

## Time-situated reasoning within tight deadlines and realistic space and computation bounds

Madhura Nirkhe

Dept. of EE and Dept. of CS

University of Maryland, College Park, MD 20742

madhura@cs.umd.edu

advisor: Donald Perlis (perlis@cs.umd.edu)

collaborator: Sarit Kraus (sarit@bimacs.cs.biu.ac.il)

We develop an effective representational and inferential framework for *fully* deadline-coupled, time-situated problem solving. Our effort is to model an agent in a tight and rigid deadline situation, in need of successfully formulating and executing a deadline-feasible plan of action as the world around the agent continues to change. We highlight the severe time-pressure under which the agent must operate with a paradigmatic problem scenario: Nell & Dudley and the railroad tracks. Nell is tied to the railroad tracks as a train approaches. Dudley, our agent, must formulate a plan to save her and carry it out before the oncoming train reaches her. He must deliberate (plan) in order to decide this, yet as he does so, the train draws nearer to Nell.

While meta-planning is the usual proposal for reasoning about the reasoning process, few formalisms acknowledge that it takes time, and none to date account for *all* the time spent in the reasoning within the same framework. We cite here some key related works: (Dean & Boddy 1988) formulate an algorithmic approach to time-dependent planning problems by introducing “anytime algorithms” which capture the notion that utility is a monotonic function of deliberation time. (Pollack & Ringuette 1990) explore the relation between agent design and environmental factors. (Rosenschein & Kaelbling 1989) provide a situated automata theory highlighting the role of logic, complexity and information in situated agents. (Russell & Wefald 1991) offer an optimal design for a limited rational agent using utility-based search. Our *fully* deadline-coupled planner has an important qualification that these efforts fail to meet: in addition to determining the current time, estimating the expected execution time of partially completed plans and being able to discard alternatives that are deadline-infeasible, it also has a built-in way of accounting for all the time spent as a deadline approaches. The underlying framework is that of Step-logics (Elgot-Drapkin & Perlis 1990) (now renamed active logics) which is a mechanism for reasoning situated in time.

We have demonstrated the generality of the formal methods employed by our formalism by solving

some real-time versions of canonical temporal projection problems such as the Yale Shooting Problem. While an agent under severe time-pressure may spend substantial amount of the available time in reasoning toward and about a plan of action, in a realistic setting, the same agent must also measure up to two other crucial resource limitations as well, namely space and computation bounds. We address these concerns and offer improvements by introducing a limited short term memory combined with a primitive relevance mechanism and a limited capacity inference engine. We propose heuristics to maximize the agent’s chances of meeting the deadline in this enhanced framework with additional space and computation constraints.

We provide a modal semantics for active logics which serves as a link between active logics and existing logic approaches that deal with time. Logical omniscience, and in particular inferential closure is a computational impossibility for an agent in the real-world. We have constructed a variation on active-logics for which there is a sound and complete modal semantics. It overcomes the key obstacle of closure under consequence, and restricts closure under *valid* consequence. It illustrates important comparisons between active logic work and previous modal approaches to knowledge and belief.

### References

- Dean, T., and Boddy, M. 1988. An analysis of time-dependent planning. In *Proceedings, AAAI-88*, 49–54.
- Elgot-Drapkin, J., and Perlis, D. 1990. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence* 2(1):75–98.
- Pollack, M. E., and Ringuette, M. 1990. Introducing the tileworld: Experimentally evaluating agent architectures. In *Proceedings, AAAI-90*, 183–189.
- Rosenschein, S., and Kaelbling, L. 1989. Integrating planning and reactive control. In *Proceedings of NASA Telerobotics conference*. Pasadena, CA.
- Russell, S., and Wefald, E. 1991. *Do the right thing*. The MIT press, Cambridge, Mass.