

*Department of Computing Science
Illinois Institute of Technology*

A New Metric for Quantifying Similarity between Timing Constraint Sets in Soft Real-Time Systems

Yue Yu and Shangping Ren

Department of Computing Science
Illinois Institute of Technology
10 W. 31st Street, Chicago, IL 60616

Telephone (312) 567-5215

Facsimile (312) 567-5067

Email {yyu8, ren}@iit.edu

Technical Report CS-115-01-02-2009

February 8, 2009

Acknowledgements

We would like to thank Dr. Sharon Hu for her insightful comments and suggestions on this work and would also like to thank Thidapat Chantem for her help with ILOG CPLEX[®] while solving the MILP formulation in Section 6.3.

The research is supported by NSF under CAREER grant CNS 0746643

Abstract

Real-time systems are systems that their timing behaviors must satisfy a specified set of timing constraints and they often operate in a real-world environment with scarce resources. As a result, real-world performance of these systems may deviate from the design, either inevitably due to unpredictable factors, or by intention in order to improve system's other quality-of-service (QoS) properties. In this paper, we first introduce a new metric, *constraint set similarity*, to quantify the resemblance between two different timing constraint sets. Because directly calculating the exact value of the metric involves calculating the size of a polytope which is a $\#P$ -hard problem, we instead introduce an efficient method for estimating its bound. We further illustrate how this metric can be exploited for improving system predictability and evaluating trade-offs between timing constraint compromises and system's other QoS property gains.

Contents

Acknowledgements	i
Abstract	ii
1 Introduction	1
2 Motivation and Related Work	2
3 Timing Constraint Set Normal Form	3
4 Similarities between Timing Constraint Sets	8
4.1 Similarities between Constraint Sets	8
4.2 Discussions	13
5 Application 1: Predicting Tracking Error Rate Based on Constraint Similarity	16
6 Application 2: Improving Systems' QoS Properties with Constraint Similarity Guarantees	18
6.1 System and Task Model	20
6.2 Reducing Total Energy Consumption	20
6.3 Determining Constraint Relaxations	22
7 Conclusion	24

List of Figures

1	The feasible region of constraint $0ms < t(f_j) - t(s_j) \leq 22ms$	4
2	The feasible region of a constraint set (1).	5
3	The feasible regions satisfying constraint $0 < t(f_j) - t(s_j) \leq 22$ and $0 < t(f_j) - t(s_j) \leq 25$	8
4	The feasible regions satisfying constraint sets (1) (bold lines) and (14) (light lines), and their intersection (the shaded region).	9
5	Fixing $t(e_3)$ at 5 in (14) can be interpreted as using the plane $t(e_3) = 5$ to cut the feasible region of (14) in Fig. 4 and view the slice in the $\mathbf{x}_1\mathbf{x}_2$ ($t(e_1)t(e_2)$) plane.	11
6	Feasible region similarities of non-uniformly distributed timed data streams.	14
7	Similarity relation is not transitive.	15
8	Similarity between general timing constraint sets.	16
9	Four sensors tracking a moving target.	17
10	Actual trajectories and sensed coordinates of the moving target before and after the data consistency constraints are modified.	19

List of Tables

1 Introduction

Software for real-world systems often operates in an unpredictable environment and interacts with physical machineries. Hence, for most of these software systems, it is difficult and unrealistic to implement them in such a way that they behave *precisely* as specified due to the following facts:

- **System Complexity** The ever-increasing complexities of software systems have made guarantees of *exact* system behavior impractically expensive, if not impossible. For example, advances in computer architecture and software have made it difficult to predict the execution time of software, and networking techniques further introduce variability and stochastic behavior into the system [?].
- **Unpredictable Operating Environment** The intrinsically unpredictable nature of the environments in which software systems operate determines that even though software operates precisely as designed, its interactions with the outer world may not be totally expected. For example, Jackson et al. have shown that several aircraft accidents have been attributed to “mode confusion”, where the software operated as designed but not as expected by the pilots [?].
- **Computational Intractability** From a theoretical point of view, achieving exactness in the verification of system properties is sometimes intractable. For example, Alur and Dill have shown that the satisfiability of a very simple class of real-time properties such as “every p -state is followed by a q -state *precisely* 5 time units later” turns out to be undecidable in a continuous time model [?]. Although several real-time logics are decidable under discrete approximations of real time [?] or under interval timing constraints [?], these models unfortunately prohibit infinite precision.

Moreover, real-time and embedded systems often face trade-offs between time and limited resources. Therefore, even if a system can be implemented precisely as specified, relaxing some of the specifications may reduce resource consumptions: for hard real-time systems, all timeliness requirements must be met and thus optimizing other properties such as minimizing energy consumption must not violate timing constraints; for soft real-time systems, on the other hand, the requirement for timing constraint satisfaction guarantees is not as stringent. Such timing flexibility allowed by soft real-time systems can often be utilized to improve system’s other QoS properties, such as reduce total energy consumption. A challenging task in investigating the trade-offs between timing constraint satisfaction and other QoS properties is how to quantify the degree of timing constraint satisfaction. That is, how do we measure the level of satisfaction for some given timing behavior with respect to a set of timing constraints? Another closely related challenge is to determine which timing constraints to be relaxed and by how much in order to achieve certain other QoS objectives, e.g., energy consumption bound. Though some researchers have studied problems that are somewhat related to the above problems (to be discussed in the next section), we contend that there exists no systematic approach for tackling these challenges.

In this paper, we propose a framework for measuring timing constraint satisfaction which can be used to address the above challenges. Specifically, we introduce a novel metric, i.e., *constraint set similarity*, to capture the resemblance between two timing constraint sets. It is defined in terms of the common feasible region of two systems constrained by the two given timing constraint sets. This value reflects the probability of timing constraint satisfaction when the original timing constraints are violated or intentionally modified for improving QoS properties.

However, directly calculating the exact value of similarity between two sets of timing constraints is computationally intractable. To overcome this difficulty, we leverage the concept of similarity bound and derive a closed form formula for computing a tight similarity bound. This bound can be used to guide the design process and provide confidence guarantees on certain QoS properties.

To show how one may use the timing constraint similarity metric to predict real-world performances and guide design processes of real-time embedded systems, we

1. study the similarities between timing constraint sets in the specification of an object tracking system and its real-world deviation, and use the similarity to infer the relationships on other properties fulfilled by different but similar systems;
2. give a detailed design example in which a set of soft real-time tasks are executed on a multiprocessor system-on-chip (MPSoC) and the goal is to trade timing constraint satisfaction for reducing energy consumption. This concrete example serves as a demonstration that the similarity metric provides an effective tool to measure and guide the trade-offs between different QoS properties.

The rest of this paper is organized as follows. Next section provides a motivating example and reviews related work. Section 3 introduces a timing constraint set normal form. It is used to establish the constraint similarity metric. Section 4 presents the similarity metric that quantifies how much one set of timing constraints resembles another. Section 5 applies the theory of timing constraint similarities to an object tracking system for predicting tracking error rates. Section 6 studies how the theory of timing constraint similarities can be utilized to reduce the total energy consumption of an MPSoC system with minimal changes to the satisfaction of original timing constraints. Finally, we conclude and point out future work in Section 7.

2 Motivation and Related Work

To be able to quantify the level at which a timing constraint is satisfied in a soft real-time system has several important implications. For instance, it provides a systematic way to compare different system implementations when none of them can strictly meet the given timing constraints. In addition, it allows studies of “what if” scenarios where certain timing constraints are relaxed to some extent in order to improve other QoS properties. Furthermore, it can be used to judiciously decide design specifications.

One intuitive way to quantify the level of timing constraint satisfaction is to measure the probability with which a system satisfying a set of modified timing constraints still satisfies the original timing constraints. With such a probability, design alternatives with different timing behavior can be easily compared. We use a simple example to illustrate this point.

Example 1 *Consider scheduling a task j with a relative deadline of 22ms on an MPSoC with three cores m_1 , m_2 , and m_3 . The worst-case execution times (WCETs) of j on m_1 , m_2 , and m_3 are 20ms, 25ms, and 30ms with peak powers 10W, 7W, and 6W, respectively. For simplicity, we also assume that the actual execution times are uniformly distributed between 5ms and respective WCETs. Now, if we need to limit the peak power to be less than 8W, but allow some deadline misses, we can schedule the task on either m_2 or m_3 . If we schedule the task on m_2 , for instance, what we can guarantee is the satisfaction of a constraint with a relative deadline of 25ms, rather than 22ms. Similarly, with the task on m_3 , we can guarantee the satisfaction of a deadline of 30ms. In other words, in this example, to maintain the peak power below 8W, we have two different approaches. Now, the question is from timing perspective, which one is a better option?*

If task j is executed on m_2 , the probability of the system satisfying the original timing constraint of 22ms is $\frac{22-5}{25-5} = 85\%$. The probability reduces to $\frac{22-5}{30-5} = 68\%$ if task j is executed on m_3 . So for this simple example, the answer to the question above is obvious. That is, from the timing perspective, using m_2 is better than m_3 . Note that this conclusion coincides with the intuition that 25ms is ‘closer’ to 22ms than 30ms. However, this may not always be true — One could easily see this by considering the extreme case where the best-case execution time of j on m_2 is greater than 22ms. \square

From the above simple example, one can see that the probability with which a system satisfying a set of modified timing constraints still satisfies the original timing constraints can be used effectively to compare design alternatives with different timing behaviors. Now the challenge is how to measure such a probability when there are more complex timing constraints involved. Furthermore, given the timing constraint satisfaction as one of the system comparison criteria, how can we find a subset of constraints from a given constraint set and modify them so that the required non-timing properties (e.g., power consumption) are satisfied, but the timing property change is minimal, or the timing property is the most similar (closest) to the original one? The goal of this paper is to address these questions by introducing a new metric.

As related work, many researchers have studied feasibility probabilities for tasks with varying execution times. Tia et al. [?] propose a way to find the probability of a single task meeting its timing constraint, referred to as task feasibility probability. Kalavade et al. [?] present an approach to compute the probability of any single task delay exceeding its deadline, which is equivalent to the task feasibility probability. However, Hu et al. point out in [?] that the probability of each individual task meeting its timing constraint is not sufficient in several situations since there often exists strong correlation among events of tasks meeting their deadlines. The authors give a new metric that considers the overall system probabilistic behavior where tasks have their individual deadlines and the correlations between tasks are captured by precedence constraints. With this metric in the system-level design exploration process, one

can readily compare the probabilistic timing performance of alternative designs. Based on [?], Wang et al. [?] define a design metric called performance yield, which is the probability of the assigned schedule meeting the predefined performance constraints. However, none of these works consider the problem of measuring the level of timing constraint satisfaction when the original timing constraints cannot be satisfied or are intentionally modified.

Our study, on the other hand, focuses on a more generalized constraint model where correlation between tasks are treated as linear timing constraints. More specifically, we study similarities between two different timing constraint sets and use the similarity value to infer constraint satisfaction probability of a system that satisfies one set of timing constraints satisfies the other. Note that some of the feasibility research results can be used in combination with our proposed approach.

Many notions on similarities have been defined in the literature for process models. Gupta et al. [?] give a pseudometric analogue of bisimulation for generalized semi-Markov processes and show that two metrically similar processes have similar observable quantitative properties. Thorsley et al. [?] use Wasserstein pseudometrics to quantify the similarities between stochastic processes and introduce an algorithm to approximate the pseudometrics directly from sampled data rather than from process models themselves. The notion of similarity on other models are also studied, e.g., in [?, ?, ?]. However, the pseudometrics proposed in these works are used to compare processes. Though there are similarities between the idea of introducing quantitative metrics to measure two non-equivalent processes or constraints, the metrics introduced in this paper not only measures the resemblance between two sets of timing constraints, but also provides quantitative design guidance in deciding the trade-offs between constraint satisfaction and other QoS properties.

Trading one QoS property for another has been studied in various contexts. For example, reducing energy consumption through compromising system performance has been considered in a wide spectrum of computing. To name a few, Moscibroda et al. discuss the trade-off between energy efficiency and rapidity of event dissemination in ad hoc and sensor networks [?]; in high performance computing, Feng et al. analyzed NAS and SPEC suites to determine the relationship between frequency and voltage settings and execution time, and show that a significant decrease in energy is possible with a small increase in time [?]. In fact, for real-time and embedded system, dynamic voltage scaling techniques, which reduce system supply voltage for lower operation frequencies, has been extensively used in various power management schemes [?, ?, ?]. However, to our best knowledge, there is no quantitative study of trading timing constraint satisfaction in soft real-time systems for other QoS properties.

