance of (20) is that it can be used to identify a set of admissible weighting factors whose selection ensures $T_i \geq T_{i_{\min}}$. Indirectly, (20) also ensures that the resultant task periods will be greater than zero.

4. Period selection with additional deadline constraints

In this section, we consider the case where task deadlines are less than task periods. This more general model is useful, as there are situations where it is desirable for a task to finish executing early (before its period ends). By using the optimization framework introduced in Section 3.1, we again formulate the period selection problem as a constrained optimization problem. We also propose a novel heuristic based on the task compression algorithm. The algorithm is guaranteed to find a local optimal solution to the problem, if one exists and is efficient enough for online use.

4.1. Simplify feasibility condition

Baruah et al. considered the case where task deadlines are less than or equal to task periods and derived a set of sufficient and necessary conditions for EDF schedulability [3], which are later improved in [6]. The conditions can be restated in the following theorem.

Theorem 3 *Given a periodic task set with $D_i \leq T_i$, the task set is schedulable if and only if the following constraint is satisfied $\forall L \in \{kT_i + D_i \leq \min(B_p, H)\}$ and $k \in \mathcal{N}$ (the set of natural numbers including 0), where B_p and H denote the busy period and hyperperiod, respectively,*

$$L \geq \sum_{i=1}^N \left(\left\lfloor \frac{L - D_i}{T_i} \right\rfloor + 1 \right) C_i \quad (21)$$

Based on Theorem 3, the period determination problem can be formulated as follows:

$$\min: \quad E(U_1, \dots, U_N) = \sum_{i=1}^N w_i (U_{i0} - U_i)^2 \quad (22)$$

$$\text{s.t.}: \quad L \geq \sum_{i=1}^N \left(\left\lfloor \frac{L - D_i}{T_i} \right\rfloor + 1 \right) C_i \quad (23)$$

$$L \in \{kT_i + D_i \leq \min(B_p, H)\}, k \in \mathcal{N} \quad (24)$$

$$U_i \geq U_{i_{\min}} \quad \text{for } i = 1, 2, \dots, N \quad (25)$$

$$U_i \leq U_{i_{\max}} \quad \text{for } i = 1, 2, \dots, N \quad (26)$$

Solving the above constrained optimization problem can be extremely time consuming. Hence, we investigate solving the problem approximately with an efficient algorithm below. An approximate solution is both acceptable and preferred, as a rapid response allows the system to degrade gracefully instead of going into catastrophic states caused by some dynamic perturbations. Since verifying the constraint in (21) for all L values is the main source of high complexity, we consider simplifying the feasibility test by using the following stronger schedulability condition,

$$L \geq \sum_{i=1}^N \left(\frac{L - D_i}{T_i} + 1 \right) C_i \quad (27)$$

It is not difficult to see that if the inequality in (27) is satisfied then the original inequality in (21) must also be satisfied. What makes (27) an excellent candidate for online use is that the schedulability of a task set can be determined based on a single L value, L^* . Below, we introduce several lemmas and a theorem to support this claim.

For simplicity, we denote the set of all possible values of L by a distinct ordered set $\mathcal{L} = \{L_0, L_1, \dots\}$ where $L = kT_i + D_i, \forall k \in \mathcal{N}$ and $L \leq \min(B_p, H)$.

Lemma 3 Given a set Γ of N tasks with $D_i \leq T_i$, let L_j and $L_{j+1} \in \mathcal{L}$ and let $L_j < L_{j+1}$. If the constraint in (27) is satisfied for L_j , then it is satisfied for L_{j+1} .

Proof: By regrouping the terms in (27), we can rewrite the inequality as follows.

$$L \geq \frac{\sum_{i=1}^N C_i - \sum_{i=1}^N U_i D_i}{1 - \sum_{i=1}^N U_i} \quad (28)$$

Given that L_j satisfies the constraint in (28) and $L_{j+1} > L_j$, it immediately follows that L_{j+1} satisfies the constraint in (28). \square

Based on the above lemma, we can conclude that if the constraint in (27) is satisfied for L_j , then it is also satisfied for all $L_k \in \mathcal{L}$, where $L_k > L_j$. It may then seem natural to simply set L^* to be the minimum of all L values in \mathcal{L} . However, such a choice can be extremely pessimistic, often resulting in finding no feasible solutions to the problem. To avoid being too pessimistic, we introduce the next lemma, which identifies useful necessary conditions for any feasible task set. The lemma helps to eliminate pessimistic choices of L^* .

