COMMENTARY: On "Strategies for Global Optimization of Temporal Preferences" by Morris, et al.

Mark S. Boddy
Adventium Labs
Minneapolis, MN
USA
mark.boddy@adventiumlabs.org

1 Introduction

This paper is the latest in an occasional series on the topic, written by various combinations of the current authors, sometimes in collaboration with others. The overall aim of the work is to provide an effective means for formulating and solving problems in temporal reasoning involving preferences expressed as *soft constraints*, imposing a cost on deviating from a preferred value for a given temporal variable. The framework is quite general, the results so far are fairly preliminary. In this commentary, I will briefly summarize the different definitions of global preference, comment on the new results presented, and discuss some further places to take this work, including the ones provided by the authors of the paper.

2 Preference Models

A central issue in the paper is defining precisely what kind of global preference is to be applied. The authors define four, three in previous work and one introduced here. The three previous global preference criteria are:

Weakest Link in which the least-preferred time is maximized. As an optimization problem, this can be described as $\max_{\mathcal{S}}(\min_{S}(x))$, where $\mathcal{S}$ is the set of all solutions, and $S \in \mathcal{S}$.

Pareto in which the preference values of the individual constraints are compared independently, consistent with the usual definition of Pareto optimality.

Utilitarian in which constraint values are summed to compute the global value.

The fourth preference model, introduced in this paper, is "*stratified egalitarian* optimality" which characterizes a subset of Pareto-optimal solutions. An algorithm is given for computing solutions that are optimal according to this model.

The subset of Pareto-optimal solutions computed by this algorithm can be intuitively described as the maximally preferred solutions according to a specific ordering on preference constraints. Solutions which have better values for earlier constraints in the ordering will be preferred. The ordering involved, termed "WLO ordering" is derived in the solution process used in the algorithm WLO+, defined in the paper, and is loosely related to a sort by ascending order for the preference values for each constraint achievable in a solution to the STPP.[1]

This subset of the Pareto-optimal solutions is described in the paper as being useful in some way beyond their being computed by a tractable algorithm, but how it is useful is not described.

3 Real Applications

There is more general point here, with regard to the paper as a whole: while the authors state their intent to address real, or at least realistic, problems, the preference models they define are only loosely tied back to real problems. For example, the real breadth of what can be expressed with STPPs is not really explored. This may well be because the authors are so familiar with their own constructs that to them it is obvious, but several times in the course of reading

[1] Simple Temporal Problem with Preferences

the paper, I would come to a tentative conclusion that some form of preference could not be represented, only to realize after more thought that, in fact, it could.

The running example of the rover is helpful with the basic mechanics of STPs and STPPs, but so simple as to not provide useful intuitions regarding the properties of and tradeoffs among the four different preference models in practical use. For example, a careless reader might get the impression from the discussion of Pareto-optimal solutions to the rover problem that in general Pareto-optimality is a sufficient critierion for minimizing summed durations (CPU operation, in that example), which it is not.

For another, the additional breadth possible with the introduction of a linear solver is not addressed. Having established that the STPP with a Utilitarian preference model can be expressed as a linear program, it would be interesting to go on and talk about how much farther a linear solver can take you. For example, it becomes possible to talk about functional dependence of one duration on another, or of a duration on the value assigned to a time point, or in some cases to relate durations and time-point variable assignments to other quantities, such as resource usage. Weakest Link Optimization can be expressed as a linear program, as well.

All of these are generalizations that can be accomplished strictly within a linear model, which extend the application substantially in the direction of being able to address real problems (the rover domain cited in the paper, to pick one instance, has constraints on power and energy that interact in interesting ways with when and how fast the rover moves, and potentially even with in which direction).

4 Presentation Issues

The paper has a few issues with notation and semantics. For example on page two, the discussion of the Simple Temporal Problem (STP) and Simple Temporal Problem with Preferences (STPP) hops back and forth between the conventional definition of an STP as a graph of variables representing time points, with labeled arcs providing bounds on the difference in variable values, and an STPP defined by the pair (V, C), in which V is a set of variables ranging over temporal *distances*, thus defining something like the dual of a fully-connected STP.

The definition of SE-domination is sloppy with respect to notation, leaving variables free that should be explicitly quantified. One way to rewrite the definition would be the following:

- $\exists i \| U_S^i < x$
- $\forall i, (U_S^i < x) \rightarrow (U_{S'}^i \geq U_S^i)$
- $U_S^i \geq x \rightarrow U_{S'}^i \geq x$

Finally, the definition of the linear program in Section 4 is not quite right. In particular, line 3 should read $Z_{ij} \leq f_{ij}^k(S_{ij})$. Also, the degree to which this is actually a linear program depends on the form of f, so line 3 might better be written as something like: $Z_{ij} \leq f_{ij}^k * S_{ij}$.

5 Summary and Conclusion

This paper is the latest in a series of papers addressing the issue of defining and solving for global preferences over temporal constraints. The authors present some new results, in particular a more precise characterization of a subset of Pareto-optimal solutions for which they have a tractable algorithm, and describe a (fairly obvious) encoding of Utilitarian optimization of an STPP as a linear program.

There is clearly a lot to be done, as of yet. In some ways, the most interesting aspect of this line of work is in the generation and analysis of non-Utilitarian global preference models. It would be interesting to see more of them, and in particular to see more discussion of where these models might be useful in applications. Are there places where Weakest Link models usefully apply? What can be done with the authors' insight that Stratified Egalitarianism can be viewed as a spectrum between Weakest Link and Pareto optimality?

It will also be interesting to see the results of the ongoing comparison between WLO+ and Utilitarian solutions described in the Conclusion.