

The E-T-C Scheduling Model - A new framework for the specification of Real-Time Scheduling Problems

K. Subramani¹

Department of Computer Science and Electrical Engineering,
West Virginia University,
Morgantown, WV, USA
ksmani@csee.wvu.edu

Abstract. Scheduling in Real-time systems differs from scheduling in conventional models in two principal ways: (a) Parameter variability, (b) Existence of complex constraints between jobs. Our focus has been on variable execution times. Whereas traditional models assume fixed values for job execution time, we model execution times of jobs through convex sets. The second feature unique to real-time systems, is the presence of temporal relationships that constrain job execution. Consider for instance the requirement that job 1 should conclude 10 units before job 2. This can be modeled through a simple, linear relationship, between the start and execution times of jobs 1 and 2. In real-time scheduling, it is important to guarantee a priori, the scheduling feasibility of the system. Depending upon the nature of the application involved, there are different schedulability specifications viz. Static, Co-Static and Parametric. Each specification comes with its own set of flexibility issues. In this paper, we present a framework that enables the specification of real-time scheduling problems and discuss the relationship between flexibility and complexity in the proposed model.

In this paper, we describe the features of our real-time scheduling framework called the E-T-C (Execution-Time-Constraints) Real-Time Scheduling model. The Scheduling model is composed of 3 sub-models, viz. the Job model, the Constraint model and the Query model. The Job model describes the type of jobs that we are interested in scheduling. The Constraint model is concerned with the nature of relationships constraining the execution of the jobs. The query model specifies what it means for a set of jobs to be schedulable, subject to constraints imposed as per the Constraint model. *An instance of a problem in the E-T-C model is specified by instantiating the variables in the sub-models.*

We focus on the following issues:

- (a) Designing a framework that enables specification of real-time scheduling problems, and
- (b) Studying instantiations of interest in this framework.

The rest of this paper is organized as follows: Section §1 describes the Job model within the E-T-C scheduling framework. The Constraint model is discussed in the succeeding section viz. Section §2. Section §3 details the Query model and presents the 3 types of queries that we consider in this thesis. A classification scheme for Scheduling

problems in the E-T-C model is introduced in §4. Section §6 provides a complexity picture of various instantiations in the E-T-C model. We conclude in Section §7 by summarizing the discussion in this paper.

1 Job Model in E-T-C

Assume an infinitely extending time axis, starting at time $t = 0$. This axis is divided into intervals of length L ; these intervals are ordered and each interval is called a scheduling window e.g. $[0, L]$ represents the first scheduling window, $[L, 2.L]$ represents the second scheduling window and in general, $[(i - 1).L, L]$ represents the i^{th} scheduling window. We are given a set of *ordered, non-preemptive* jobs $\mathcal{J} = \{J_1, J_2, \dots, J_n\}$, with start times $\{s_1, s_2, \dots, s_n\}$ and execution times $\{e_1, e_2, \dots, e_n\}$. L is the period of the job-set and all jobs execute periodically in each scheduling window.

2 Constraint Model in E-T-C

The executions of the jobs in the job-set \mathcal{J} (discussed in the above section) are constrained through relationships that exist between their start times and execution times. In the E-T-C model, we permit only linear relationships; thus the constraint system on the job-set is expressed in matrix form as :

$$\mathbf{A} \cdot [\vec{s}, \vec{e}] \leq \vec{b}, \quad (1)$$

where,

- $\vec{s} = [s_1, s_2, \dots, s_n]$ is an n -vector, representing the start times of the jobs;
- $\vec{e} = [e_1, e_2, \dots, e_n]$ is an n -vector representing the execution times of the jobs;
- \mathbf{A} is a $m \times 2.n$ matrix of rational numbers, called the *constraint matrix*;
- $\vec{b} = [b_1, b_2, \dots, b_m]$ is an m -vector of rational numbers,

Observe that System (1) can be rewritten in the form:

$$\mathbf{G} \cdot \vec{s} + \mathbf{H} \cdot \vec{e} \leq \vec{b}, \quad (2)$$

where,

$$\mathbf{G} \cdot \vec{s} + \mathbf{H} \cdot \vec{e} = \mathbf{A} \cdot [\vec{s}, \vec{e}]$$

We can also use the finish times f_i of jobs in relationships. Since the jobs are non-preemptive, the relation: $s_i + e_i = f_i$ holds for all jobs J_i and hence our expressiveness is not enhanced by the inclusion.