Lemma 4 Let D_i be the deadline of task τ_i in a given task set Γ . Further, let the tasks in Γ be ordered in a non-decreasing order of deadlines and suppose that D_{\min} is unique. Regardless of the choices of periods, any task set that is schedulable must satisfy the following property:

$$\sum_{i=1}^j C_i \leq D_j, \forall j = 1, \dots, N \quad (29)$$

Proof: We can prove (29) by contradiction. Suppose the task set is schedulable and $\sum_{i=1}^j C_i > D_j$ for some j . By D_j , an instance of τ_1, \dots, τ_j must each be executed. Moreover, The total processor demand at D_j is at least $\sum_{i=1}^j C_i$ time units. However, since there are only D_j time units available and $\sum_{i=1}^j C_i > D_j$, at least one instance of a task must miss its deadline. This is a contradiction to the assumption that the task set is schedulable. Hence, (29) must be true for all $j = 1, \dots, N$. \square

We are now ready to introduce two lemmas which form the basis for our selection of L^* .

Lemma 5 Consider a set Γ of N tasks that satisfies the condition in Lemma 4. Let the tasks in Γ be sorted in a non-decreasing order of deadlines. If $D_1 + kT_1 \leq D_2$, $k = 0, 1, \dots$, and $L^* = D_2$ satisfies the inequality constraint in (27), then the task set is guaranteed to be schedulable.

Proof: Given that the task set satisfies Lemma 4, the inequality in (21) is automatically satisfied for $L_j = D_1 + kT_1$ such as $L_j \leq D_2$, where $j = 0, 1, \dots$ and $k = 0, 1, \dots$, since $C_1 \leq D_1$. From this point, $L_2^* = D_2$ is the minimum of all the remaining L values. Hence, if $L^* = D_2$ satisfies (27), the remaining L values must also satisfy (27) following Lemma 3. Since (27) is a stronger condition than (21), (21) will be satisfied for all L values. \square

Lemma 6 Consider a set Γ of N tasks that satisfy the condition in Lemma 4. Let the tasks in Γ be sorted in a non-decreasing order of deadlines. If $D_1 + T_1 > D_2$, and $L^* = \min_{i=1}^N (T_i + D_i)$ satisfies the inequality constraint in (27), then the task set is guaranteed to be schedulable.

Proof: Recall that \mathcal{L} is a distinct ordered set. There are two possible types of values that L^* could take: either L^* is set to some deadline D_k where $D_k < \min_{i=1}^N (T_i + D_i)$, or L^* is set to $\min_{i=1}^N (T_i + D_i)$. For any $D_k < \min_{i=1}^N (T_i + D_i)$ the inequality in (21) is automatically satisfied since the task set satisfies the conditions in Lemma 4. Therefore, $\min_{i=1}^N (T_i + D_i)$ is the smallest L amongst all the remaining L values that are not inherently covered by Lemma 4. Hence, if $L^* = \min_{i=1}^N (T_i + D_i)$ satisfies (27), all the remaining L values must also satisfy (27) following Lemma 3. Since (27) is a stronger condition than (21), (21) will be satisfied for all L values. \square

Based on Lemmas 5 and 6, we have a new schedulability test which is stated in the following theorem. The proof for the theorem can be found in [12].

Theorem 4 Consider a set Γ of N tasks that satisfies the condition in Lemma 4. Let the tasks in Γ be sorted in a non-decreasing order of deadline. A given task set is schedulable if

$$L^* \geq \sum_{i=1}^N \left(\frac{L^* - D_i}{T_i} + 1 \right) C_i \quad (30)$$

where

$$L^* = \begin{cases} D_2 & : D_1 + kT_1 \leq D_2 \\ \min_{i=1}^N (T_i + D_i) & : \text{otherwise} \end{cases}$$

where $k = 0, 1, \dots$

Attentive readers may have already noticed that, according to Lemma 4 and Theorem 4, at most $N + 1$ checks are required to test the feasibility of any given task set with N tasks without any previous assumptions.

The above theorem paves the way to finding a simpler constrained optimization problem formulation for the purpose of period determination. We present the actual problem formulation in the following subsection.

