Joint Chance Constrained Probabilistic Simple Temporal Networks via Column Generation (Extended Abstract)

Andrew Murray[1], **Michael Cashmore**[1], **Ashwin Arulselvan**[1], **Jeremy Frank**[2]

[1] University of Strathclyde, Glasgow, UK
[2] NASA AMES Research Center, Moffett Field, CA
{a.murray, michael.cashmore, ashwin.arulselvan}@strath.ac.uk, jeremy.d.frank@nasa.gov

Abstract

Probabilistic Simple Temporal Networks (PSTN) are used to represent scheduling problems under uncertainty. In a temporal network that is Strongly Controllable (SC) there exists a concrete schedule that is robust to any uncertainty. We solve the problem of determining Chance Constrained PSTN SC as a Joint Chance Constrained optimisation problem via column generation, lifting the usual assumptions of independence and Boole's inequality typically leveraged in PSTN literature. Our approach offers on average a 10 times reduction in cost versus previous methods.

Introduction

Probabilistic Simple Temporal Networks (PSTN) are networks made up of time-points (nodes) and constraints (edges) used to represent and reason over scheduling problems involving uncertain durations (Tsamardinos 2002). Strong Controllability (SC) (Vidal and Fargier 1999) is a property of temporal networks which implies the existence of a single schedule that is robust to any uncertainty in the problem. However in PSTNs this characteristic is rarely applicable due to the unbounded nature of continuous probability distributions.

SC can be imposed on a PSTN through *truncating* the probability distributions over durations (Fang, Yu, and Williams 2014; Santana et al. 2016). However, this increases the risk of the schedule. Another approach would be to *relax* some of the constraints (Yu, Fang, and Williams 2015) and pay a cost relative to the amount of relaxation. With the Chance Constrained PSTN (CC-PSTN) it is possible to minimize the relaxation cost subject to a user-defined bound on risk (Fang, Yu, and Williams 2014; Yu, Fang, and Williams 2015). However, previous approaches at tackling this problem use Boole's inequality, which is a loose upper bound on the true risk, and are therefore not guaranteed to find the optimal schedule minimising cost. Furthermore they assume independence of uncontrollable outcomes, which does not always hold. For example if we consider vehicle routing, then clearly a correlation exists between the travel times - driven by uncertainty factors such as traffic at a particular time of day. In this extended abstract we present an approach that overcomes both of these limitations based on established techniques from the field of convex optimisation.

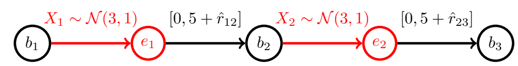

Figure 1: Relaxable CC-PSTN example. Uncertain durations are continuous probability distributions (X_1, X_2). Constraints are lower and upper bounds between time-points.

Motivating Example

We motivate our problem by considering the relaxable CC-PSTN introduced in Yu et al. (2015). In the example in Figure 1, we have two adjacent uncontrollable constraints $c(e_1, b_2)$ and $c(e_2, b_3)$, with relaxable upper bounds.

$$c(e_1, b_2) : 0 \leq b_2 - b_1 - X_1 \leq 5 + \hat{r}_{12}$$
$$c(e_2, e_3) : 0 \leq b_3 - b_2 - X_2 \leq 5 + \hat{r}_{23}$$

We consider that the decision maker has defined a risk bound of 0.3, the random variables X_1 and X_2 are independent and there is a schedule such that $b_2 - b_1 = b_3 - b_2 = 7$. Since the normally distributed random variables have identical distributions, we can write $X_1 = X_2 = X$. Finally, we note that the total relaxation cost is given by: $\hat{r}_{12} + \hat{r}_{23}$.

Boole's Inequality Using Boole's inequality the chance constraint is formulated as a sum of the two risks:

$$(1-P(2-\hat{r}_{12} \leq X \leq 7))+(1-P(2-\hat{r}_{23} \leq X \leq 7)) \leq 0.3$$

We can choose $r = \text{argmin}_{r \in \{\hat{r}_{12}, \hat{r}_{13}\}}\{1 - P(2 - r \leq X \leq 7)\}$ and rewrite the sum of probabilities as:

$$2(1 - P(2 - r \leq X \leq 7)) \leq 0.3$$

Therefore the cumulative probability function $F_X(2 - r) \leq F_X(7) - 0.85$ and subsequently $r \geq 0.04$, i.e. to satisfy the chance constraint we need to relax at least one of the bounds by at least 0.04 and therefore the cost, $\hat{r}_{12} + \hat{r}_{23} \geq 0.04$.