3 Timing Constraint Set Normal Form

In this section, we introduce the geometric foundations for characterizing timing constraint sets. The constraint normal form defined in this section will be used to establish constraint similarity metrics in Section 4.

In our system model, we take a commonly used approach in that system behaviors (or computations) are represented as *data streams*, i.e., a sequence of event occurrences (e_1, e_2, \dots, e_n) [?], and a *timed data stream* is formed by pairing each event e_i with its corresponding occurrence time $t(e_i)$, as defined below [?]:

Definition 1 (Timed Data Stream) *A timed data stream (TDS) is a sequence $((e_1, t(e_1)), (e_2, t(e_2)), \dots, (e_n, t(e_n)))$ where $(t(e_1), t(e_2), \dots, t(e_n))$ is a monotonically increasing sequence with elements in $\mathbb{R}^+ \cup \{+\infty\}$. Geometrically, a TDS is represented as a point in $|E|$ -dimensional space where each axis represents an event and the projection of the point on the axis represents the occurrence time of the corresponding event. \square*

Without timing constraints, events can occur at any time instances and thus the set of all TDS's occupies the entire nonnegative portion of the $|E|$ -dimensional space. However, when a set of timing constraints of the form $t(e_i) - t(e_j) \leq d (d \in \mathbb{R}^+ \cup \{+\infty\})$ exists, the set of TDS's satisfying the set of timing constraints is only a convex region in the $|E|$ -dimensional space and we call it *feasible region* throughout the paper. Feasible regions are the key in comparing timing constraint sets and we illustrate them in Example 2 and 3.




Figure 1: The feasible region of constraint $0ms < t(f_j) - t(s_j) \leq 22ms$.

Example 2 (2-Dimensional Feasible Region) Let s_j and f_j be the events that task j starts and finishes, the feasible region of the relative deadline constraint $0 < t(f_j) - t(s_j) \leq 22$ in Example 1 is shown in Fig. 1 (shaded area)

In the figure, TDS $((s_j, 20), (f_j, 38))$ in the feasible region satisfies the relative deadline constraint, while TDS's $((f_j, 16), (s_j, 28))$ and $((s_j, 8), (f_j, 40))$ outside the feasible region violates causality $t(s_j) - t(f_j) < 0$ and deadline $t(f_j) - t(s_j) \leq 22$, respectively. \square

The dimension of feasible regions becomes higher when the number of constrained events increases. Consider the following example:

Example 3 (3-Dimensional Feasible Region) Let the set of timing constraints that specify the relative time spans among three events be

$$\left\{ \begin{array}{ll} t(e_1) - t(e_2) \leq 6, & t(e_2) - t(e_1) \leq 6, \\ t(e_1) - t(e_3) \leq 7, & t(e_3) - t(e_1) \leq 3, \\ t(e_2) - t(e_3) \leq 9, & t(e_3) - t(e_2) \leq 14 \end{array} \right\} \quad (1)$$

Each timing constraint confines a half space in the 3-dimensional space and the intersection of such half spaces is the feasible region. The feasible region of (1) is shown in Fig. 2 with its boundaries marked as bold lines.

In the figure, the pentagonal prism circumscribed by all but the plane representing the constraint $t(e_3) - t(e_2) \leq 14$ characterizes the feasible region, i.e., each point $(t(e_1), t(e_2), t(e_3))$ in the region satisfies constraint set (1). \square

From Example 2 and 3, we can see that a feasible region characterizes all valid execution time traces, i.e., a system's valid timing behaviors under a set of timing constraints. However, when the dimension of a feasible region becomes higher, its shape becomes more complex and makes the graphical representation difficult. In order to compare feasible regions, alternative ways to represent high dimensional feasible regions are needed.

We introduce an algebraic representation to describe feasible regions such that comparisons can be directly made between feasible regions. This representation builds on the concept of the most stringent constraints, which we explain below by using Example 3 again. Examine the feasible region of Example 3 shown in Fig. 2. Note that the shape of the feasible region of (1) does not change when the constraint $t(e_3) - t(e_2) \leq 14$ is changed to $t(e_3) - t(e_2) \leq 9$ (or any other constraint value larger than 9). In fact, $t(e_3) - t(e_2) \leq 9$ is the most stringent timing constraints between event e_3 and e_2 which can be implied by the given constraint set.

For a given set of timing constraints, we can find the most stringent constraint set by leveraging the approach of finding all-pairs shortest paths. Specifically, we construct a constraint graph G by defining the vertex set of G as the set of events in the timing constraint set; for every two vertices e_i, e_j in G , there is an edge from e_i to e_j with weight d if there is a constraint $t(e_i) - t(e_j) \leq d$. The most stringent timing constraint implied by the given constraint set between every pair of events, $t(e_i) - t(e_j) \leq d_{i,j}^*$,




Figure 2: The feasible region of a constraint set (1).

can hence be derived from applying the Floyd-Warshall all-pairs shortest paths algorithm on G [?]. The most stringent constraint set has an important property which is summarized in the following lemma which shows that the feasible region of a set of real-time constraints does not change when constraints between all event pairs are replaced by the corresponding most stringent constraints derived from the Floyd-Warshall algorithm.

Lemma 1 *Given a set of m timing constraints of the form $t(e_i) - t(e_j) \leq d_k$ among n events, $\mathbf{A}\mathbf{t} \leq \mathbf{d}$, where \mathbf{A} is an $m \times n$ matrix, $\mathbf{t} = [t(e_1) \dots t(e_n)]^T$, and $\mathbf{d} = [d_1 \dots d_m]^T$. We have $\{\mathbf{t} \mid \mathbf{A}\mathbf{t} \leq \mathbf{d}\} = \{\mathbf{t} \mid \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}}\}$, i.e., the set of solutions of $\mathbf{A}\mathbf{t} \leq \mathbf{d}$ is the same as the set of solutions of $\tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}}$ where*

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & -1 & & & \\ 1 & & -1 & & \\ \vdots & & & \ddots & \\ 1 & & & & -1 \\ \hline -1 & 1 & & & \\ & 1 & -1 & & \\ & \vdots & & \ddots & \\ & 1 & & & -1 \\ \hline \vdots & & \dots & & \vdots \\ -1 & & & & 1 \\ & -1 & & & 1 \\ & & & \ddots & \vdots \\ & & & -1 & 1 \end{bmatrix} \quad \text{and} \quad \tilde{\mathbf{d}} = \begin{bmatrix} d_{1,2}^* \\ d_{1,3}^* \\ \vdots \\ d_{1,n}^* \\ \hline d_{2,1}^* \\ d_{2,3}^* \\ \vdots \\ d_{2,n}^* \\ \hline \vdots \\ \hline d_{n,1}^* \\ d_{n,2}^* \\ \vdots \\ d_{n,n-1}^* \end{bmatrix} \quad (2)$$

and $d_{i,j}^*$, $i \neq j$ are the shortest path weights.

Proof:

$$(i) \{ \mathbf{t} | \mathbf{A}\mathbf{t} \leq \mathbf{d} \} \supseteq \{ \mathbf{t} | \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}} \}$$

This directly follows from the fact that \mathbf{A} contains some rows of $\tilde{\mathbf{A}}$ and the corresponding d 's in \mathbf{d} is no less than those in $\tilde{\mathbf{d}}$ (the shortest path weights).

$$(ii) \{ \mathbf{t} | \mathbf{A}\mathbf{t} \leq \mathbf{d} \} \subseteq \{ \mathbf{t} | \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}} \}$$

Assume to the contrary that there is a vector $\mathbf{t}' = [t_1 \dots t_n]^T$ s.t. $\mathbf{t}' \in \{ \mathbf{t} | \mathbf{A}\mathbf{t} \leq \mathbf{d} \} \wedge \mathbf{t}' \notin \{ \mathbf{t} | \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}} \}$. This implies that the following set of linear inequalities has no solution

$$\begin{bmatrix} \mathbf{I} \\ -\mathbf{I} \\ \tilde{\mathbf{A}} \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \mathbf{t}' \\ -\mathbf{t}' \\ \tilde{\mathbf{d}} \end{bmatrix} \quad (3)$$

Based on Farkas' Lemma, together with the infeasibility of (3), we have that there exists an $(n^2 + n)$ -vector $[\mathbf{t}_1^T \quad \mathbf{t}_2^T \quad \mathbf{t}_3^T]^T$ where \mathbf{t}_1 and \mathbf{t}_2 are two n -vector and \mathbf{t}_3 is a $(n^2 - n)$ -vector, such that (4), (5), and (6) hold

$$\begin{bmatrix} \mathbf{I} & -\mathbf{I} & \tilde{\mathbf{A}}^T \end{bmatrix} \begin{bmatrix} \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \end{bmatrix} = \mathbf{0} \quad (4)$$

$$\begin{bmatrix} \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \end{bmatrix} \geq \mathbf{0} \quad (5)$$

$$\begin{bmatrix} \mathbf{t}'^T & -\mathbf{t}'^T & \tilde{\mathbf{d}}^T \end{bmatrix} \begin{bmatrix} \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \end{bmatrix} < 0 \quad (6)$$

From (4) we have that

$$\mathbf{t}_1 - \mathbf{t}_2 = -\tilde{\mathbf{A}}^T \mathbf{t}_3 \quad (7)$$

Insert (7) into (6) we have that

$$-\mathbf{t}'^T \tilde{\mathbf{A}}^T \mathbf{t}_3 + \tilde{\mathbf{d}}^T \mathbf{t}_3 = (\tilde{\mathbf{d}}^T - \mathbf{t}'^T \tilde{\mathbf{A}}^T) \mathbf{t}_3 < 0 \quad (8)$$

Therefore, it must be that

$$\exists i, j : d_{i,j}^* < t_i - t_j \quad (9)$$

since otherwise $\tilde{\mathbf{d}}^T - \mathbf{t}'^T \tilde{\mathbf{A}}^T \geq \mathbf{0}$ together with (5) would imply $(\tilde{\mathbf{d}}^T - \mathbf{t}'^T \tilde{\mathbf{A}}^T) \mathbf{t}_3 \geq 0$ which contradicts (8). However, (9) contradicts the fact that $d_{i,j}^*$ is the optimal solution to the linear program

$$\begin{aligned} & \mathbf{maximize} && t(e_i) - t(e_j) \\ & \mathbf{subject\ to} && \mathbf{A}\mathbf{t} \leq \mathbf{d} \end{aligned} \quad (10)$$

i.e., $d_{i,j}^*$ is the shortest path weight. Therefore, we have $\{ \mathbf{t} | \mathbf{A}\mathbf{t} \leq \mathbf{d} \} \subseteq \{ \mathbf{t} | \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}} \}$ and thus $\{ \mathbf{t} | \mathbf{A}\mathbf{t} \leq \mathbf{d} \} = \{ \mathbf{t} | \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}} \}$. \square

An important implication of Lemma 1 is that the shape of the feasible region is determined solely by the most stringent timing constraints between all *pairs* of events. Therefore, the constraint matrix that represents the most stringent constraints among all pairs of events uniquely characterizes the shape of the feasible region. We define this as the normal form of the constraint set.

Definition 2 (Constraint Set Normal Form) *Given a timing constraint set C and the corresponding constraint graph G , its all-pairs shortest paths matrix, denoted as \mathbf{D}^* , where*

$$\mathbf{D}^* = \begin{bmatrix} 0 & d_{1,2}^* & \cdots & d_{1,n}^* \\ d_{2,1}^* & 0 & \cdots & d_{2,n}^* \\ \vdots & \vdots & \ddots & \vdots \\ d_{n,1}^* & d_{n,2}^* & \cdots & 0 \end{bmatrix} \quad (11)$$

and $d_{i,j}^*$ is the shortest path weight between $t(e_i)$ and $t(e_j)$ in the constraint graph G . \mathbf{D}^* is called constraint set C 's normal form. \square

With Definition 2, the inclusion relation of two feasible regions defined by two timing constraint sets can be validated by comparing the constraint sets' normal forms.