4.2. Minimize utilization perturbation with deadline constraints

By using Theorem 4, we can express the period determination problem where task deadlines are less than task periods as the following constrained optimization problem.

$$\min: \quad E(U_1, \dots, U_N) = \sum_{i=1}^N w_i (U_{i0} - U_i)^2 \quad (31)$$

$$\text{s.t.}: \quad L \geq \sum_{i=1}^N \left(\frac{L - D_i}{T_i} + 1 \right) C_i \quad (32)$$

$$L = \begin{cases} D_2 & : D_1 + k \frac{C_1}{U_1} \leq D_2 \\ \min\left(\frac{C_i}{U_i} + D_i\right) & : \text{otherwise} \end{cases} \quad (33)$$

$$U_i \geq U_{i \min} \quad \text{for } i = 1, 2, \dots, N \quad (34)$$

$$U_i \leq U_{i0} \quad \text{for } i = 1, 2, \dots, N \quad (35)$$

where $k = 0, 1, \dots$

Solving the above optimization problem with a general optimization software can still be rather inefficient in terms of both processor time and memory usage. Consequently, the challenge is how to solve the problem efficiently as to allow the system to respond fast to dynamic changes. We leverage the task compression algorithm to tackle the challenge as follows.

Let $k_i = L - D_i$ and $U'_i = k_i U_i$. By reorganizing the terms, the optimization problem in (31)–(35) can be rewritten as follows:

$$\min: \quad E(U_1, \dots, U_N) = \sum_{i=1}^N \frac{w_i}{k_i^2} (k_i U_{i0} - U'_i)^2 \quad (36)$$

$$\text{s.t.}: \quad \sum_{i=1}^N U'_i \leq L - \sum_{i=1}^N C_i \quad (37)$$

$$U'_i \geq k_i U_{i \min} \quad \text{for } i = 1, 2, \dots, N \quad (38)$$

$$U'_i \leq k_i U_{i0} \quad \text{for } i = 1, 2, \dots, N \quad (39)$$

Note that the above constrained optimization problem would have exactly the same format as the QP problem in (1)–(4) if L and k_i can be treated as constants.

Consider the case where $D_1 + k \frac{C_1}{U_1} \leq D_2$ ($k = 0, 1, \dots$). According to Lemma 5, we only need to check $L^* = D_2$ for schedulability, which indeed leads to a constant L value in (37). It follows that we can solve the optimization problem efficiently by using the following theorem, which we conjecture to be correct but have no formal proof at the time of this writing.

Theorem 5 *Given the constrained optimization problem as specified in (31)–(35), for $L = D_2$ and $\sum_{i=1}^N k_i U_i > L -$*

$\sum_{i=1}^N C_i$, a solution, U_i^* , is optimal if and only if

$$U_i^* = \begin{cases} \frac{D_2 - \sum_{j=1}^N C_j - \sum_{j=3}^N k_j U_{j0}}{D_2 - D_1} & : i = 1 \\ U_{i0} & : \text{otherwise} \end{cases}$$

The above theorem immediately leads to an efficient algorithm to solve the optimization problem in (31)–(35) when $D_1 + k T_1 \leq D_2$, $k = 0, 1, \dots$. Moreover, Theorem 5 also implies that if $U_1 < U_{1 \min}$ then the task set is infeasible.

Let us now consider the case where $D_1 + T_1 > D_2$. According to Lemma 6, one needs to check whether $L^* = \min_{i=1}^N (T_i + D_i)$ satisfies (27) to determine feasibility. Since the value of L^* may change as T_i changes, the constraint in (37) is no longer linear, which can greatly increase the complexity of the optimization problem. Since our aim is to have an efficient online algorithm, we propose a heuristic to solve the problem. The following theorem forms the basis for our heuristic approach. The theorem can be proved using the Kuhn-Tucker conditions as in Lemmas 1 and 2 and is omitted.