System (1) is a convex polyhedron in the $2.n$ dimensional space, spanned by the start time axes $\{s_1, s_2, \dots, s_n\}$ and the execution time axes $\{e_1, e_2, \dots, e_n\}$.

The execution times are independent of the start times of the jobs; however they may have complex interdependencies among themselves. This interdependence is expressed by setting

$$\vec{e} \in \mathbf{E} \quad (3)$$

where \mathbf{E} is an arbitrary convex set. We regard the execution times as n -vectors belonging to the set \mathbf{E} .

The ordering on the jobs is obtained by imposing the constraints:

$$s_i + e_i \leq s_{i+1}, \forall i = 1, \dots, n - 1.$$

The ordering constraints are included in the \mathbf{A} matrix in (1).

Goal: *We wish to determine a start time vector \vec{s} , in each scheduling window, such that the constraint system (1) holds (is not violated) at run-time for any execution time vector $\vec{e} \in \mathbf{E}$.*

The above specification (called the *schedulability specification*) is rather vague and is intended to be so; in Section §3 we shall present three different formalizations of the above specification. Each formalization is characterized by a distinct set of complexity and flexibility concerns. We also use the terms *schedulability query* and *schedulability predicate* to refer to the schedulability specification.

The Constraint model can be adapted to special situations by restricting either \mathbf{E} or \mathbf{A} or both. The following advantages result from such restrictions:

- A model that more accurately describes the requirements of the current situation,
- Faster algorithms for schedulability queries, and
- More efficient dispatching schemes.

In §2.1 we discuss a restriction to \mathbf{E} , while §2.2 and §2.3 deal with restrictions to the constraint matrix \mathbf{A} .

2.1 The Axis-parallel Hyper-rectangle domain

As specified above, the set \mathbf{E} in the Constraint model can be an arbitrary convex domain. One domain that finds wide applicability is the axis-parallel hyper-rectangle domain (henceforth abbreviated as aph). Figure (1) displays a 3-dimensional aph.

The Maruti Operating System [LTCA89,MAT90,MKAT92] estimates running times of jobs by performing repeated *runs* so as to determine upper and lower bounds on their execution time. Accordingly, the running time of job J_i , viz. e_i , belongs to the interval $[l_i, u_i]$, where l_i and u_i denote the lower and upper bounds on the execution time as determined by empirical observation. These independent range variations are the only constraints on the execution times. Observe that during actual execution, e_i can take any value in the range $[l_i, u_i]$.

The aph domain possesses two useful features:

- A specification that is tractable for this domain is also tractable for arbitrary convex domains,
- A specification that is provably “hard” for arbitrary convex domains is also “hard” for this domain.

Thus when proving complexity results (especially *hardness results*), it suffices to focus on the aph domain only.




Fig. 1. Axis-parallel Hyper-rectangle for 3 dimensions

2.2 Standard Constraints

The class of “standard constraints” was introduced in [Sak94], as a restriction to the constraint matrix \mathbf{A} for which the Parametric Schedulability query (see Section §3.3) could be decided efficiently.

Definition 1. A constraint is said to be a standard constraint, if it can be put in the following form:

$$a \cdot s_i + b \cdot e_i \leq c \cdot s_j + d \cdot e_j + k, \quad (4)$$

where,

- $\{a, b, c, d\} \in \{0, 1, -1\}$,
- $a = 0 \Rightarrow b = 0$
- $c = 0 \Rightarrow d = 0$
- $(b \neq 0) \Rightarrow b = a$,
- $(d \neq 0) \Rightarrow d = c$,
- $(a \neq 0) \wedge (c \neq 0) \Rightarrow a = c$
- $k \in \mathfrak{R}$ (For computational purposes, k is any rational number).

