

PhD in Computer Science

School of Computing, Informatics, and Decision Systems Engineering

Title:

When is Temporal Planning *Really* Temporal?

Ph.D. Defense:
William Cushing

Friday, November 16, 2012 - 10:00 am
BYENG 528

Committee:

Prof. Subbarao Kambhampati (*Chair*)
Prof. Chitta Baral
Prof. Hasan Davalcu
Dr. David E. Smith
Prof. Daniel S. Weld

Dissertation Abstract

The *simplest* temporal extensions of Classical Planners have swept clean the Temporal Planning tracks of the International Planning Competitions, from their inception in 2002 through today. As a practical matter the data is great news. We should always celebrate being able to perform well on hard problems by using only simple-minded approaches. To be realistic though: the result is *too good to be true*. Broadly speaking, my aim is to separate the fact from the fiction.

In other words, the aim is to better understand the computational relationship, in theory and practice, between Classical and Temporal Planning. A key notion is that, while *theoretical expressiveness* and *practical efficiency* are generally opposed, there are opportunities for exceptions. Such exceptions are well worth discovering and documenting in detail. To be more specific, the thesis I defend is as follows.

We should understand the precise interpretation of *concurrency* as a crucial feature separating the more from less expressive forms of Temporal Planning. Especially I call attention to three general approaches to formalizing: in *Sequential Planning* concurrency is forbidden, in *Conservative Temporal Planning* concurrency is strictly optional, and in *Interleaved Temporal Planning* concurrency is requirable.