Theorem 6 *Given the constrained optimization problem as specified in (31)–(35), for a fixed value of L (where $L = \min\{\frac{C_i}{U_i} + D_i\}, \forall i = 1, \dots, N$) and $\sum_{i=1}^N k_i U_i > L - \sum_{i=1}^N C_i$, the solution, U_i^* , is optimal if and only if*

$$U_i^* = U_{i0} - \frac{k_i}{w_i \sum_{U_j^* \neq U_{j \min}} (k_j^2 / w_j)} \times \left(\sum_{U_j^* \neq U_{j \min}} k_j U_{j0} - (L - \sum_{i=1}^N C_j) + \sum_{U_j^* = U_{j \min}} k_j U_j \right) \quad (40)$$

Our proposed algorithm aims to find a feasible solution to the optimization problem defined in (22)–(25) in an efficient manner. The solution may not be optimal but it is guaranteed to be schedulable by the EDF policy. The algorithm adopts an iterative approach. After iteration h , a set of periods $T_i(h)$ is found and the algorithm checks to see whether the constraint in (27) is satisfied. If this is the case and if $T_i(h)$ also minimize the objective function till now, then the algorithm keeps $T_i(h)$ as the current best solution to the problem. The detailed algorithm is given in Figure 1.

In each iteration h , we fix the value of L as either $L(h) = D_2$ if $T_1(h-1) + D_1 \leq D_2$ or $L(h) = \min_{i=1}^N (T_i(h-1) + D_i)$ otherwise. Task utilization, $U_i(h)$, can then be determined using Theorem 5 or Theorem 6, respectively. In the case of $L(h) = \min_{i=1}^N (T_i(h-1) + D_i)$, as shown in Figure 1, U_i is obtained by using a slightly modified task compression algorithm. (To save space, we omit the modified task compression algorithm.) The following modifications were made to the original task compression algorithm: (i) the inputs to the task compression algorithm are

task set Γ and $L(h)$, instead of Γ and U_d , and (ii) the equation $U_i = U_{i0} - (U_{v0} - U_d + U_f) E_i/E_v$ in the original algorithm is replaced by (40). For the case where $L(h) = D_2$, Theorem 5 is applied straightforwardly.

To determine convergence, a user-defined parameter, Δ , is included as a stopping criterion; if the difference between U_i found in the current iteration and U_i found in the last iteration is smaller than Δ for all i , the algorithm terminates and returns the best set of periods it has encountered. To handle the case where task periods do not converge to some fixed values (or when it may take too long for the solution to converge), the algorithm uses another user-defined parameter, $maxIter$, to limit the maximum number of iterations.

An additional challenge is how to assign the initial value of L . We propose to set the initial value of L to $\min_{i=1}^N (T_{i_{max}} + D_i)$. In this way, if the task set is found to be infeasible, then the algorithm immediately exits since the task set cannot be made schedulable without violating the given period bounds. The following lemma serves to support our choice of the initial value as well as the iterative approach. The proof of the lemma is omitted due to page limit and can be found in [12].

Lemma 7 *Let T_i for $1 \leq i \leq N$ be a set of periods that satisfies the constraint in (21). Then the set of $T'_i \geq T_i$ also satisfy the constraint in (21).*

The above lemma has two significant consequences. First, if our algorithm cannot find a feasible solution when setting $L(0) = \min_{i=1}^N (T_{i_{max}} + D_i)$, it is not fruitful to continue with the algorithm as any smaller T_i 's would not satisfy the constraints in (21). Second, to guarantee feasibility, it is advantageous to use larger T_i 's. However, larger T_i 's tends to increase the objective function. Therefore, by setting $L(h) = \min_{i=1}^N (T_{i_{h-1}} + D_i)$, we hope to achieve a good balance.

The following theorem states the correctness and complexity of our proposed algorithm. The proof of the theorem follows directly from the lemmas in this and previous subsections and is omitted due to page limit. However, the proof can be found in [12].

Theorem 7 *Given a set of N periodic tasks with deadlines less than periods, the algorithm in Figure 1 takes $O(N^2 \cdot maxIter)$ time to output a set of task periods. Moreover, if the solution set is non-empty, the task set with the new periods is guaranteed to be schedulable by the EDF policy.*