Theorem 1 *Given two sets of real-time constraints $C = \mathbf{A}\mathbf{t} \leq \mathbf{d}$ and $C' = \mathbf{A}'\mathbf{t} \leq \mathbf{d}'$ on the same set of events¹. Let their corresponding normal forms be \mathbf{D}^* and \mathbf{D}'^* , respectively. The feasible region of C' is included within that of C if and only if $\mathbf{D}^* \geq \mathbf{D}'^*$, i.e., $\forall i, j : d_{i,j}^* \geq d'_{i,j}^*$.*

Proof:

Note that the feasible region of $\mathbf{A}'\mathbf{t} \leq \mathbf{d}'$ is included in that of $\mathbf{A}\mathbf{t} \leq \mathbf{d}$ iff the feasible region of $\begin{bmatrix} \mathbf{A} \\ \mathbf{A}' \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \mathbf{d} \\ \mathbf{d}' \end{bmatrix}$ is same as that of $\mathbf{A}'\mathbf{t} \leq \mathbf{d}'$. Hence, we can instead prove that the feasible regions of $\begin{bmatrix} \mathbf{A} \\ \mathbf{A}' \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \mathbf{d} \\ \mathbf{d}' \end{bmatrix}$ and $\mathbf{A}'\mathbf{t} \leq \mathbf{d}'$ are the same iff $\mathbf{D}^* \geq \mathbf{D}'^*$. In the following, for a constraint normal form \mathbf{D}^* , we use its equivalent system of linear inequalities representation $\tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}}$ where $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{d}}$ are defined in (2).

(i) $\mathbf{D}^* \geq \mathbf{D}'^* \Rightarrow C$ includes C' :

Suppose we have $\mathbf{D}^* \geq \mathbf{D}'^*$, i.e., $\tilde{\mathbf{d}} \geq \tilde{\mathbf{d}'}$, it is not hard to see that $\begin{bmatrix} \tilde{\mathbf{A}} \\ \tilde{\mathbf{A}} \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \tilde{\mathbf{d}} \\ \tilde{\mathbf{d}'} \end{bmatrix}$ has the same feasible region as $\tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}'}$. Then, from Lemma 1, $\begin{bmatrix} \mathbf{A} \\ \mathbf{A}' \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \mathbf{d} \\ \mathbf{d}' \end{bmatrix}$ has the same feasible region as $\mathbf{A}'\mathbf{t} \leq \mathbf{d}'$.

(ii) C includes $C' \Rightarrow \mathbf{D}^* \geq \mathbf{D}'^*$:

If $\begin{bmatrix} \mathbf{A} \\ \mathbf{A}' \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \mathbf{d} \\ \mathbf{d}' \end{bmatrix}$ has the same feasible region as $\mathbf{A}'\mathbf{t} \leq \mathbf{d}'$, then from Lemma 1, $\begin{bmatrix} \tilde{\mathbf{A}} \\ \tilde{\mathbf{A}} \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \tilde{\mathbf{d}} \\ \tilde{\mathbf{d}'} \end{bmatrix}$ has the same feasible region as $\tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}'}$. To prove the necessary condition, we should have $\mathbf{D}^* \geq \mathbf{D}'^*$, i.e., $\tilde{\mathbf{d}} \geq \tilde{\mathbf{d}'}$. Assume to the contrary that there is some $d_{i,j}^*$ in $\tilde{\mathbf{d}}$ and $d'_{i,j}^*$ in $\tilde{\mathbf{d}'}$ such that $d_{i,j}^* < d'_{i,j}^*$. Since $d'_{i,j}^*$ is the optimal solution to the linear program

$$\begin{aligned} & \text{maximize} && t(e_i) - t(e_j) \\ & \text{subject to} && \tilde{\mathbf{A}}\mathbf{t} \leq \tilde{\mathbf{d}'} \end{aligned} \tag{12}$$

and thus the optimal solution to the linear program (13)

$$\begin{aligned} & \text{maximize} && t(e_i) - t(e_j) \\ & \text{subject to} && \begin{bmatrix} \tilde{\mathbf{A}} \\ \tilde{\mathbf{A}} \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \tilde{\mathbf{d}} \\ \tilde{\mathbf{d}'} \end{bmatrix} \end{aligned} \tag{13}$$

However, the optimal solution to (13) can be at most $d_{i,j}^*$ when the solution set of $\begin{bmatrix} \tilde{\mathbf{A}} \\ \tilde{\mathbf{A}} \end{bmatrix} \mathbf{t} \leq \begin{bmatrix} \tilde{\mathbf{d}} \\ \tilde{\mathbf{d}'} \end{bmatrix}$ is not empty. The contradiction implies that $\mathbf{D}^* \geq \mathbf{D}'^*$. \square

From Theorem 1, we have the following:

$$\begin{aligned} \mathbf{D}^* = \mathbf{D}'^* & \Leftrightarrow \mathbf{D}^* \geq \mathbf{D}'^* \wedge \mathbf{D}'^* \geq \mathbf{D}^* \\ & \Leftrightarrow \text{feasible region of } C \text{ include that of } C' \\ & \quad \wedge \text{feasible region of } C' \text{ include that of } C \\ & \Leftrightarrow \text{feasible regions of } C \text{ and } C' \text{ are identical} \end{aligned}$$

In other words, there is a one-to-one correspondence between a timing constraint normal form and a feasible region. Therefore, the constraint normal form bridges the geometric problem of a feasible region and their corresponding algebraic problem of linear inequalities and can serve as the algebraic representation that we stated earlier in this section. We can hence derive the relationship between feasible regions of two different constraint sets by studying the constraint normal forms.

¹Note that the event sets of the two constraint sets need not be the same in order for the two feasible regions to be comparable. One can always extend both event sets to the same one by adding unconstrained events.

4 Similarities between Timing Constraint Sets

The example in Section 1 has shown that timing constraint changes often affect system's other QoS properties. In other words, there are trade-offs between the stringency of timing constraints and other QoS properties. For instance, relaxing some deadlines, or allowing certain probability of constraint violations, may reduce total energy consumptions. It is hence important to know how much the timing behavior compromise is in order to bring such QoS benefits.

4.1 Similarities between Constraint Sets

In this section, we focus on quantifying timing behavior similarities and we base our model on the feasible regions of timing constraint sets discussed in Section 3. The following example of the similarities between feasible regions in 2 and 3-dimensions gives the intuition. Note that in the following discussions, for simplicity, we assume that event occurrence times allowed by a set of constraints are uniformly distributed in the feasible region of the constraint set.

Example 4 (Feasible Region Similarity) *In Example 1, the original constraint was $0 < t(f_j) - t(s_j) \leq 22$ and the relaxed one is $0 < t(f_j) - t(s_j) \leq 25$. The relationship between the two corresponding feasible regions is depicted in Fig. 3.*




Figure 3: The feasible regions satisfying constraint $0 < t(f_j) - t(s_j) \leq 22$ and $0 < t(f_j) - t(s_j) \leq 25$.

As can be seen from the figure, timed data stream $((s_j, 20), (f_j, 38))$ satisfies both constraint sets while $((s_j, 8), (f_j, 32))$ satisfies only the relaxed deadline. In fact, the common area of the two feasible regions occupies $\frac{22}{25} = 88\%$ of that of the relaxed deadline 25ms.

Advancing to 3-dimensional feasible regions, consider the feasible region of the following timing constraint set that has three events:

$$\left\{ \begin{array}{ll} t(e_1) - t(e_2) \leq 5, & t(e_2) - t(e_1) \leq 7, \\ t(e_1) - t(e_3) \leq 5, & t(e_3) - t(e_1) \leq 2, \\ t(e_2) - t(e_3) \leq 10, & t(e_3) - t(e_2) \leq 5 \end{array} \right\} \quad (14)$$

The relationship between feasible regions satisfying constraint sets (1) and (14) is illustrated in Fig.4, where bold lines, light lines, and the shaded region represent constraint sets (1), (14), and the intersection between their feasible regions, respectively.

From Fig.4, we can see that although feasible regions satisfying constraint sets (1) and (14) are not identical, they share some common region. Hence, we can expect that they have some timing behaviors in common. \square

Generalizing the above discussions, we define the similarity between two timing constraint sets as the following:

Definition 3 (Constraint Set Similarity) *Let $S(C)$ denote the size of the feasible region of a timing constraint set C . Given two timing constraint sets C, C' , the similarity relation is defined as $C \sim C' = \frac{S(C \cap C')}{S(C')}$, where $C \cap C'$ is the intersection of C and C' . \square*




Figure 4: The feasible regions satisfying constraint sets (1) (bold lines) and (14) (light lines), and their intersection (the shaded region).

Intuitively, if $C \sim C' = P\%$, i.e., the intersection of the feasible regions of constraint sets C and C' occupies $P\%$ of the feasible region of C' , we know that $P\%$ of all the timed data streams satisfying C' satisfies C . Therefore, system satisfying C' will have a $P\%$ guarantee of satisfying C . Unfortunately, directly calculating the similarity between two sets of complete timing constraints is difficult. In fact, calculating the size of a polytope formed by a set of linear inequalities ($\mathcal{S}(C)$ in our context) has been shown to be $\#P$ -hard [?], and thus directly calculating the proportions of the intersection in any of the feasible regions, i.e., the similarity metric, by comparing their sizes is costly. To overcome the computational hurdle of evaluating directly the constraint set similarity between two constraint sets, we resort to finding a lower bound on the constraint set similarity such that it is easily computable and is tight. The following theorem defines such a bound.

Theorem 2 *Given two timing constraint sets C and C' , and corresponding normal forms be \mathbf{D}^* and \mathbf{D}'^* , respectively. If the feasible region of C' is not included in that of C , i.e., $\mathbf{D}^* \not\geq \mathbf{D}'^*$, then the similarity is bounded by:*

$$\left(\min_{\substack{i,j=1,\dots,n, \\ i \neq j, d_{i,j}^* \leq d_{i,j}'^*}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \right)^{|E|-1} \leq C \sim C' < 1 \quad (15)$$

where $|E|$ is the cardinality of the event set being constrained, $d_{i,j}^*$ and $d_{i,j}'^*$ are the corresponding entries in \mathbf{D}^* and \mathbf{D}'^* , respectively. The similarity reaches upper bound 1 when feasible region of C' is included in that of C , i.e., $\mathbf{D}^* \geq \mathbf{D}'^*$.

Proof Preliminary:

Before giving the formal proof for the theorem, we briefly introduce some background in computing volumes of high-dimension polytopes. In [?], Lawrence gives an algorithm for computing polytope volume based on a combinatorial form of Gram's relation. A convex polytope P is given as

$$P = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq \mathbf{0}, \mathbf{A}\mathbf{x} \leq \mathbf{b}\} \quad (16)$$

where \mathbf{A} is an $m \times n$ matrix and \mathbf{b} is a column vector in \mathbb{R}^m which has only nonnegative entries. If P is

a simple (or non-degenerate)² polytope, then the volume of P can be derived by extracting parameters from basic feasible tableaux for the following linear programming problem

$$\begin{aligned} & \text{maximize} && f(\mathbf{x}) \\ & \text{subject to} && \mathbf{Ax} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0} \end{aligned} \quad (17)$$

where $f(\mathbf{x}) = \mathbf{c}^t \mathbf{x} + \mathbf{d}$ can be any function that is not constant on any one of the hyperplanes defining P . The volume of P , denoted as $\text{vol}(P)$, is thus computed as

$$\text{vol}(P) = \sum_{\text{all vertices } v \text{ of } P} \frac{1}{n!} \frac{1}{\delta_v} \frac{\tilde{d}^n}{\gamma_{i_1} \cdots \gamma_{i_n}} \quad (18)$$

where δ_v is the cumulative product of the pivot elements, $\gamma_{i_1}, \dots, \gamma_{i_n}$ are the non-basic feasible solutions, and \tilde{d} is the value of the objective function. The pivot elements, non-basic feasible solutions, and the values of the objective function can all be retrieved from the basic feasible tableaux for (17).

To apply the volume computation algorithm in our setting, the following issues need to be addressed:

1. As can be seen in Fig. 2, the feasible region formed by the intersections of half spaces corresponding to timing constraints is not bounded. We need to find a way to bound the feasible region.
2. The volume computation algorithm crucially relies upon simplex-pivoting-based vertex enumeration algorithms. However, McMullen [?] has shown that the maximum number of vertices a polytope defined by m inequalities on n non-negative variables can have is $\binom{m+n}{n} + \binom{m+n}{n-1}$. Therefore, in (18), summing for ‘‘all vertices v of P ’’ generally has exponential cost. In fact, Dyer et al. [?] have shown that computing $\text{vol}(P)$ is $\#P$ -hard. Therefore, in our context, directly calculating the proportion of the intersection in any of the feasible regions by comparing the volumes is costly. We need to find a way to utilize the special properties of the feasible regions being compared.