Joint Outcome Instead, if we consider the joint outcome of X_1 and X_2, then the chance constraint becomes:

$$1 - P(2 - \hat{r}_{12} \leq X \leq 7)P(2 - \hat{r}_{23} \leq X \leq 7) \leq 0.3$$

Without considering any relaxation, we have: $P(2 \leq X \leq 7) = 0.84$. Since $1-0.84\times0.84 \leq 0.3$, the chance constraint is satisfied with a cost of 0.

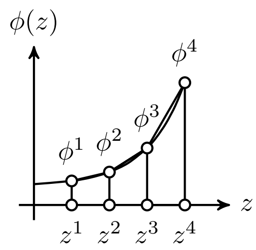

Figure 2: Inner approximation of a uni-variate function.

Method

To consider the joint outcome in our optimisation, we encode SC of CC-PSTNs as a Joint Chance Constrained (JCC) optimisation problem in the following form:

$$\min\{c^T x \mid Ax \leq l,\ P(\Lambda x + \beta \geq \Psi X) \geq 1 - \alpha\}$$

The decision variables, $x \in \mathcal{R}^n$ are the controllable time-points which can be scheduled and any problem specific variables (e.g relaxation variables in the relaxable CC-PSTN). Controllable constraints, containing only controllable time-points, are defined by the polyhedron, $Ax \leq l$. Uncontrollable constraints, containing a time-point whose outcome is dependant upon a random duration, are defined by the JCC, $P(\Lambda x + \beta \geq \Psi X) \geq 1 - \alpha$, in which $X \sim \mathcal{N}(\mu_X, \Sigma_X)$, is the multivariate normal vector with CDF $F_X(z)$ representing the uncertain durations and α is the allowable tolerance on risk.

The linearly transformed normal vector, $\xi = \Psi X$ is also a multivariate normal vector distributed according to: $\xi \sim \mathcal{N}(\mu_\xi, \Sigma_\xi)$, where $\mu_\xi = \Psi \mu_X$ and $\Sigma_\xi = \Psi \Sigma_X \Psi^T$. If the probability distribution is log-concave, then the chance constraint $\phi(z) \leq \pi$, where $\phi = -\log(F_\xi(z))$ and $\pi = -\log(1 - \alpha)$, is convex (Prékopa 1971, 1973) allowing for tractable evaluation of the global optimum. We can then express the optimisation problem as a convex one:

$$\min\{c^T x \mid Ax \leq l,\ z \leq \Lambda x + \beta,\ \Phi(z) \leq \pi\}$$

We follow the approach of Fábián et al. (Fábián et al. 2018; Fábián 2021) which forms a polyhedral approximation of the JCC using a restricted set of approximation points: $z^1, z^2, ..., z^k$ (see Figure 2). We iteratively solve the linear program (1), which minimises the cost given the current points; and the column generation problem (2), in which we find a new approximation point, z^{k+1} that refines the approximation. In the following, $u, v \leq 0$ and $\nu \in \mathcal{R}$ are the optimal dual variables associated with constraints $Z\lambda - \Lambda x \leq \beta$, $\Phi^T \lambda \leq \pi$ and $\mathbf{1}^T \lambda = 1$ respectively.

$$\{\min_{x,\lambda} c^T x \mid Ax \leq l,\ \Lambda x + \beta \geq \sum_{i=0}^{k} \lambda^i z^i,\ \sum_{i=0}^{k} \lambda^i = 1,$$