Finally, with sufficiently small $maxIter$, the time complexity makes the proposed algorithm suitable for online period adjustments. In the next section, we will provide some guidance on how to adjust such user-defined parameters based on some experimental results.

```

Algorithm Task_compress_deadline( $\Gamma, \Delta, maxIter$ ) {
  sumC = 0;
  for each ( $\tau_i \in \Gamma$ ) {
    sumC = sumC +  $C_i$ ;
    if ( $sumC > D_1$ )
      return NULL; // no feasible solution exists
  }

  bestObjF =  $\infty$ ;
  for each ( $\tau_i \in \Gamma$ ) {
    prevTi =  $T_{i_{min}}$ ;
    currTi =  $T_{i_{max}}$ ;
  }
  for ( $h = 0, h < maxIter, h = h + 1$ ) {
    if ( $D_1 + currT_1 \leq D_2$ )
      L =  $D_2$ ;
    else
      L =  $\min_{i=1}^N (currT_i + D_i)$ ;

    objF = cns = 0;
    for each ( $\tau_i \in \Gamma$ ) {
      objF = objF +  $\frac{1}{c_i} (U_{i0} - C_i / currT_i)^2$ ;
      cns = cns +  $(\frac{L - D_i}{currT_i} + 1) C_i$ ;
    }
    if ( $cns > L$ ) and ( $h = 0$ )
      return NULL; // no feasible solution can be found
    if ( $cns \leq L$ ) and ( $objF < bestObjF$ ) {
      bestObjF = objF;
      for each ( $\tau_i \in \Gamma$ )
        bestTi = currTi;
    }
    for each ( $\tau_i \in \Gamma$ ) {
      deltaTi =  $|currT_i - prevT_i|$ ;
      prevTi = currTi;
    }
    if deltaTi  $\leq \Delta$ , break; //achieved convergence
    if ( $D_1 + currT_1 \leq D_2$ ) {
      Compute currT following Theorem 5;
    }
    else
      currT = Mod_Task_compress( $\Gamma, L$ );
  }
  return bestT;
}

```

Figure 1. Period selection alg. for $D_i < T_i$

5. Experimental results

Here, we illustrate the application of our flexible scheduling framework and verify the claims made in the previous sections.

5.1. General period selection

To demonstrate that the task compression algorithm solves the optimization problem in(6)–(9), we reuse the task set provided in the experimental results section of [9] (reproduced below in Table 1). The task compression algorithm was written in C++, while the results obtained from

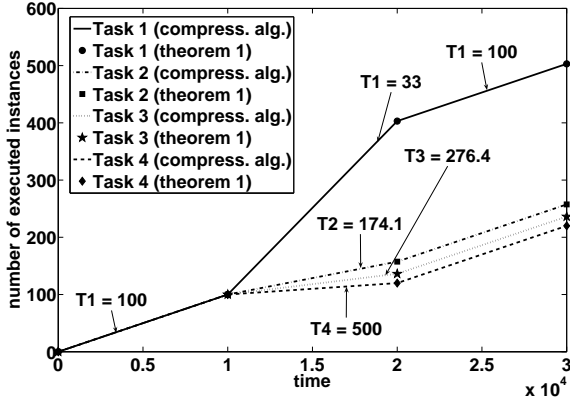


Figure 2. Utilization perturbation example

the constrained optimization problems and the closed-form expression were obtained from `MatLab`. In this experi-

Table 1. Task set parameters used

Task	C_i	T_{i0}	$T_{i\min}$	$T_{i\max}$	e_i
τ_1	24	100	30	500	1
τ_2	24	100	30	500	1
τ_3	24	100	30	500	1.5
τ_4	24	100	30	500	2

ment, all tasks start at time 0 with an initial period of 100 time units. 0.192. Since the current utilization is 0.96, the task set is schedulable under EDF. Assume that, at time 10000, τ_1 needs to reduce its period to 33 time units, perhaps due to some changes in system dynamics not experienced by other tasks. To accomplish this, the period of tasks τ_2 , τ_3 , and τ_4 must increase for the system to remain schedulable. At time 20000, τ_1 goes back to its original period. Figure 2 shows the cumulative number of executed instances for each task as its period changes over time. First of all, the data verifies that the results obtained from the task compression algorithm and those from Theorem 1 match perfectly.

Furthermore, it can be seen from the graph that the number of executed instances of a task is inversely proportional to its elastic coefficient. Recall that the weight of a task is the inverse of its elastic coefficient. From Figure 2, τ_2 , τ_3 , and τ_4 all have the same computation time, initial period, and period range, but τ_2 is determined to have the smallest (e.g. best) sampling period because of its weight. On the other hands, τ_4 has the largest sampling period because it is considered to be of least importance. Due to page limit, we omit the graph that shows the resultant periods of the same task set, only this time with the objective of minimizing task periods. The actual graph can be found in [12]. The results from (17)–(18) matched those from the closed-form formula presented in Theorem 2, which verifies our previ-

ous claim that both are indeed equivalent.