The first issue can be addressed by ‘‘fixing’’ one of the n timers of events e_1, \dots, e_n . As the selection of the ‘‘fixed’’ one does not change the ratio of the intersection in any of the feasible regions, without loss of generality, we choose to fix $t(e_n)$ at $d = \max_{i=1, \dots, n-1} d_{n,i}^*$ (in fact, as can be seen from (2), d could be any value larger than $\max_{i=1, \dots, n-1} d_{n,i}^*$). The geometric interpretation for this is that the hyperplane $t(e_n) = d$ is used to ‘‘cut’’ the feasible region so that the resulting region is bounded in \mathfrak{R}^{n-1} .

Example 5 For example, in (14), if we let $t(e_3) = 5$, the constraint set becomes

$$\left\{ \begin{array}{l} t(e_1) - t(e_2) \leq 5, \quad t(e_2) - t(e_1) \leq 7, \\ t(e_1) \leq 10, \quad t(e_1) \geq 3, \quad t(e_2) \leq 15 \end{array} \right\} \quad (19)$$

which is illustrated as dash line segments in Fig. 5.

Since the volume computation algorithm requires that the origin in \mathfrak{R}^n is a vertex of the polytope³, we shift the feasible region to the origin (illustrated as the gray region in Fig. 5), the constraint set will become

$$\left\{ \begin{array}{l} t(e_1) - t(e_2) \leq 2, \quad t(e_2) - t(e_1) \leq 10, \\ t(e_1) \leq 7, \quad t(e_2) \leq 15 \end{array} \right\} \quad (20)$$

As can be easily seen from the figure, the area of the feasible region is 80. To gain some insights into the volume computation algorithm which will facilitate our proof of Theorem 2, we illustrate the derivation of the area using the algorithm as follows:

Adopting the volume computation algorithm on the bounded feasible region, we choose the objective function $f((t(e_1), t(e_2))^T) = t(e_1) + t(e_2)$ and have the initial tableau

$$\left[\begin{array}{cccccc|cccc} \boxed{I} & -1 & 1 & 0 & 0 & 0 & 2 & = & d_{3,1}^* & + & d_{1,2}^* & - & d_{3,2}^* \\ -1 & 1 & 0 & 1 & 0 & 0 & 10 & = & d_{2,2}^* & + & d_{2,1}^* & - & d_{3,1}^* \\ 1 & 0 & 0 & 0 & 1 & 0 & 7 & = & d_{3,1}^* & + & d_{1,3}^* & & \\ 0 & 1 & 0 & 0 & 0 & 1 & 15 & = & d_{3,2}^* & + & d_{2,3}^* & & \\ \hline -1 & -1 & 0 & 0 & 0 & 0 & 0 & & & & & & \end{array} \right] \quad (21)$$

²An n -dimensional simple (or non-degenerate) polytope is a polytope where each vertex is the intersection of exactly n hyperplanes (defined by n of the m inequalities in (16) with \leq replaced by $=$). This requirement is inherited from vertex enumeration algorithms used in the volume computation algorithm. Although feasible regions in this paper are sometimes not non-degenerate, the requirement can be dropped by perturbing the auxiliary hyperplanes by a very small amount and the same result holds.

³This requirement can be discarded by lexicographic techniques for handling primal degeneracy in linear programming.




Figure 5: Fixing $t(e_3)$ at 5 in (14) can be interpreted as using the plane $t(e_3) = 5$ to cut the feasible region of (14) in Fig. 4 and view the slice in the $\mathbf{x}_1\mathbf{x}_2$ ($t(e_1)t(e_2)$) plane.

where we have $n = 3 - 1 = 2$, $\delta_v = 1$ (the initial pivot element enclosed in a box in (21)), the non-basic feasible solutions are the two nonnegative numbers $(-1, -1)$ in the lower left partition of the tableau, and $\tilde{d} = 0$ (the lower right partition of the tableau, i.e., the value of the objective function at the current vertex $(0,0)$). Then the first element of the summation in (18) is $\frac{1}{2!} \frac{1}{1} \frac{0^2}{(-1)(-1)} = 0$. Continuing pivoting at $(t(e_1), t(e_2)) = (2, 0)$ using the 1st row, we have

$$\left[\begin{array}{cccccc|cccc} 1 & -1 & 1 & 0 & 0 & 0 & 2 & = & d_{3,1}^* & + & d_{1,2}^* & - & d_{3,2}^* \\ 0 & 0 & 1 & 1 & 0 & 0 & 12 & = & d_{1,2}^* & + & d_{2,1}^* & & \\ 0 & \boxed{1} & -1 & 0 & 1 & 0 & 5 & = & d_{1,3}^* & + & d_{3,2}^* & - & d_{1,2}^* \\ 0 & 1 & 0 & 0 & 0 & 1 & 15 & = & d_{3,2}^* & + & d_{2,3}^* & & \\ \hline 0 & -2 & 1 & 0 & 0 & 0 & 2 & = & d_{3,1}^* & + & d_{1,2}^* & - & d_{3,2}^* \end{array} \right] \quad (22)$$

and $\frac{1}{2!} \frac{1}{1} \frac{2^2}{(-2)(1)} = -1$. Pivoting at $(t(e_1), t(e_2)) = (7, 5)$ using the 3rd row, we have

$$\left[\begin{array}{cccccc|cccc} 1 & 0 & 0 & 0 & 1 & 0 & 7 & = & d_{1,3}^* & + & d_{3,1}^* & & \\ 0 & 0 & 1 & 1 & 0 & 0 & 12 & = & d_{1,2}^* & + & d_{2,1}^* & & \\ 0 & 1 & -1 & 0 & 1 & 0 & 5 & = & d_{1,3}^* & + & d_{3,2}^* & - & d_{1,2}^* \\ 0 & 0 & \boxed{1} & 0 & -1 & 1 & 10 & = & d_{1,2}^* & + & d_{2,3}^* & - & d_{1,3}^* \\ \hline 0 & 0 & -1 & 0 & 2 & 0 & 12 & = & 2d_{1,3}^* & + & d_{3,1}^* & + & d_{3,2}^* & - & d_{1,2}^* \end{array} \right] \quad (23)$$

and $\frac{1}{2!} \frac{1}{1} \frac{12^2}{(-1)(2)} = -36$. Pivoting at $(t(e_1), t(e_2)) = (7, 15)$ using the 4th row, we have

$$\left[\begin{array}{cccccc|cccc} 1 & 0 & 0 & 0 & 1 & 0 & 7 & = & d_{1,3}^* & + & d_{3,1}^* & & \\ 0 & 0 & 0 & 1 & \boxed{1} & -1 & 2 & = & d_{2,1}^* & + & d_{1,3}^* & - & d_{2,3}^* \\ 0 & 1 & 0 & 0 & 0 & 1 & 15 & = & d_{2,3}^* & + & d_{3,2}^* & & \\ 0 & 0 & 1 & 0 & -1 & 1 & 10 & = & d_{1,2}^* & + & d_{2,3}^* & - & d_{1,3}^* \\ \hline 0 & 0 & 0 & 0 & 1 & 1 & 22 & = & d_{1,3}^* & + & d_{3,1}^* & + & d_{2,3}^* & + & d_{3,2}^* \end{array} \right] \quad (24)$$

and $\frac{1}{2!} \frac{1}{1} \frac{22^2}{(1)(1)} = 242$. Pivoting at $(t(e_1), t(e_2)) = (5, 15)$ using the 2nd row, we have

$$\left[\begin{array}{cccccc|cccc} 1 & 0 & 0 & -1 & 0 & \boxed{1} & 5 & = & d_{2,3}^* & + & d_{3,1}^* & - & d_{2,1}^* \\ 0 & 0 & 0 & 1 & 1 & -1 & 2 & = & d_{2,1}^* & + & d_{1,3}^* & - & d_{2,3}^* \\ 0 & 1 & 0 & 0 & 0 & 1 & 15 & = & d_{2,3}^* & + & d_{3,2}^* & & \\ 0 & 0 & 1 & 1 & 0 & 0 & 12 & = & d_{2,1}^* & + & d_{1,2}^* & & \\ \hline 0 & 0 & 0 & -1 & 0 & 2 & 20 & = & 2d_{2,3}^* & + & d_{3,2}^* & + & d_{3,1}^* & - & d_{2,1}^* \end{array} \right] \quad (25)$$

and $\frac{1}{2!} \frac{1}{1} \frac{20^2}{(-1)(2)} = -100$. Pivoting at $(t(e_1), t(e_2)) = (0, 10)$ using the 1st row, we have

$$\left[\begin{array}{cccccc|cccc} 1 & 0 & 0 & -1 & 0 & 1 & 5 & = & d_{2,3}^* & + & d_{3,1}^* & - & d_{2,1}^* \\ 1 & 0 & 0 & 0 & 1 & 0 & 7 & = & d_{3,1}^* & + & d_{1,3}^* & & \\ -1 & 1 & 0 & \boxed{1} & 0 & 0 & 10 & = & d_{3,2}^* & + & d_{2,1}^* & - & d_{3,1}^* \\ 0 & 0 & 1 & 1 & 0 & 0 & 12 & = & d_{2,1}^* & + & d_{1,2}^* & & \\ \hline -2 & 0 & 0 & 1 & 0 & 0 & 10 & = & d_{3,2}^* & + & d_{2,1}^* & - & d_{3,1}^* \end{array} \right] \quad (26)$$

and $\frac{1}{2!} \frac{1}{1} \frac{10^2}{(-2)(1)} = -25$. Pivoting at $(t(e_1), t(e_2)) = (0, 0)$ using the 3rd row, we have

$$\left[\begin{array}{cccccc|cccc} 0 & 1 & 0 & 0 & 0 & 1 & 15 & = & d_{3,2}^* & + & d_{2,3}^* & & \\ 1 & 0 & 0 & 0 & 1 & 0 & 7 & = & d_{3,1}^* & + & d_{1,3}^* & & \\ -1 & 1 & 0 & 1 & 0 & 0 & 10 & = & d_{3,2}^* & + & d_{2,1}^* & - & d_{3,1}^* \\ 1 & -1 & 1 & 0 & 0 & 0 & 2 & = & d_{3,1}^* & + & d_{1,2}^* & - & d_{3,2}^* \\ \hline -1 & -1 & 0 & 0 & 0 & 0 & 0 & = & 0 & & & & \end{array} \right] \quad (27)$$

and $\frac{1}{2!} \frac{1}{1} \frac{0^2}{(-1)(-1)} = 0$. Therefore, from (18), the area of the slice is $0 - 1 - 36 + 242 - 100 - 25 + 0 = 80$, which conforms to our early observation. Similarly, the corresponding areas of the feasible region of (1) and its intersection with that of (14) are 112 and 73, respectively.

Note: As can be seen from (21) to (27), the linear combinations of $d_{i,j}^*$'s on the right side of each tableau correspond to cycles in the constraint graph of the timing constraint set. In fact, as shown in Provan's algorithm [?] for enumerating vertices of a polytope related to a network linear program, the hyperplanes of the polytope P adjacent to a vertex v is in one-to-one correspondence with simple cycles of a directed graph modified (with respect to v) from the directed graph defined by the network linear program. Therefore, although a network linear program is significantly simpler than the general linear program as in (17), the number of terms in the summation (18) is still generally exponential. Therefore, deriving exact similarity ratio is of exponential cost. \square

From Example 5, we have the following observations that are crucial in our proof of Theorem 2:

Observation 1 The pivoting operation is essentially a Gaussian elimination, thus the value of the objective function at any pivoting step must be a linear combination of $d_{i,j}^*$'s.

Observation 2 As the corresponding hyperplanes of different timing constraint sets are parallel, the pivoting sequences for enumeration vertices of the two feasible regions are the same regardless of the value of $d_{i,j}^*$'s. This property overcomes the second issues identified above.

We now prove Theorem 2 based on these observations.

