Flexibility and Decoupling in the Simple Temporal Problem[1]

Michel Wilson[a] Tomas Klos[a] Cees Witteveen[a] Bob Huisman[a,b]

[a]*Delft University of Technology*, [b]*NedTrain, Utrecht*

Introduction

Scheduling problems occur in many diverse application domains such as transportation, process industry, health and education. A scheduling problem arises if we have some temporal variables and constraints between them, and we have to construct a *schedule* σ, an assignment of a value $\sigma(t)$ to each event t, that satisfies all constraints. A more robust approach in the face of uncertainty assigns to each event an *interval* in which it can start, so that we can quickly reschedule in case disturbances occur. In this case it is important to quantify how much flexibility such an 'interval schedule' offers. A second reason to be concerned about quantifying and optimizing schedule flexibility occurs in *multi-agent* scheduling, where the scheduling problem is decomposed into individual agents' subproblems. A *decoupling* procedure can be used to assure global consistency: Any combination of solutions to individual agents' subproblems is then also globally feasible. As pointed out by others [1, 2, 4, 5], when using existing flexibility metrics, this procedure can affect the total flexibility that can be achieved. Here we look at optimizing the decoupling with respect to flexibility and, if necessary, investigate the loss of flexibility due to decoupling.

We use the Simple Temporal Problem (STP) [3] as our framework for analyzing temporal scheduling problems. An instance of the STP is a pair $S = (T, C)$, where $T = \{t_1, \ldots, t_n\} \cup \{t_0\}$ is a set of time point variables (events) and C is a set of binary constraints on T, each having the form $t_j - t_i \leq c$ for some c. The time point t_0, often denoted by z, is added to express absolute time constraints and is assigned the value 0.

Flexibility in the STP

Flexibility refers to the freedom of choice we have in assigning values to events in T. A naive measure of one's flexibility in scheduling event t is $\mathit{flex}_N(t) = \mathit{lst}(t) - \mathit{est}(t)$, where $\mathit{est}(t)$ and $\mathit{lst}(t)$ are the earliest and latest starting times we can assign to t, respectively. Then we can define the total flexibility of instance S as $\mathit{flex}_N(S) = \sum_{t \in T} \mathit{flex}_N(t)$. For an *individual* event t, $\mathit{flex}_N(t)$ gives an exact measure of flexibility, since for any value $v \in [\mathit{est}(t), \mathit{lst}(t)]$, there exists a schedule σ such that $\sigma(t) = v$. A simple example shows, however, that flex_N overestimates the available flexibility for *combinations* of events:

Example 1. *Consider STP instance S_1, in which events t_1, t_2, and t_3 are to be scheduled in the interval $[0, 5]$. Now $\mathit{est}(t_i) = 0$ and $\mathit{lst}(t_i) = 5$ for all i, such that $\mathit{flex}_N(S_1) = 3 \cdot (5 - 0) = 15$. If we modify S_1 to instance S_2 by adding constraints specifying that $t_1 \leq t_2 \leq t_3$, we still have $\mathit{est}(t_i) = 0$ and $\mathit{lst}(t_i) = 5$ for all i, so $\mathit{flex}_N(S_2) = 15$ as well. But every solution for S_2 must sastisfy $0 \leq \mathit{lst}(t_1) \leq \mathit{lst}(t_2) \leq \mathit{lst}(t_3) = 5$. Therefore, the flexibility of t_1, t_2 and t_3 should rather be measured as $(\mathit{lst}(t_1) - 0)$, $(\mathit{lst}(t_2) - \mathit{lst}(t_1))$, and $(5 - \mathit{lst}(t_2))$, respectively, so the flexibility of S_2 should be $(\mathit{lst}(t_1) - 0) + (\mathit{lst}(t_2) - \mathit{lst}(t_1)) + (5 - \mathit{lst}(t_2)) = 5$. Thus, even though adding constraints between events in S_2 certainly affects the flexibility in S_2, this is not captured by the flex_N measure, which gives $\mathit{flex}_N(S_2) = \mathit{flex}_N(S_1) = 15$.*

[1]This is an abstract of our paper which received a distinguished paper award at the 2013 International Joint Conference on AI (IJCAI).