5.2. Period selection with additional deadline constraints

Since the algorithm proposed in Section 4 for period selection when task deadlines are less than task periods uses a heuristic, it is quite difficult to quantify its exact performance. However, in this subsection, we attempt to approximate the goodness of the algorithm as compared to the optimal solution by performing an experiment consisting of 50 randomly generated task sets¹. Each set contains of 5 tasks. As before, task computation times and deadlines were kept constant. Moreover, based on initial parameters, each task set was guaranteed not be schedulable. However, if the task periods were set to $T_{i\max}$, then a feasible schedule exists based on the constraint in (21).

The proposed algorithm was written in C++, whereas the optimization problem from (31)–(33) was solved using `Loqo`, a non-linear solver. The proposed algorithm found a convergent set of periods for 43 out of 50 sets. The user-defined parameters $maxIter$ and Δ were set to 200 and 1×10^{-10} , respectively. The reason why Δ was so small is to show that task periods found by the heuristic indeed converged to some fixed values. On the other hand, `Loqo` only found an optimal solution to 16 sets provided that a maximum of 500 iterations were run. The solver was not run longer, as, according to [28], it is unlikely that an optimal solution will be found after the 100th iteration. Table 2 shows the difference in total utilization obtained from using the proposed algorithm and that from `Loqo` for the task sets that were solved optimally by `Loqo`.

Table 2. Difference in total utilization

Benchmark	Iterations needed (Heuristic)	Total utilization difference (%)
2	3	0.000210
7	39	0.000054
10	200+	8.761857
12	2	0.000023
20	2	0.005683
25	21	5.657413
28	2	0.000099
29	200+	5.201867
32	17	0.059623
33	200+	3.088146
37	200+	0.925411
40	32	0.154522
42	200+	1.421277
44	49	2.745394
46	4	7.517813
48	2	0.000011

As can be seen by Table 2, 11 task sets are directly comparable. In other words, 11 task sets were optimally solved by `Loqo` and a convergent set of periods were found by

¹The authors thank Bren Mochocki for providing this tool to accomplish this step

the algorithm for these same sets. Amongst them, 7 sets have the same solutions, characterized by a difference of less than 0.01% in total utilization. The other 4 task sets were found to have different solutions where the maximum discrepancy was less than 8%. We suspect that such discrepancy was a result of the algorithm finding a local minimizer. Interestingly, for the task sets that the algorithm could not find a convergent set of periods, the maximum difference in utilization was less than 9%.

Based on the above results, it seems that the convergent solution set found by the proposed algorithm is in general on par with the optimal solution. In several instances, the heuristic successfully returns a set of feasible solutions while it was not possible to solve the optimization problem directly. Since the complexity of the heuristic is lower, the proposed algorithm is preferable for online period adjustment. Lastly, the experiment suggests that the maximum number of iterations, $maxIter$, need not be greater than 100. On the other hand, Δ can be set to be in the order of time granularity used by the operating system.

6. Conclusion and future work

By examining the elastic task model more closely, we were able to create a general framework where the elastic task model can be treated as a special case. Our flexible framework allows for formulating a problem in a systematic way, which can make it easier to develop an efficient algorithm or expression that will optimize a specific performance measure. In this framework, trade-offs can be viewed as optimization problems with potentially infinite performance criterion. By further studying the task compression algorithm, we were able to gain insights regarding some relevant performance criteria such as the minimization of perturbation in task utilizations as well as in task periods. In addition, as shown in the paper, the task compression algorithm can not only solve a QP problem, but also be used as a powerful component of the proposed heuristic.

The framework can be extended to treat cases where the computation times are variable. For example, the proposed framework can be used to adjust task periods when the computation time changes.

Since the algorithm presented in Section 4 is best-effort, it would be interesting to study whether there exists a way to select the value of L at every iteration such that the solution found will always be optimal.

Finally, we plan on exploring different classes of objective functions and constraints that may be even harder to solve. The case for aperiodic tasks may be worth investigating, since they impose a different type of constraints such as response times.