Proof:

Given constraint sets C and C' whose normal forms are \mathbf{D}^* and \mathbf{D}'^* , respectively. Let C^\cap denote the intersection of C and C' and $\mathbf{D}^{\cap*}$ denote its normal form. We define a new constraint set C'' whose normal form \mathbf{D}''^* is

$$\mathbf{D}''^* = \sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \cdot \mathbf{D}^{\cap*} = \frac{1}{d_{r_{\inf}}(C, C')} \cdot \mathbf{D}^{\cap*} \quad (28)$$

Since the constraint graph of C^\cap comprises of edges from either C or C' , for any entry $d_{i,j}^{\cap*}$ of $\mathbf{D}^{\cap*}$, we have

$$d_{i,j}^{\cap*} = d_{i,k_1}^{(\prime)*} + d_{k_1,k_2}^{(\prime)*} + \dots + d_{k_{t-1},k_t}^{(\prime)*} + d_{k_t,j}^{(\prime)*} \quad (29)$$

where $d_{k_s,k_{s+1}}^{(\prime)*}$ denotes either $d_{k_s,k_{s+1}}^*$ or $d_{k_s,k_{s+1}}'^*$. Therefore, from (28), when $\sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} = \frac{1}{d_{r_{\inf}}(C, C')} \geq 1$, we have

$$d_{i,j}''^* = \left(d_{i,k_1}^{(\prime)*} + \dots + d_{k_t,j}^{(\prime)*} \right) \cdot \sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \quad (30)$$

If $d_{k_s,k_{s+1}}^{(\prime)*}$ is $d_{k_s,k_{s+1}}'^*$, since $\sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \geq \frac{d_{k_s,k_{s+1}}^*}{d_{k_s,k_{s+1}}'^*}$ we have

$$d_{k_s,k_{s+1}}^{(\prime)*} \cdot \sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \geq d_{k_s,k_{s+1}}^* \quad (31)$$

and if $d_{k_s,k_{s+1}}^{(\prime)*}$ is $d_{k_s,k_{s+1}}^*$, as $\sup_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'^*} \right\} \geq 1$, (31) also holds. Thus, from (30) and (31), we have

$$d_{i,j}''^* \geq d_{i,k_1}^* + d_{k_1,k_2}^* + \dots + d_{k_{t-1},k_t}^* + d_{k_t,j}^* \geq d_{i,j}^* \quad (32)$$

i.e., $\mathbf{D}''^* \geq \mathbf{D}^*$. Therefore, from Theorem 1, the feasible region of C is included within that of C'' and thus

$$\text{vol}(C'')/\text{vol}(C) \geq 1 \quad (33)$$

We now use the volume computation algorithm to calculate the ratio between $\text{vol}(C^\cap)$ and $\text{vol}(C'')$. From Observation 1, at each pivoting step, the value of the objective function is a linear combination of $d_{i,j}^*$'s, i.e., at the k 'th pivoting step, the values of the objective functions for $\text{vol}(C^\cap)$ and $\text{vol}(C'')$ are

$$\sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}^{\cap(k)} \cdot d_{i,j}^{\cap*} \quad \text{and} \quad \sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}''^{(k)} \cdot d_{i,j}''^* \quad (34)$$

respectively. From Observation 2, since the pivoting sequence for deriving $vol(C^\cap)$ and $vol(C'')$ are the same, we have

$$\forall i, j, k : a_{i,j}^{\cap(k)} = a_{i,j}''^{(k)} \quad (35)$$

Moreover, from (28), we have

$$\forall i, j = 1, \dots, n, i \neq j : \frac{d_{i,j}^{\cap*}}{d_{i,j}''*} = d_{r_{\text{inf}}} (C, C') \quad (36)$$

Therefore, from (35) and (36), we have

$$\begin{aligned} \frac{\sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}^{\cap(k)} \cdot d_{i,j}^{\cap*}}{\sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}''^{(k)} \cdot d_{i,j}''*} &= \frac{d_{1,2}^{\cap*}}{d_{1,2}''*} = \dots = \frac{d_{n,n-1}^{\cap*}}{d_{n,n-1}''*} \\ &= d_{r_{\text{inf}}} (C, C') \end{aligned} \quad (37)$$

From Observation 2, we also know that at each pivoting step, the cumulative products of pivoting elements and the non-basic feasible solutions for deriving $vol(C^\cap)$ and $vol(C'')$ are the same. Therefore, from (18) and (37), we have

$$\begin{aligned} \frac{vol(C^\cap)}{vol(C'')} &= \frac{\sum_{\forall v \in C^\cap} \frac{1}{n!} \frac{1}{\delta_v} \frac{\left(\sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}^{\cap(k)} \cdot d_{i,j}^{\cap*} \right)^{|E|-1}}{\gamma_{i_1} \dots \gamma_{i_n}}}{\sum_{\forall v \in C''} \frac{1}{n!} \frac{1}{\delta_v} \frac{\left(\sum_{\substack{i,j=1,\dots,n, \\ i \neq j}} a_{i,j}''^{(k)} \cdot d_{i,j}''* \right)^{|E|-1}}{\gamma_{i_1} \dots \gamma_{i_n}}} \\ &= (d_{r_{\text{inf}}} (C, C'))^{|E|-1} \end{aligned} \quad (38)$$

Finally, from (33) and (38), we have

$$\frac{vol(C^\cap)}{vol(C)} = \frac{vol(C^\cap)}{vol(C'')} \cdot \frac{vol(C'')}{vol(C)} \geq (d_{r_{\text{inf}}} (C, C'))^{|E|-1} \quad (39)$$

□

From Theorem 2, one can see that the similarity lower bound can be calculated easily once the normal forms of the constraint sets are available. Comparing similarities of different constraint set pairs then can be indirectly achieved through evaluating their similarity bounds. Before discussing various implications of using the similarity bound in Section 4.2, we demonstrate the use of Theorem 2 on the constraint sets given in Example 4. From Theorem 2, the ratio of the common region between (1) and (14) to the feasible region of (14) is bounded by $[\frac{36}{49}, 1]$ where $\frac{36}{49} = (\min\{\frac{6}{7}, \frac{9}{10}\})^{3-1}$. Therefore, assuming a uniform distribution of the event timing behavior in the feasible regions, Theorem 2 guarantees that at least $\frac{36}{49} = 73\%$ timed data streams that satisfy (14) also satisfy (1). This gives us a quantitative measure of the resemblance between systems constrained by (1) and (14), respectively. Actually, as shown in Example 5, the exact ratio of the common region between (1) and (14) to the feasible region of (14) is $\frac{73}{80} = 91.25\%$.

4.2 Discussions

Timed data stream distribution in the feasible region

In the above discussions, we assume that timed data streams are uniformly distributed in the feasible region of the constraint set. The bound given in Theorem 2 is based on the assumption. However, the definition of constraint feasible region similarities can be extended to non-uniform cases. For example, consider two 2-dimensional feasible regions of constraint sets $C = \{t(e_1) - t(e_2) \leq 5, t(e_2) - t(e_1) \leq 15\}$ and $C' = \{t(e_1) - t(e_2) \leq 15, t(e_2) - t(e_1) \leq 9\}$. Assuming timed data streams are not uniformly distributed in the regions, but are as shown in Fig. 6(a) and 6(b), respectively. Obviously, in order to

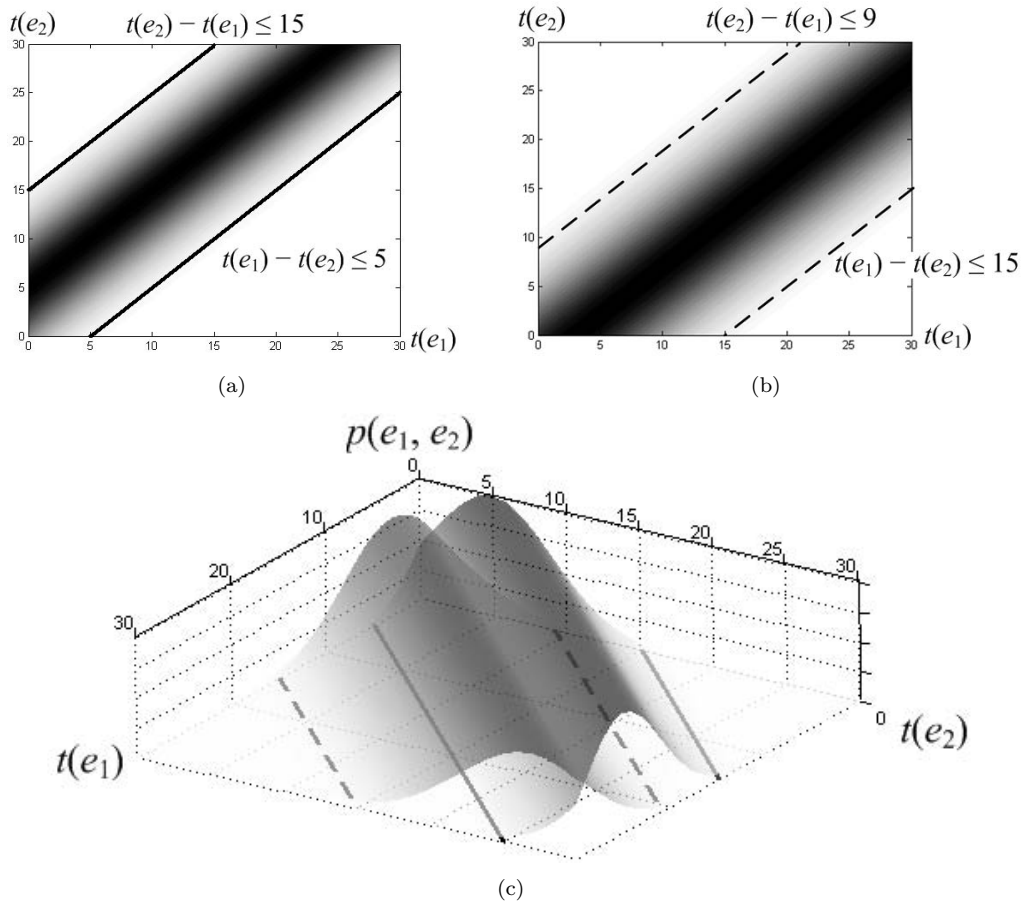


Figure 6: Feasible region similarities of non-uniformly distributed timed data streams.

compare their similarities, not only their areas but also the densities within the areas must be considered. For instance, the intersection of the feasible regions of C and C' is denser than the complements of the regions as depicted in Fig. 6(c). Therefore, the concept $\mathcal{S}(C)$ in Definition 3 are to be extended to weighted sizes.

In a soft real-time system, the distribution of timing values (such as the completion time of a task) can be evaluated by methods presented in existing work, e.g., [?, ?, ?, ?]. The distribution can then be used in combination with our proposed similarity bound concept to compare timing behaviors of different designs. The detail of this is beyond the scope of this paper.

Symmetry and transitivity of constraint set similarity

It is worth pointing out that the constraint set similarity relation is neither symmetric nor transitive. From Definition 3, it is not hard to see that in general $C \sim C' \neq C' \sim C$. For instance, for constraint sets $C = \{0 < t(f_j) - t(s_j) \leq 22\}$ and $C' = \{0 < t(f_j) - t(s_j) \leq 25\}$ as given in Example 4, $C \sim C' = 88\%$, while $C' \sim C = 1$.

Similarly, neither can we infer $C \sim C''$ from $C \sim C'$ and $C' \sim C''$. Figure 7 shows an example. In the figure, the feasible regions of three constraint sets C , C' , and C'' are represented as a tetragon, a pentagon, and a hexagon, respectively. The similarity between C and C' ($C \sim C'$) is the same for both figure Fig. 7(a) and Fig. 7(b). However, depending on the positions from which C'' similar to C' , C and C'' can be either very similar (as shown in Fig. 7(b)) or very dissimilar (as shown in Fig. 7(a)).




Figure 7: Similarity relation is not transitive.

The tightness of the similarity bound

From Theorem 2, it is easy to see that as the dimension of feasible regions gets higher, the similarities between their corresponding constraint sets decrease significantly due to the exponent $|E| - 1$. This is quite intuitive since, on one hand, as more events and constraints get involved, the chance of timed data streams satisfying one constraint set but violating the other gets bigger; on the other hand, from a geometric point of view, the volume of a polytope is exponential to its dimension, and so does the similarity between two polytopes.

Dealing with unconstrained event pairs in a constraint set

In Example 4, we illustrate the similarities between timing constraint sets where there is a constraint, either explicit or implicit, for every pair of events. However, there are cases where there are event pairs which are not constrained. For example, for constraint sets $C_1 = \{-5 \leq t(e_2) - t(e_1) \leq 22\}$ and $C_2 = \{t(e_2) - t(e_1) \leq 25\}$, the similarity $C_1 \sim C_2$ is close to 0 since in C_2 we implicitly have $t(e_1) - t(e_2) \leq +\infty$ and the feasible region is not bounded on the corresponding direction. In this case, the similarity relation stated in Theorem 2 still applies, but it approaches to 0 ($C_1 \sim C_2 = \min\left\{\frac{22}{25}, \frac{5}{+\infty}\right\} \rightarrow 0^+$), such 0 similarities render the metric too coarse. Hence, a refinement that considers unconstrained events is needed.