These constraints are also known as *monotone constraints* in the literature [HN94]. Standard constraints serve to model relative positioning requirements between two jobs and absolute constraints on a single job. When the constraints are standard, the matrix \mathbf{G} in System (2) is network unimodular [Dan63,NW88] and hence the constraint system can be represented as a network graph [SA00a,CLR92]. The following example serves to illustrate the relationship between standard constraints and network graphs.

Example 1. Consider a three job set $S = \{J_1, J_2, J_3\}$ with the following constraints imposed on the execution of the jobs:

- (a) Job J_1 finishes before J_2 starts: $s_1 + e_1 \leq s_2$
- (b) Job J_2 finishes before J_3 starts: $s_2 + e_2 \leq s_3$
- (c) Job J_2 commences within 4 units of J_1 concluding: $s_2 \leq s_1 + e_1 + 4$
- (d) Job J_3 finished before time 12: $s_3 + e_3 \leq 12$
- (e) $e_1 \in [2, 4], e_2 \in [1, 4], e_3 \in [2, 5]$.

We can represent the above constraint set in matrix form as:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} \leq \begin{bmatrix} 0 \\ 0 \\ 4 \\ 12 \end{bmatrix} \quad (5)$$

The dual of this constraint system is the network graph shown in Figure (2).




Fig. 2. Relationship between standard constraints and network graphs

The advantage of the network representation is that certain feasibility queries in the primal system can be expressed as shortest-path queries in the corresponding dual network [CLR92].

2.3 Network Constraints

Network constraints are a straightforward generalization of standard constraints.

Definition 2. A constraint is said to be a network constraint, if it can be put in the following form:

$$a \cdot s_i + b \cdot s_j \leq c \cdot e_i + d \cdot e_j + k, \quad (6)$$

where $a, b, c, d, k \in \mathbb{R}$.

Network constraints can also be represented as graphs [HN94,AS80]; however the relationships between adjacent vertices form a polyhedron and are not adequately represented through edges, as in the case of standard constraints. Once again, the advantage of the graph representation is the existence of faster algorithms for feasibility checking as opposed to general constraints.

3 Query model in E-T-C

In this section, we discuss three different schedulability specifications in the E-T-C model. Each specification has a different notion of what it means for a job-set to be schedulable. *However, in all the specifications the guarantees provided are absolute i.e. if the schedulability query is decided affirmatively, then the constraint set will not be violated at run time.*

3.1 Static Scheduling

Static scheduling is concerned with deciding the following predicate:

$$\mathcal{P}_f \equiv \exists \vec{s} = [s_1, s_2, \dots, s_n] \forall \vec{e} = [e_1, e_2, \dots, e_n] \in \mathbf{E} \quad \mathbf{A} \cdot [\vec{s}, \vec{e}] \leq \vec{\mathbf{b}} \quad ? \quad (7)$$

In other words, the goal is to determine the existence of a single start-time vector $\vec{s} \in \mathbb{R}^n$, such that the constraint system represented by (1) holds. The only information that is available prior to the dispatching of jobs in the i^{th} scheduling window is the knowledge of the execution time domain \mathbf{E} .

In [SA00c], we showed that the above proposition can be decided efficiently for arbitrary convex domains. From a computational perspective, query (7) is the easiest to answer.