References

- [1] P. Antsaklis and J. B. (editors). Special issue on networked control systems. *Trans. on Automatic Control*, 49:1421–1423, 2004.
- [2] H. Aydin, R. Melhem, D. Mosse, and P. Alvarez. Optimal reward-based scheduling for periodic real-time tasks. *RTSS*, 1999.
- [3] S. Baruah, L. Rosier, and R. Howell. Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor. *Real-Time Systems*, Nov 1990.
- [4] G. Bernat, A. Burns, and A. Llamas. Weakly hard real-time systems. *Trans. on Computers*, 2001.
- [5] E. Bini and M. D. Natale. Optimal task rate selection in fixed priority systems. *RTSS*, pp. 399–409, 2005.
- [6] G. Buttazzo. *Hard real time computing systems: predictable scheduling algorithms and applications*. Springer, 2005.
- [7] G. Buttazzo and L. Abeni. Smooth rate adaptation through impedance control. *ECRTS*, Jun 2002.
- [8] G. Buttazzo, G. Lipari, and L. Abeni. Elastic task model for adaptive rate control. *RTSS*, 1998.
- [9] G. Buttazzo, G. Lipari, M. Caccamo, and L. Abeni. Elastic scheduling for flexible workload management. *Trans. on Computers*, 2002.
- [10] M. Caccamo, G. Buttazzo, and L. Sha. Elastic feedback control. *ECRTS*, 2000.
- [11] A. Cervin, J. Eker, B. Bernhardsson, and K.-E. Årzén. Feedback-feedforward scheduling of control tasks. *Real-Time Systems*, Jul 2002.
- [12] T. Chantem, X. Hu, and M. Lemmon. Generalized elastic scheduling. Technical report, U. of Notre Dame, 2006.
- [13] J.-Y. Chung, J. W. Liu, and K.-J. Lin. Scheduling periodic jobs that allow imprecise results. *Trans. on Computers*, 1990.
- [14] J. Eker, P. Hagander, and K.-E. Årzén. A feedback scheduler for real-time controller tasks. *Control Engineering Practice*, Dec 2000.
- [15] P. Gai, L. Abeni, M. Giorgi, and G. Buttazzo. A new kernel approach for modular real-time systems development. *ECRTS*, Jun 2001.
- [16] M. Hamdaoui and P. Ramanathan. A dynamic priority assignment technique for streams with (m,k)-firm deadlines. *Trans. on Computers*, Dec 1995.
- [17] M. Hamdaoui and P. Ramanathan. Evaluating dynamic failure probability for streams with (m,k)-firm deadlines. *Trans. on Computers*, Dec 1997.
- [18] G. Koren and D. Shasha. Skip-over: Algorithms and complexity for overloaded systems that allow skips. *RTSS*, Dec 1995.
- [19] X. Koutsoukos, R. Tekumalla, B. Natarajan, and C. Lu. Hybrid supervisory utilization control of real-time systems. *RTAS*, pp. 12–21, 2005.
- [20] T.-W. Kuo and A. Mok. Load adjustment in adaptive real-time systems. *RTSS*, 1991.
- [21] C. Liu and J. Layland. Scheduling algorithms for multiprogramming in a hard real-time environment. *Journal of the Assoc. for Computing Machinery*, 20(1), 1973.

- [22] D. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley Publishing Co., Reading, MA, 1989.
- [23] A. Mok and W. Wang. Window-constraint real-time periodic task scheduling. *RTSS*, Dec 2001.
- [24] G. Quan and X. Hu. Enhanced fixed-priority scheduling with (m,k)-firm guarantee. *RTSS*, Nov 2000.
- [25] P. Ramanathan. Overload management in real-time control applications using (m,k)-firm guarantee. *Trans. on Parallel and Distributed Systems*, Jun 1999.
- [26] D. Seto, J. Lehoczky, and L. Sha. Task period selection and schedulability in real-time systems. *RTSS*, 1998.
- [27] D. Seto, J. Lehoczky, L. Sha, and K. Shin. On task schedulability in real-time control systems. *RTSS*, 1996.
- [28] R. Vanderbei. Loqo. <<http://www.princeton.edu/rvdb/>>.
- [29] R. West and C. Poellabauer. Analysis of a window-constrained scheduler for real-time and best-effort packet streams. *RTSS*, Nov 2000.