For most real-time applications, we observe that events often form groups such that those within the same group are pairwise constrained either explicitly or implicitly as shown in Section 4.1, and the timing relations between groups are either nonexistent or constrained by unidirectional constraints such as precedence constraints or delays. Therefore, given two timing constraint sets C and C' on the same set of events E , in order to take the unconstrained event pairs into consideration, we take the following steps

I. Partition E by strongly connected components of constraint graphs of C and C' . We only consider the case where both partitions are the same. It is not hard to see that each pair of events in a partition is explicitly or implicitly constrained.

II. Let E_1, \dots, E_K denote the K partitions and C_1, \dots, C_K and C'_1, \dots, C'_K denote the constraints of C and C' within the partitions, respectively. Then $C \sim C'$ is bounded by

$$C \sim C' \geq \min_{k=1, \dots, K} \{C_k \sim C'_k\} \quad (40)$$

$$\geq \min_{k=1, \dots, K} \left\{ \min_{\substack{i, j=1, \dots, n, \\ i \neq j, d_{k,i,j}^* \leq d_{k,i,j}'^*}} \left\{ \frac{d_{k,i,j}^*}{d_{k,i,j}'^*} \right\}^{|E_k|-1} \right\} \quad (41)$$

By partitioning events as well as the constraints among them, we reduce the dimensions of feasible regions of a constraint sets, filter out constraints that are irrelevant to the measurement of similarities, and thus get a more fine-grained view of similarities between the constraint sets.

We demonstrate the approach through a simple example. Consider vote-and-decide applications where several groups of voters vote within groups and a decision unit collects decisions from all groups. A typical constraint set constrains events within each voting group by relative deadlines to guarantee voting consistency and defines certain delays for the decision unit to make decision after all votes are collected. Figure 8 shows the timing constraint graphs of two timing constraint sets. According to strongly connected components, we partition the events into $E_1 = \{e_1, e_2, e_3\}$, $E_2 = \{e_4, e_5\}$, and $E_3 = \{e_6\}$, where partitions E_1 and E_2 are events from the corresponding voting groups, and partition E_3 is the deciding event. The similarity between the two sets of constraints, $C \sim C'$, is then lower bounded by $\min\{\frac{36}{49}, \frac{9}{13}, 1\} \approx 69\%$.




Figure 8: Similarity between general timing constraint sets.

5 Application 1: Predicting Tracking Error Rate Based on Constraint Similarity

In this section, we use a simplified real-time target tracking system to illustrate how a timing constraint set in a real-world setting may differ from the one under ideal assumptions. We further illustrate that despite partially-known and changing environments, the constraint similarity study given in Section 4

can be utilized to infer the tracking error rate and thus improve the predictability of the system in real-world environment.

Example 6 As an illustrative example, we consider a target tracking system presented in *Enviro-Track* [?]. To simplify our discussion, we assume a rectangular grid of sensors that periodically report to a control center their distances to a moving target. Each sensor s_i has a bounded unknown delay, denoted by $t_i \in [0, d]$, from sensing the data to reporting the data. A target that moves in the area is detected by the sensors and the final coordinates of the target are decided by the control center aggregating the sensed data from them. The system's QoS is measured by the tracking error rate which is the ratio of inaccurately reported data due to inconsistent data from sensors to the total number of data reported. The tracking error rate is determined by the sensor data freshness and sensor data consistency.

In order to have fresh data, it is desirable for the control center to have short decision times. In addition, the sensor data aggregated at the control center should belong to the same sampling period to minimize data inconsistency among different sensors. More specifically, let $t(e_i(k))$ denote the time that sensor s_i reports to the control center of its distance to the target in the k 'th detecting period, under the ideal assumption that a sensor locates the target at the beginning of the k 'th detecting period. We have

$$t(e_i(k)) = kT + t_i \quad (42)$$

where T is the detecting period for all sensors. Hence, for data consistency, we require $|t(e_i(k)) - t(e_j(k))| = |t_i - t_j| \leq d$, where $t_i, t_j \in [0, d]$. However, if we consider the signal transmission time, or possible objects in between the sensors and target, (42) should be changed into

$$t(e_i(k)) = kT + t_i + \delta_i(k) \quad (43)$$

where $\delta_i(k)$ is the time discrepancy caused by the Euclidean distance between the sensor s_i and the location of the target at the k 'th period. The constraint thus becomes $|t(e_i(k)) - t(e_j(k))| \leq d + |\delta_i(k) - \delta_j(k)|$.

To simplify our discussion, we restrict our attention to an $r \times r$ ($r = 1m$) square field with four ultrasonic sensors (whose detecting signal speed is $V = 340m/sec$) located on the corners of the field. The detecting radius of each sensor is assumed to be large enough to cover the entire field.




Figure 9: Four sensors tracking a moving target.

From time 0 and for every $T = 0.2sec$, all sensors try to measure their distances to a moving target, and after their distinct bounded unknown delays $t_i \in [0, d]$ ($d = 0.025sec$), the sensors report the distances to a control center which decides the coordinates of the target by aggregating the distance values from the four sensors. Under the ideal assumption that a sensor measures the distance at the beginning of the k 'th detecting period, the time that the new distance value for sensor s_i is available for reporting is $t(e_i(k)) = kT + t_i$ as given in (42). As mentioned earlier in Example 6, to guarantee the consistency of the reported data, the control center requires that the time distances between every two reporting events from two sensors, be bounded by $|t(e_i(k)) - t(e_j(k))| = |t_i - t_j| \leq d$. This results in a constraint set C and its constraint matrix \mathbf{D} is given in (44). In case of a constraint violation, i.e., the control center does not receive data from one of the sensor(s) before the corresponding deadline, the data received in previous periods will be used in the coordinates calculations⁴.

$$\mathbf{D} = \begin{bmatrix} 0 & d & d & d \\ d & 0 & d & d \\ d & d & 0 & d \\ d & d & d & 0 \end{bmatrix} \quad (44)$$

⁴Note that if the target is restricted to move only within the square field, the control center will only need two distance values in order to decide the coordinates of the target. However, although four sensors bring some redundancy, there can still be cases where less than two distance values come before deadlines. In this case, the control center takes the values received in previous periods for approximation.

However, if we take into account the travel time of ultrasonic distance measuring signals, the time of the event that sensor s_i reports the distance value is

$$t(e_i(k)) \approx kT + t_i + 2l_i/V \quad (45)$$

where l_i is the distance from sensor s_i to the target⁵. For instance, when the target appears at the same site as sensor s_1 , we have $l_1 = 0$, $l_2 = l_3 = r$ and $l_4 = \sqrt{2}r$, and the actual data consistency requirement on $t(e_i(k))$, $i = 1, \dots, 4$, becomes C' whose constraint matrix is $\mathbf{D}' = \mathbf{D} + \mathbf{\Delta}$, where

$$\mathbf{\Delta} = \begin{bmatrix} 0 & \frac{2r}{V} & \frac{2r}{V} & \frac{2\sqrt{2}r}{V} \\ -\frac{2r}{V} & 0 & 0 & \frac{2(\sqrt{2}-1)r}{V} \\ -\frac{2r}{V} & 0 & 0 & \frac{2(\sqrt{2}-1)r}{V} \\ -\frac{2\sqrt{2}r}{V} & -\frac{2(\sqrt{2}-1)r}{V} & -\frac{2(\sqrt{2}-1)r}{V} & 0 \end{bmatrix} \quad (46)$$

Therefore, if the control center uses constraint matrix (44) to monitor the events from the sensors, some timed data streams satisfying (44) may in fact correspond to inconsistent data, i.e., distances sensed in previous periods. Moreover, different locations of the target in the field result in different actual data consistency constraint sets (similar to (46)). It is thus difficult for the control center to adjust the data consistent constraints. \square

Given the data consistency constraint matrix \mathbf{D} in (44) under ideal assumptions, and a real-world deviation \mathbf{D}' as the one in (46), some timed data streams that satisfy C' may violate C . These violations cause the imprecise coordinates (ones that have blatant regressions due to data taken from previous periods). They are circled in Fig. 10(a), 10(c), and 10(e). From Theorem 2, the proportion of timed data streams satisfying C' that violate C is bounded by $[0, 1 - (d_{r_{\text{inf}}}(C', C))^{|E|-1}]$. Again, $1 - [d_{r_{\text{inf}}}(C', C)]^{|E|-1}$ has the largest value when the target is at one of the four corners of the field, where $d_{r_{\text{inf}}}(C', C) = \inf_{\substack{i,j=1,\dots,n, \\ i \neq j}} \left\{ \frac{d_{i,j}^*}{d_{i,j}'} \right\} = \frac{d}{d+2\sqrt{2}r/V} \approx \frac{3}{4}$. Therefore, we can estimate that the ratio of timed data streams satisfying C' that violate C to be bounded by $[0, 1 - (3/4)^{(4-1)}] = [0, 58\%]$. In fact, as can be observed from Fig. 10, approximately 3 (marked with gray circles) of the entire 8 reported coordinates in Fig. 10(a), 3 of the 9 coordinates in Fig. 10(c), and 6 of the 17 coordinates in Fig. 10(e) significantly deviate from the actual target location, indicating tracking error rates of 37.5%, 33.3%, and 35.3%, respectively.

It is worth pointing out that although the bound given in Theorem 2 may not be tight for constraint satisfaction ratio, it is a quick and easy way to get an estimation of constraint similarity and of how the system's behaviors in a real-world environment confirm or deviate from the initial design.

6 Application 2: Improving Systems' QoS Properties with Constraint Similarity Guarantees

The constraint similarity study is important as it has broad applications in areas where other types of QoS requirements, such as total energy consumption, are directly affected by a system's timing behaviors. As an example, we consider the energy-aware task assignment for soft real-time applications on a multiprocessor system-on-chip (MPSoC) which is similar to the one discussed in [?]. In particular, in this section, we will demonstrate (a) given the similarity metric and its bound (Section 4), calculate the probability guarantee that the original timing constraints are still satisfied by the modified constraint set for the purpose of reducing total energy consumption; and (b) given a maximum allowed constraint comprise, determine the constraint relaxations that best reduces energy consumption.

It is worth pointing out that reducing energy consumption is used only as an example to illustrate our approach. The similarity metric and the methodologies of using the metric to guide the trade-offs between timing and other QoS properties can be applied in a broad spectrum of soft real-time applications which involve timing and limited resources.

⁵We assume the speed of the moving target v is much smaller than the speed of the detecting signal V , i.e., $v \ll V$.