3.2 Co-Static Scheduling

Static scheduling is unduly restrictive, in that even simple constraint sets will fail to have static schedules [SA00b]. The Co-Static scheduling query represents our first attempt to mitigate the inflexibility of static scheduling. In co-static scheduling, we ask whether there exists a start time vector corresponding to each execution time vectors $\vec{e} \in \mathbf{E}$. The start time vector may depend on the particular choice of \vec{e} . Accordingly, we wish to decide the following predicate:

$$\mathcal{P}_c \equiv \forall \vec{e} = [e_1, e_2, \dots, e_n] \in \mathbf{E} \quad \exists \vec{s} = [s_1, s_2, \dots, s_n] \quad \mathbf{A} \cdot [\vec{s}, \vec{e}] \leq \vec{\mathbf{b}} \quad ? \quad (8)$$

Co-static scheduling permits maximum flexibility during the dispatching phase, in that if a constraint system is not co-statically schedulable, then it is not schedulable. However, query (8) is coNP-complete for arbitrary constraint sets, as shown in [SA00b].

3.3 Parametric Scheduling

Co-static scheduling requires knowledge of the execution time vector for a particular scheduling window, prior to determining the start time vector for that window. This may not be feasible in all real-time systems. Parametric scheduling attempts to provide a balance between the Static and Co-Static scheduling modes. In a parametric schedule, the start time of a job is permitted to depend upon the start and execution times of jobs that have been sequenced before it *and only on those times*. In this mode, we restrict our discussion to aph domains, inasmuch as this simple domain preserves the hardness of schedulability queries. Thus, the parametric schedulability predicate is:

$$\mathcal{P}_{\checkmark} \equiv \exists s_1 \forall e_1 \in [l_1, u_1] \exists s_2 \forall e_2 \in [l_2, u_2] \dots \exists s_n \forall e_n \in [l_n, u_n] \mathbf{A} \cdot [\vec{s}, \vec{e}] \leq \vec{b} \quad ? \quad (9)$$

4 A Taxonomy of Scheduling problems

From the discussion in the above sections, it is clear that in order to specify an instance of a scheduling problem in the E-T-C scheduling framework, it is necessary to specify:

- The nature of the execution time domain (\mathbf{E}),
- The type of constraints on the jobs (\mathbf{A}), and
- A description of the schedulability query ($\mathcal{P}_f, \mathcal{P}_j, \mathcal{P}_{\checkmark}$).

Thus, a problem instance can be specified by instantiating the tuples in the $\langle \alpha | \beta | \gamma \rangle$ triplet, where,

- α represents the execution time domain \mathbf{E} - The following values are permissible for α :
 - aph - \mathbf{E} is an axis-parallel hyper-rectangle,
 - poly - \mathbf{E} is a polyhedron
 - arb - \mathbf{E} is an arbitrary convex domain.

Clearly aph is the weakest domain in terms of what can be specified and arb is the strongest; Figure (3) represents this relationship.

- β represents the constraint matrix $\mathbf{A}(\mathbf{G}, \mathbf{H})$ - β can assume the following values:
 - stan - The constraints are *standard* which implies that \mathbf{G} and \mathbf{H} are network, unimodular matrices.
 - net - The constraints are *network* which implies that \mathbf{G} and \mathbf{H} have at most two non-zero entries in any row
 - arb - \mathbf{G} and \mathbf{H} are an arbitrary $m \times n$ rational matrices

Once again stan is the weakest constraint class, in terms of real-time constraints that it can model, whereas arb is the strongest; Figure (4) represents this relationship.

- γ represents the schedulability predicate - The schedulability predicate specifies what it means for a set of jobs to be schedulable; the following values are permitted:
 - stat - The query is concerned with static schedulability,
 - co-stat - The query is concerned with co-static schedulability,




Fig. 3. Relationship between execution time domains




Fig. 4. Relationship between Constraint Classes

- param - The query is concerned with parametric schedulability.

Figure (5) represents the flexibility ordering between schedulability queries, with *co-stat* being the most flexible and *stat* being the least flexible.

Accordingly, $\langle aph|arb|stat \rangle$ represents an instance of a real-time scheduling problem, in which the execution time domain is an axis-parallel hyper-rectangle, the constraints are arbitrary and the schedulability predicate is static. Our notation scheme is similar to the $\langle \alpha|\beta|\gamma \rangle$ scheme for traditional scheduling models [Pin95,Bru81].