$$\lambda^i \geq 0,\ \sum_{i=0}^{k} \phi^i \lambda^i \leq \pi\} \quad (1)$$

$$\{\min_z -v\phi(z) - u^T z - \nu\} \quad (2)$$

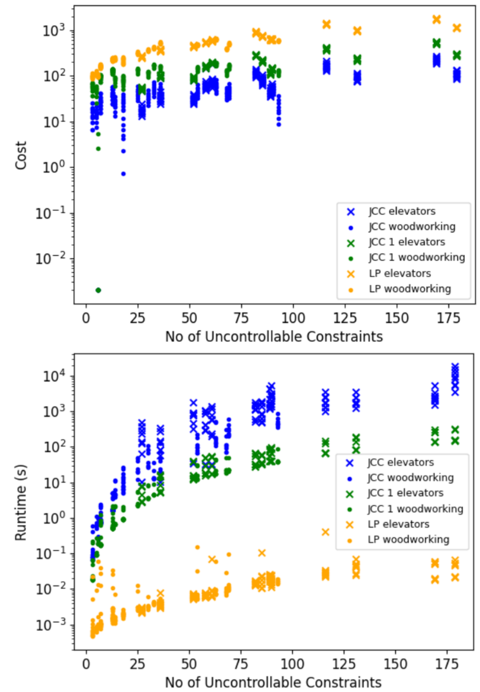

Figure 3: Results showing comparison of cost (top) and runtime (bottom).

We denote $\mathcal{M}$ and $\mathcal{C}$, the optimal objective to (1) and (2) on iteration k. From Dantzig (1963), $\mathcal{M} + \mathcal{C}$ is a valid lower bound and $\mathcal{M}$ is a valid upper bound on the optimal solution. If the difference between the upper and lower bound at iteration k is within some allowable tolerance, ε:

$$(\mathcal{M} - (\mathcal{M} + \mathcal{C}))/(\mathcal{M} + \mathcal{C}) \leq \varepsilon \quad (3)$$

then the solution so far is ε-optimal and the algorithm can be terminated.

Results and Conclusion

We tested our method against a SC linear program using Boole's inequality on a number of CC-PSTNs generated from IPC Planning domains. On average, the cost of our method is approximately 10 times less. However, the run time for the JCC approach was significantly longer for problems with a large number of probabilistic constraints - owing to the computational effort of evaluating multivariate normal probabilities and gradients within the column generation procedure (2). It's worth noting that (3) allows for a trade-off between the solution quality and runtime. We found that even after 1 iteration (presented as JCC 1 in Figure 3), the method offers improvements over the LP.

In conclusion, we recommend this approach for problems in which cost is of high priority and runtime is not restrictive. In future work we aim to investigate methods for improving efficiency. Furthermore, because our approach generalises to arbitrary covariance matrices, investigation into problems containing correlation is a logical next step.

References

Dantzig, G. B. 1963. *Linear programming and extensions*. Princeton: Princeton University Press.

Fábián, C. I. 2021. Gaining traction: on the convergence of an inner approximation scheme for probability maximization. *Central European Journal of Operations Research*, 29(2): 491–519.

Fábián, C. I.; Csizmás, E.; Drenyovszki, R.; van Ackooij, W.; Vajnai, T.; Kovács, L.; and Szántai, T. 2018. Probability maximization by inner approximation. *Acta Polytechnica Hungarica*, 15(1): 105–125.

Fang, C.; Yu, P.; and Williams, B. 2014. Chance-constrained probabilistic simple temporal problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28.

Prékopa, A. 1971. Logarithmic concave measures with applications to stochastic programming. *Acta Scientiarum Mathematicarum*, 32(1): 301–316.

Prékopa, A. 1973. On logarithmic concave measures and functions. *Acta Scientiarum Mathematicarum*, 34(1): 335–343.

Santana, P.; Vaquero, T.; Toledo, C.; Wang, A.; Fang, C.; and Williams, B. 2016. Paris: A polynomial-time, risk-sensitive scheduling algorithm for probabilistic simple temporal networks with uncertainty. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 26.

Tsamardinos, I. 2002. A probabilistic approach to robust execution of temporal plans with uncertainty. In *Hellenic Conference on Artificial Intelligence*, 97–108. Springer.

Vidal, T.; and Fargier, H. 1999. Handling contingency in temporal constraint networks: From consistency to controllabilities. *J. Exp. Theor. Artif. Intell.*, 11: 23–45.

Yu, P.; Fang, C.; and Williams, B. 2015. Resolving overconstrained probabilistic temporal problems through chance constraint relaxation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.