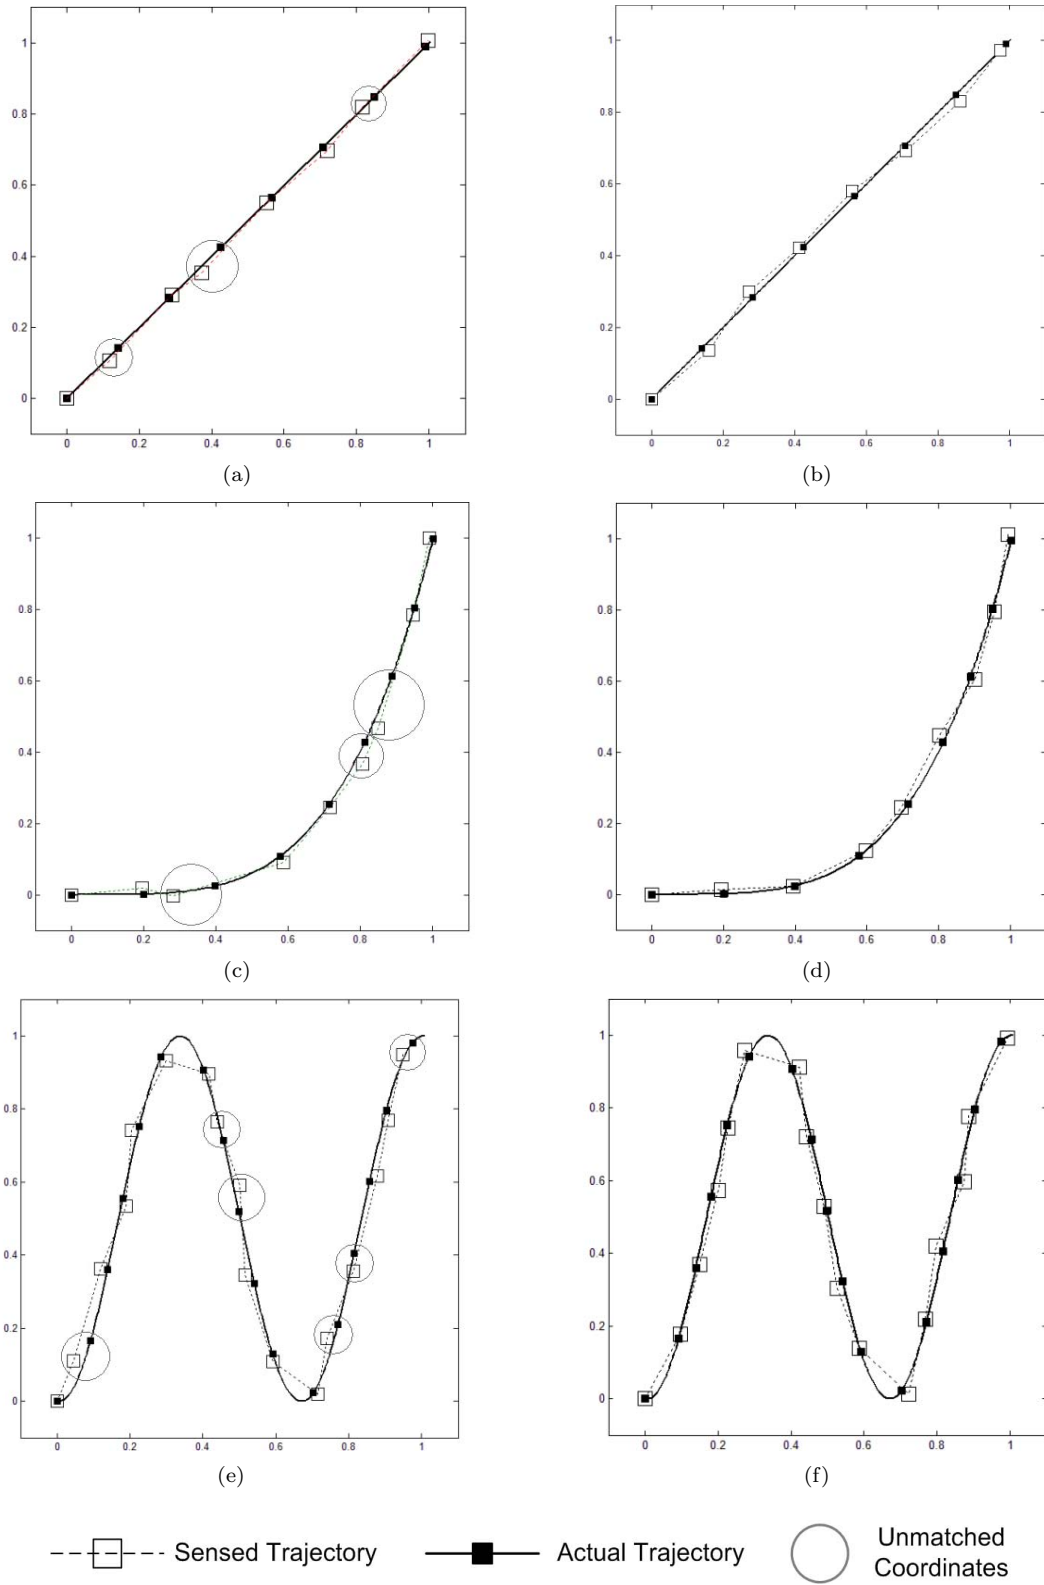


Figure 10: Actual trajectories and sensed coordinates of the moving target before and after the data consistency constraints are modified.

6.1 System and Task Model

The MPSoC under consideration consists of a set of heterogeneous cores M . Let J be the set of tasks to be executed on M . For each task $j \in J$, the following parameters are used in our discussions:

- $EX(j, m)$: j 's *worst-case* execution time on core m ,
- $ex(j, m)$: j 's *actual* execution time when running on core m , $ex(j, m) \in (0, EX(j, m)]$,
- d_j : the *relative* deadline of j ,
- s_j : the start event of task j ,
- f_j : the finish event of task j , $t(f_j) = t(s_j) + ex(j, m)$,
- $P(j, m)$: the power consumption of core $m \in M$ when task j executes on m .

The goal is to determine a static assignment of tasks to cores to further reduce the energy consumption while ensuring the required probability of constraint satisfactions guarantees. The hard real-time version of the problem, where a 100% deadline satisfaction must be ensured, is discussed in [?]. From the constraint satisfaction perspective, a deadline miss indicates that an execution trace falls outside of the feasible region defined by the given timing constraint set. When we allow a certain percentage of deadline misses, we actually include some execution traces outside the original feasible region, or in other words, the feasible region is expanded. The expanded feasible region can be considered as a relaxed constraint set. The constraint similarity study discussed in Section 4 allows us to quantitatively compare the deviations of the changed constraint from its original set, and hence to select which constraint(s) to relax based on a quantitative measure.

6.2 Reducing Total Energy Consumption

As shown in [?], the problem of minimizing total energy consumption for the MPSoC is to minimize $\sum_{j \in J} \sum_{m \in M} P(j, m) \cdot EX(j, m) \cdot \delta(j, m)$ where

$$\delta(j, m) = \begin{cases} 1 & \text{if } j \text{ is assigned to } m \\ 0 & \text{otherwise} \end{cases} \quad (47)$$

However, in our case, the actual execution time $ex(j, m)$ is not a constant value, and we assume it follows a certain probability distribution over the interval $(0, EX(j, m)]$. Therefore, the goal is to minimize the *expectation* of the total energy consumption and the objective function thus becomes minimizing $\sum_{j \in J} \sum_{m \in M} P(j, m) \cdot E[ex(j, m)] \cdot \delta(j, m)$.

Below, we demonstrate through an example how to use the similar bound to reduce total energy consumption by relaxing timing constraints.

Example 7 Consider two tasks j_1 and j_2 with relative deadline constraints $d_{j_1} = d_{j_2} = 20ms$ and synchronization constraints $|t(s_{j_1}) - t(s_{j_2})| \leq 5ms$. We thus have the following set of constraints:

$$\left\{ \begin{array}{ll} t(f_{j_1}) - t(s_{j_1}) \leq 20, & t(s_{j_1}) - t(s_{j_2}) \leq 5, \\ t(f_{j_2}) - t(s_{j_2}) \leq 20, & t(s_{j_2}) - t(s_{j_1}) \leq 5, \\ t(s_{j_1}) - t(f_{j_1}) \leq \epsilon, & t(s_{j_2}) - t(f_{j_2}) \leq \epsilon \end{array} \right\} \quad (48)$$

where $t(s_{j_1}) - t(f_{j_1}) \leq \epsilon (\epsilon \rightarrow 0^-)$ guarantees causality. The normal form of the constraint set (indexed by $t(s_{j_1}), t(f_{j_1}), t(s_{j_2}), t(f_{j_2})$) is given by (49).

$$\begin{bmatrix} 0 & \epsilon & 5 & 5 + \epsilon \\ 20 & 0 & 25 & 25 + \epsilon \\ 5 & 5 + \epsilon & 0 & \epsilon \\ 25 & 25 + \epsilon & 20 & 0 \end{bmatrix} \quad (49)$$

Now, consider the scheduling problem of task j_1 and j_2 on the following MPSoC with 4 cores:

$$\begin{array}{ccccc}
& & 10W & 10W & \\
20ms & \boxed{m_1} & & \boxed{m_2} & 20ms \\
22ms & \boxed{m_3} & & \boxed{m_4} & 25ms \\
& & 7W & 5W &
\end{array}$$

where $P(j_1, m_1) = P(j_2, m_1) = 10W$, $EX(j_1, m_1) = EX(j_2, m_1) = 20ms$, etc.

To satisfy the constraint set (48), j_1 and j_2 can only be assigned to m_1 and m_2 , respectively. Assuming the actual execution time is uniformly distributed in the interval $(0, Ex(j_1, m_1)]$, the expected total energy consumption is $10W \times 10ms + 10W \times 10ms = 200W \cdot ms$.

If we are willing to compromise the timing constraints, the deadline constraint of j_1 can be relaxed to $d_{j_1} = 22ms$ from $20ms$, the new constraint set becomes.

$$\left\{ \begin{array}{l} t(f_{j_1}) - t(s_{j_1}) \leq 22, \quad t(s_{j_1}) - t(s_{j_2}) \leq 5, \\ t(f_{j_2}) - t(s_{j_2}) \leq 20, \quad t(s_{j_2}) - t(s_{j_1}) \leq 5, \\ t(s_{j_1}) - t(f_{j_1}) \leq \epsilon, \quad t(s_{j_2}) - t(f_{j_2}) \leq \epsilon \end{array} \right\} \quad (50)$$

with normal form

$$\begin{bmatrix} 0 & \epsilon & 5 & 5 + \epsilon \\ 22 & 0 & 27 & 27 + \epsilon \\ 5 & 5 + \epsilon & 0 & \epsilon \\ 25 & 25 + \epsilon & 20 & 0 \end{bmatrix} \quad (51)$$

Based on Theorem 2, the similarity between these two constraint sets is lower-bounded by $(\frac{20}{22})^{4-1} \approx 75\%$. In other words, a system that satisfies the new constraints (50) has at least 75% guarantee of satisfying the initial system constraints (48). The benefit of relaxing the constraint is that we can now use m_3 to schedule j_1 or j_2 and the expected total energy consumption is thus reduced to $177W \cdot ms$, a 11.5% reduction.

Similarly, if we further relax the deadline constraint of j_2 to $d_{j_2} = 25ms$, one can easily verify that the similarity between the original and the modified constraint sets is bounded by $[51.2\%, 1]$ ($(\frac{20}{25})^{4-1} = 51.2\%$). In other words, systems that satisfy the modified constraints still have at least 50% chance to satisfy the original one. However, with such deadline relaxation, we can now schedule tasks j_1 and j_2 on m_3 and m_4 , respectively, with the corresponding expected total energy consumption reduced to $139.5W \cdot ms$, a 30.25% reduction.

Suppose we now have another job j_3 with a relative deadline of $22ms$. New constraints $t(f_{j_3}) - t(s_{j_3}) \leq 22ms$ and $t(s_{j_3}) - t(f_{j_3}) \leq \epsilon$ need to be inserted into (48). Since j_3 has no timing relations with j_1 and j_2 , based on Section 4.2, we partition the constraint set into two smaller normal forms.

$$\begin{bmatrix} 0 & \epsilon & 5 & 5 + \epsilon \\ 20 & 0 & 25 & 25 + \epsilon \\ 5 & 5 + \epsilon & 0 & \epsilon \\ 25 & 25 + \epsilon & 20 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & \epsilon \\ 22 & 0 \end{bmatrix} \quad (52)$$

For (52), the most energy-efficient assignment is to assign j_1 , j_2 , and j_3 to m_1 , m_2 , and m_3 , respectively, with a total expected energy consumption of $277W \cdot ms$. If the deadlines for j_1 and j_3 are reduced to $22ms$ and $25ms$, respectively, the corresponding normal forms are changed from (52) to (53)

$$\begin{bmatrix} 0 & \epsilon & 5 & 5 + \epsilon \\ 22 & 0 & 27 & 27 + \epsilon \\ 5 & 5 + \epsilon & 0 & \epsilon \\ 25 & 25 + \epsilon & 20 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & \epsilon \\ 25 & 0 \end{bmatrix} \quad (53)$$

We can then assign j_1 , j_2 , and j_3 to m_3 , m_2 , and m_4 , respectively, reducing the total expected energy consumption to $239.5W \cdot ms$, 14% reduction. The similarity between (52) and (53) is bounded by $\min \left\{ \left(\frac{20}{22}\right)^{4-1}, \left(\frac{22}{25}\right)^{2-1} \right\} \approx 75\%$. In other words, we have at least 75% guarantee to satisfy the initial constraints with the relaxed constraint set. \square

The above examples show that understanding the implication of constraint changes both from the system timing property and non-timing properties points of view plays a key role in conducting design tradeoffs. The similarity metric provides a quantitative measure about this implication in terms of timing constraint satisfaction. Specifically, the similarity bound between the original constraint set and that the modified one quantifies the maximal timing constraint satisfaction compromise in order to achieve certain desired QoS improvements. It thus allows us to make well-found decisions.