For the rest of this thesis, we use the term *domain* to refer to the execution time domain E and the term *constraint set* to refer to the constraint matrix $A(G, H)$.

5 Offline Analysis versus Online Dispatching

Scheduling algorithms in the E-T-C model possess an offline schedulability analyzer and an online dispatching component (See Figure (6)). The analyzer examines the constraints on the system and the type of schedulability query involved, to determine whether a feasible schedule is possible. *This analysis is always carried out offline.* The dispatching component is concerned with determining the exact start times of the jobs in the current scheduling window. *Dispatching is always carried out online.*




Fig. 5. Flexibility ordering between schedulability queries

For a given instance of a scheduling problem, the offline analyzer is executed exactly once. If the schedulability query is decided affirmatively, the online dispatcher is executed in every scheduling window.

6 Complexity of some instantiations

Figure (7) represents the complexity picture of some interesting instantiations as per our research.

7 Summary

In this paper, we presented a framework called the E-T-C scheduling model, for specifying scheduling problems in real-time systems. The E-T-C model possesses the following characteristics that distinguish it from a traditional scheduling model [Pin95]:

- Explicit accommodation of execution time interdependence - As specified, the execution time domain \mathbf{E} can be any arbitrary convex set. This enables us to model various situations such as execution time variability ($\alpha\phi$) and the existence of relationships between execution times.
- Unrestricted constraint matrix - As specified, the constraint matrix \mathbf{A} can be an arbitrary rational matrix. The generality of the constraint set permits the specification of relationships that cannot be expressed through precedence graphs or deadline and ready-time requirements.
- Decoupling of Schedulability analysis and Dispatching - This feature permits the analysis phase to be carried out offline, thereby reducing the online computation needed to determine the dispatch vectors in each scheduling window.

References

- [AS80] Bengt Aspvall and Yossi Shiloach. A fast algorithm for solving systems of linear equations with two variables per equation. *Linear Algebra and its Applications*, 34:117–124, 1980.