The failure of flex_N is due to the fact that it sums flexibility over time intervals which are interdependent, and several other flexibility measures proposed in the literature suffer similar flaws. A measure that more accurately computes the flexibility of a system S should be based upon 'uncorrelated' starting time intervals.

Definition 1. *Given an STP instance $S = (T, C)$, a set of intervals $I_S = \{[\ell_t, u_t]\}_{t \in T}$ is* uncorrelated *iff for every $t \in T$ and every $v_t \in [\ell_t, u_t]$, the assignment σ, given by $\sigma(t) = v_t$, is a schedule for S.*

To compute the flexibility inherent in a system $S = (T, C)$, we find a set I_S of uncorrelated starting time intervals $[\ell_t, u_t]$ for the events in T, that *maximizes* the sum of the intervals $\sum_{t \in T}(u_t - \ell_t)$. To find such a set, we make use of the following proposition. (Proofs for all our theorems are in the full paper.)

Proposition 1. *Let $S = (T, C)$ be an STP instance. A set of intervals $I_S = \{[\ell_i, u_i]\}_{t_i \in T}$ is uncorrelated if for every pair $(t_i, t_j) \in T^2$, it holds that if $(t_j - t_i \leq c) \in C$ then $u_j - \ell_i \leq c$.*

So the difference between t_j's upper bound and t_i's lower bound should not exceed c, if there is a constraint specifying that $t_j - t_i \leq c$. Using this proposition, we construct a special STP instance S' to find intervals for the events in S: The events in S' are the interval endpoints for the events in S.

Proposition 2. *Given STP instance $S = (T, C)$, consider instance $S' = (T', C')$, derived from S as follows.*

- $T' = \{t^-, t^+ \mid t \in T\} \cup \{z\}$ *(where $z^- = z^+ = z$);*
- $C' = \{t^+ - t'^- \leq c \mid t - t' \leq c \in C\} \cup \{t^- \leq t^+ \mid t \in T\}$.

Now for every solution σ for S', the set $\{[\sigma(t^-), \sigma(t^+)]\}_{t \in T}$ is an uncorrelated set of intervals for S.

To determine the flexibility of S, we find a solution σ for S' that maximizes the sum of the sizes of the intervals. The following LP formulation precisely models this task.

Theorem 1. *Given an STP instance $S = (T, C)$, $\mathit{flex}(S)$ can be computed by solving the following LP.*

$$\begin{array}{lll} \textit{maximize} & \sum_{t \in T} (t^+ - t^-) & \\ \textit{subject to} & t^- \leq t^+ & \forall\, t \in T \\ & t^+ - t'^- \leq c & \forall\, (t - t' \leq c) \in C. \end{array}$$

Decoupling without loss of flexibility

The STP can model distributed scheduling [4, 1, 2], by partitioning the set of events T across k agents $A_1, \ldots, A_k$. Each agent A_i wants to independently make a schedule σ_i, but the union $\sigma = \bigcup_{i=1}^{k} \sigma_i$ should be a schedule for the total instance. To ensure this, *decoupling* proposes tightening some intra-agent constraints so inter-agent constraints become implied. It is often supposed that this comes at the cost of some flexibility, but we show that we can efficiently find an optimal decoupling with no loss of flexibility.

Theorem 2. *Let $\{S_i\}_{i=1}^{k}$ be an optimal decoupling of an STP instance S. Then $\sum_{i=1}^{k} \mathit{flex}(S_i) = \mathit{flex}(S)$.*

References

[1] J.C. Boerkoel and E.H. Durfee. Distributed algorithms for solving the multiagent temporal decoupling problem. In *AAMAS*, 2011.

[2] A. Brambilla. *Artificial Intelligence in Space Systems: Coordination Through Problem Decoupling in Multi Agent Planning for Space Systems*. Lambert Academic Publishing, 2010.

[3] R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49:61–95, 1991.

[4] L. Hunsberger. Algorithms for a temporal decoupling problem in multi-agent planning. In *AAAI*, 2002.

[5] L. Planken, M. de Weerdt, and C. Witteveen. Optimal temporal decoupling in multiagent systems. In *AAMAS*, 2010.