For the above examples, we manually picked some timing constraints to relax and calculated the similarity between the resultant constraint set and the original one. Under the same setting given in Example 7, a more interesting problem is: suppose we are allowed to relax the predefined constraints by certain amounts, can we determine which constraints to relax and how to relax them in order to find an assignment that further reduces expected total energy consumption?

6.3 Determining Constraint Relaxations

As we have seen from Section 6.2, relaxing timing constraints can further reduce total energy consumption, and Theorem 2 gives the bound of similarity between the modified constraint set and the original one. However, for real systems with a large number of events and constraints, there are possibly infinite ways even to relax a single timing constraint, not to mention there are combinatorial choices of constraints to relax. Therefore, relaxing constraints through exhaustive search is not realistic. Below, we consider one type of design problems and provide a systematic approach.

Given an application with both timing requirements and a desired QoS property, suppose that the design problem is formulated as an optimization of some QoS property under multiple types of constraints (including timing constraints). The goal is to find *appropriate* timing constraints and relax them to *appropriate* degrees so that the desired QoS property can be further improved while the initial timing constraints are still at least $P\%$ satisfied. We introduce the following steps for solving the problem.

Step 1: Based on given timing constraints, construct the corresponding timing constraint graph G . Partition G by strongly connected components. And for each strongly connected component, compute its normal form.

Step 2: Modify the original timing constraints such that each event pair of a constraint within a partition is constrained by a variable deadline (instead of the original deadline). Add new constraints to constrain the newly introduced deadline variables based on the specified similarity bound $P\%$.

Step 3: Solve the modified optimization problem using standard algorithms. The optimization solution contains the optimized value of the objective function which is the improved QoS property value, and the variable assignments which define the necessary timing constraint relaxations.

In the following, we illustrate the use of the above general steps through the example given in Section 6.2. More specifically, consider the specific example of assigning a set of five tasks j_1, \dots, j_5 to the MPSoC illustrated under the following timing constraints:

1. The relative deadlines of all tasks are $20ms$, i.e., $d_{j_1} = d_{j_2} = d_{j_3} = d_{j_4} = d_{j_5} = 20ms$;
2. There are synchronization constraints between j_1 and j_2 , and between j_3 and j_4 , i.e., $|t(s_{j_1}) - t(s_{j_2})| \leq 5ms$ and $|t(s_{j_3}) - t(s_{j_4})| \leq 5ms$;
3. Task j_3 and j_4 should start no later than $10ms$ after t_5 finishes, i.e., we have constraints $t(s_{j_3}) - t(f_{j_5}) \leq 10$ and $t(s_{j_4}) - t(f_{j_5}) \leq 10$.

Chantem et al. [?] formulate the problem as an MILP to optimize expected total energy consumption as following:

minimize

$$\sum_{j \in J} \sum_{m \in M} P(j, m) \cdot E[ex(j, m)] \cdot \delta(j, m) \quad (54)$$

subject to

$$\forall j \in J: \quad t(f_j) = t(s_j) + \sum_{m \in M} \delta(j, m) \cdot EX(j, m) \quad (55)$$

$$\forall j \in J: \quad \sum_{m \in M} \delta(j, m) = 1 \quad (56)$$

$$\forall e_i, e_j \in E: \quad t(e_i) - t(e_j) \leq d_{k_{i,j}} \quad (57)$$

where $E = \{s_j, f_j | j \in J\}$, and (57) generalizes timing constraints to a pairwise form ($d_{k_{i,j}}$ are constants obtained from the original constraints, for events that are not constrained, $d_{k_{i,j}} = +\infty$).⁶

⁶Note that the constraints to guarantee that all tasks execute for their durations without overlap [?] are omitted from the formulation for clarity of presentation.

Solving the MILP gives the non-preemptive schedule of tasks on the cores such that all timing constraints are met and the total energy consumption is minimized. Now, if we allow timing constraint relaxations but require a 75% constraint satisfaction guarantee, the original MILP needs to be modified based on the steps given above. In particular,

Step 1: For the constraint set given in (57), construct its corresponding constraint graph and partition the event set E into E_1, \dots, E_K based on the graph's strongly connected components. Only timing constraints *within* partitions are possible candidates for relaxations. Note that for any $j \in J$, s_j and f_j must be in the same partition since they are strongly connected by the relative deadline of j , i.e., $t(f_j) - t(s_j) \leq d_j$ and $t(s_j) - t(f_j) \leq \epsilon$. Therefore, all relative deadlines are possible to be relaxed.

For $\forall k = 1, \dots, K$, derive the constraint normal form \mathbf{D}_k^* for constraints among E_k , i.e., for $\forall e_i, e_j \in E_k$, $t(e_i) - t(e_j) \leq d_{k,i,j}^*$. For this example, we have partitions $E_1 = \{s_{j_1}, f_{j_1}, s_{j_2}, f_{j_2}\}$, $E_2 = \{s_{j_3}, f_{j_3}, s_{j_4}, f_{j_4}\}$, and $E_3 = \{s_{j_5}, f_{j_5}\}$. The constraint normal forms \mathbf{D}_1^* , \mathbf{D}_2^* , and \mathbf{D}_3^* on these partitions are

$$\mathbf{D}_1^* = \mathbf{D}_2^* = \begin{bmatrix} 0 & \epsilon & 5 & 5 + \epsilon \\ 20 & 0 & 25 & 25 + \epsilon \\ 5 & 5 + \epsilon & 0 & \epsilon \\ 25 & 25 + \epsilon & 20 & 0 \end{bmatrix}, \mathbf{D}_3^* = \begin{bmatrix} 0 & \epsilon \\ 20 & 0 \end{bmatrix} \quad (58)$$

respectively.

Step 2: For constraints within partitions, modify (57) in the MILP formulation to

$$\forall e_i, e_j \in E_k, k = 1, \dots, K : t(e_i) - t(e_j) \leq d'_{k,i,j} \quad (59)$$

$$\forall e_i, e_j \in E_k, k = 1, \dots, K : d'_{k,i,j} \leq \left\lfloor \frac{d_{k,i,j}^*}{|E_k|^{-1} \sqrt{P\%}} \right\rfloor \quad (60)$$

where $d'_{k,i,j}$ is the newly introduced variable for constraint relaxations. In the modified MILP, (59) and (60) are responsible for the selection and relaxation of constraints. From (60), we have

$$\left(\frac{d_{k,i,j}^*}{d'_{k,i,j}} \right)^{|E_k|-1} \geq \left(\frac{d_{k,i,j}^*}{d'_{k,i,j}} \right)^{|E_k|-1} \geq P\% \quad (61)$$

where $d_{k,i,j}^*$ is the corresponding entry in the normal form of the relaxed constraints and thus $d'_{k,i,j} \leq d_{k,i,j}^*$. According to Theorem 2 and Section 4.2, the probability that the system satisfying the relaxed constraint set also satisfies the original constraint set is no less than $P\%$.

For example, for constraint $t(s_{j_1}) - t(s_{j_3}) \leq 5$, we derive two constraints, i.e., $t(s_{j_1}) - t(s_{j_3}) \leq d'_{s_{j_1} s_{j_3}}$ and $d'_{s_{j_1} s_{j_3}} \leq \lfloor 5 / \sqrt[3]{0.75} \rfloor$; for constraint $t(s_{j_3}) - t(f_{j_5}) \leq 10$, since s_{j_3} and s_{j_5} belong to different partitions, the constraint is still in the modified MILP but cannot be relaxed. Specifically, (57) in the MILP is replaced by the following constraints

$$\begin{aligned} t(s_{j_1}) - t(f_{j_1}) &\leq d'_{s_{j_1} f_{j_1}} & , & \quad d'_{s_{j_1} f_{j_1}} \leq \lfloor \epsilon / \sqrt[3]{0.75} \rfloor & , \\ t(s_{j_1}) - t(s_{j_3}) &\leq d'_{s_{j_1} s_{j_3}} & , & \quad d'_{s_{j_1} s_{j_3}} \leq \lfloor 5 / \sqrt[3]{0.75} \rfloor & , \\ &\dots & & \dots & \\ t(f_{j_5}) - t(s_{j_5}) &\leq d'_{f_{j_5} s_{j_5}} & , & \quad d'_{f_{j_5} s_{j_5}} \leq \lfloor 20 / \sqrt[3]{0.75} \rfloor & , \\ t(s_{j_3}) - t(f_{j_5}) &\leq 10 & , & \quad t(s_{j_4}) - t(f_{j_5}) \leq 10 & \end{aligned}$$

Step 3: Solve the modified MILP using an MILP solver (such as ILOG CPLEX[®]). The solution contains the minimum expected total energy consumption and the assigned value of $d'_{k,i,j}$ which is the new constraint values in the correspondingly relaxed constraints. In this example, solving the modified instance of the MILP formulation, we have an optimal solution of $416.5W \cdot ms$, with $\delta(1, 1) = 1$, $\delta(2, 3) = 1$, $\delta(3, 2) = 1$, $\delta(4, 3) = 1$, and $\delta(5, 4) = 1$. The corresponding schedule is to run j_1 , j_2 , and j_5 on core m_1 , m_3 , and m_4 from time 0, respectively, with their new relative deadlines being $20ms$, $22ms$, and $26ms$, respectively. Since j_2 and j_4 are both assigned to core m_3 , to void overlap, from time $22ms$, j_3 and j_4 are scheduled to run on m_2 and m_3 , with their new relative deadlines being $20ms$ and $22ms$, respectively. The total execution time in this case is $44ms$ with all constrains satisfied. However, with the original MILP, we can only schedule all five tasks on m_1 and m_2 , with a minimum total execution time of $60ms$ and expected energy consumption of $500W \cdot ms$. Therefore, by compromising no more

than 25% of satisfaction guarantees of the original constraints, we gain a reduction of expected energy consumption and total execution time by 16.7% and 26.7%, respectively.

Through the above example of reducing total energy consumption with constraint similarity guarantees, we have demonstrated that when we do not require 100% constraint satisfaction guarantees, which is often the case for soft real-time applications, the flexibility allowed can be used to improve system's other QoS properties. We have further illustrate the detailed steps in obtaining better system QoS properties while still maintaining the required system's timing behavior resemblance.

7 Conclusion

Real-world real-time and embedded systems may behave differently from specification in the time domain: some systems deviate from the designs due to unpredictable factors; some other soft real-time systems allow certain timing flexibilities that can often be utilized to improve QoS properties of the systems. These deviations need to be exploited in a quantitative and predictable manner. Specifically, if a set of timing constraints are subject to imprecision or allowed to be modified, we need to measure how much the deviation differs from the origin set. Based on this need, in this paper, we introduce a quantitative metric to compare the similarity between two timing constraint sets. We based our study on feasible regions and proved that for a set of timing constraints, its feasible region is uniquely characterized by the constraint normal form. The similarity metric is then defined based on the common feasible region of the given two timing constraint sets, and reflects their mutual satisfactions. Since directly calculating the similarity metric is computationally intractable, we give a similarity bound based on the normal form. To demonstrate the capability of the new metric, we apply the theory of timing constraint similarities to an object tracking system for predicting tracking error rates; moreover, we use an MPSoC system to illustrate how we may use the similarity metric to guide the design phases for reducing system energy consumption. This example leads to a more general conclusion that the similarity metric between timing constraint sets can be used to guide the trade-offs between different QoS properties.

As future work, we plan to investigate the effect of non-uniformly distributed timed data streams on the evaluation of the similarity metric and its bound. Specifically, we will consider combining our earlier work on non-uniformly distributed interval-based events [?] with the computation of the similarity bounds. Intuitively, a set of interval-based events $\{I_1 = [\min(I_1), \max(I_1)], \dots, I_n = [\min(I_n), \max(I_n)]\}$, can be represented as a hypercube in the n -dimensional space whose density is determined by the joint distribution of all events. It will be revealing to understand the relationship between this hypercube with the hyperprism of a timing constraint set feasible region. This research is significant in deciding the satisfaction of timing constraints by events of a more practical model. Regarding the quality of the similarity bound, we realize that our bound may not be as tight, especially for higher dimension cases. We will further examine and improve the quality of the similarity bound.