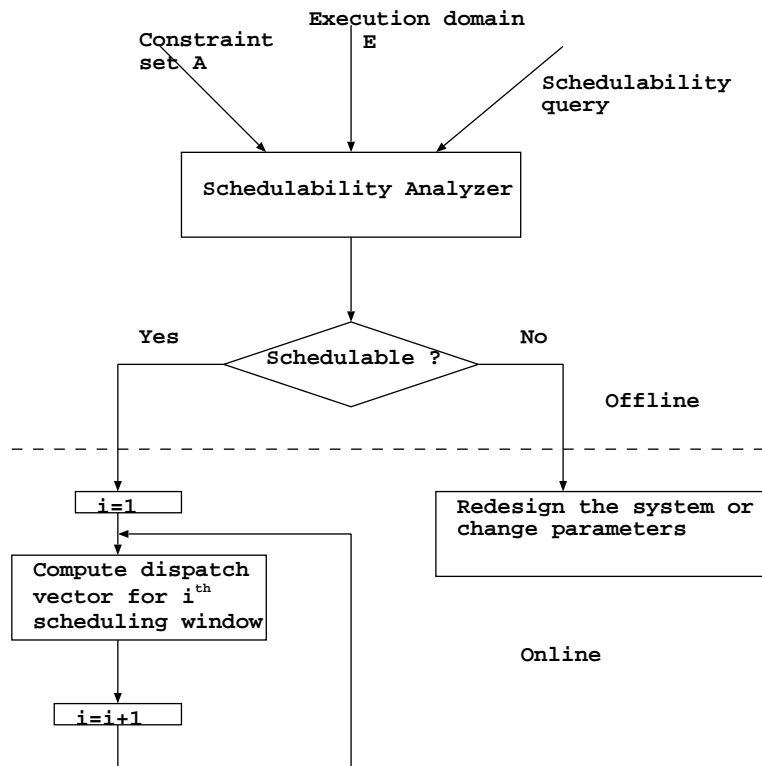


Fig. 6. Offline Analysis v/s Online Dispatching

- [Bru81] P. Brucker. *Scheduling*. Akademische Verlagsgesellschaft, Wiesbaden, 1981.
- [CLR92] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to algorithms*. MIT Press and McGraw-Hill Book Company, 6th edition, 1992.
- [Dan63] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ, 1963.
- [HN94] Dorit S. Hochbaum and Joseph (Seffi) Naor. Simple and fast algorithms for linear and integer programs with two variables per inequality. *SIAM Journal on Computing*, 23(6):1179–1192, December 1994.
- [LTCA89] S. T. Levi, S. K. Tripathi, S. D. Carson, and A. K. Agrawala. The Maruti Hard Real-Time Operating System. *ACM Special Interest Group on Operating Systems*, 23(3):90–106, July 1989.
- [MAT90] D. Mosse, Ashok K. Agrawala, and Satish K. Tripathi. Maruti a hard real-time operating system. In *Second IEEE Workshop on Experimental Distributed Systems*, pages 29–34. IEEE, 1990.
- [MKAT92] D. Mosse, Keng-Tai Ko, Ashok K. Agrawala, and Satish K. Tripathi. Maruti: An Environment for Hard Real-Time Applications. In Ashok K. Agrawala, Karen D. Gordon, and Phillip Hwang, editors, *Maruti OS*, pages 75–85. IOS Press, 1992.
- [NW88] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, New York, 1988.

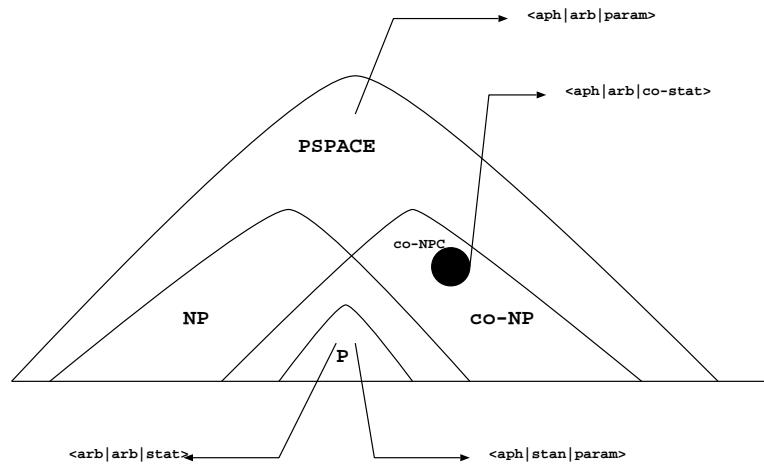


Fig. 7. Schedulability Complexity of some instantiations

- [Pin95] M. Pinedo. *Scheduling : theory, algorithms, and systems*. Prentice-Hall, Englewood Cliffs, 1995.
- [SA00a] K. Subramani and A. K. Agrawala. A dual interpretation of standard constraints in parametric scheduling. Technical Report CS-TR-4112, University of Maryland, College Park, Department of Computer Science, March 2000. Accepted at FTRTFT 2000.
- [SA00b] K. Subramani and A. K. Agrawala. The parametric polytope and its applications to a scheduling problem. Technical Report CS-TR-4116, University of Maryland, College Park, Department of Computer Science, March 2000. Submitted to the 7th International Conference on High Performance Computing (HIPC) 2000.
- [SA00c] K. Subramani and A. K. Agrawala. The static polytope and its applications to a scheduling problem. 3rd *IEEE Workshop on Factory Communications*, September 2000.
- [Sak94] Manas Saksena. *Parametric Scheduling in Hard Real-Time Systems*. PhD thesis, University of Maryland, College Park, June 1994.