Interactive Planning and Monitoring

Gary Borchardt

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: Planning and monitoring are interdependent tasks that arise in a wide range of contexts, from small-scale corporate purchasing efforts to large-scale civil engineering projects. These tasks are also complex and difficult, involving many subtasks. The question is: how best should computer power be applied to assist humans in accomplishing planning and monitoring?

Motivation: Humans and computers exhibit broadly different strengths in performing reasoning tasks. For example, humans often excel when insight and judgment are required, while computers often excel when the focus is on scanning and testing of information sets. Both of these broad categories of tasks arise during planning and monitoring, and thus a good approach would seem to be one of constructing an interactive, shared-initiative environment for planning and monitoring—one in which humans and computers can each contribute according to their relative strengths.

Previous Work: Traditional artificial intelligence approaches have placed computers at the center of planning and monitoring activity [4], while traditional management science approaches such as software-supported GANTT and PERT charts have placed humans at the center of this activity [2]. In each case, the non-central participant, human or computer, is relegated to a significantly diminished role. More recently, mixed-initiative planning systems from artificial intelligence and decision-support systems from management science have attempted to increase the degree of human-computer collaboration during planning and monitoring [3, 2]. However, the underlying representations that are used in these approaches—symbolic, numerical, text-based, and so forth—can cause unnatural divisions of labor due to the relative incomprehensibility of these representations to either human or computer participants.

Approach: An interactive planning and monitoring system called IMPACT has been constructed, based on two key notions. The first is the use of a cognitively-inspired representation for events, called *transition space*[1]. This representation is designed to be both computer-manipulable and human-understandable, making possible a flexible division of labor as human and computer participants in the planning and monitoring process each focus on those subtasks best suited to their relative skills. The second notion is a technique, called "split confirmation," that oversees this division of labor by allowing human and computer participants each to vouch for the likelihood of occurrence for different *parts* of planned and observed events, based on whether intuition/judgment or scanning/testing is the most suitable source of support for such vouching.

The key insight behind the split-confirmation technique is that many subtasks within planning and monitoring involve *differential* treatment of the beginnings, middle portions and endings of various events. Typically, one part of an event is carefully checked to see if its occurrence follows from independent information, while other portions are simply assumed to occur given the carefully scrutinized portion. For example, during planning, the initial portions of events are scrutinized—the preconditions—while the remaining portions are simply assumed to occur given establishment of the preconditions. During monitoring, however, the situation is often reversed, with ends of events being explicitly verified and beginnings assumed given occurrence of the endings. Other divisions are possible, too. For example, one form of prediction involves independently verifying an initial portion of an event—more than simply the preconditions and typically involving an initial set of *changes*—then assuming occurrence of the final portion given the initial portion. With split-confirmation, the human participant can apply intuition and judgment in deciding which portions of particular events must be carefully verified in a temporal context, and which are best left to assumption. The computer then carries out the requested verifications, and the human supplies additional judgment by deciding on the suitability of assuming the remaining portions of individual events.

Abstractly, the IMPACT system operates like a "spreadsheet for events," with temporal information presented graphically at multiple levels of abstraction, and with various modifications to events and other components of temporal information resulting in an automatic propagation of implications to related components of represented temporal information. Users of IMPACT can create detailed specifications of planned and/or observed activities, add reports of event occurrences, request a testing of preconditions for planned events, posit and vouch for possible consequences of events, and invoke automated recognition of event occurrences from low-level information about circumstances and changes. IMPACT is implemented in C and uses an X/Motif interface and stylized English to interact with its users.

Difficulty: For IMPACT, one difficulty lies in translating numerous currently-used formats for specifying timestamped information into a common format—here, stylized English—for entry into the system. Another hurdle involves negotiation of terminology—object types and attributes—for describing events involved in planning and monitoring. Finally, careful consideration must be given to the role of computer inference in the system, so that it does not encroach on the realm of human intuition and judgment.

Impact: IMPACT can be operated in three successive levels of complexity. In its simplest mode of operation, IMPACT can be used as a substrate for maintaining the consistency of collected reports of timestamped factual knowledge. At a more complex level, IMPACT can be used as an environment for performing automated and semi-automated recognition of event occurrences from low-level timestamped information. At the most complex level, IMPACT can be used as a general-purpose planning and monitoring environment, with facilities for considering extraneous event occurrences, causal consequences of planned and observed events, and alternate hypothetical scenarios of events.

Future Work: In the past year, we have begun implementation of a new version of IMPACT, in which the underlying temporal information is maintained in a relational database. This implementation will make it easier to access information already maintained in relational databases and exchange the results of IMPACT's processing with other software systems.

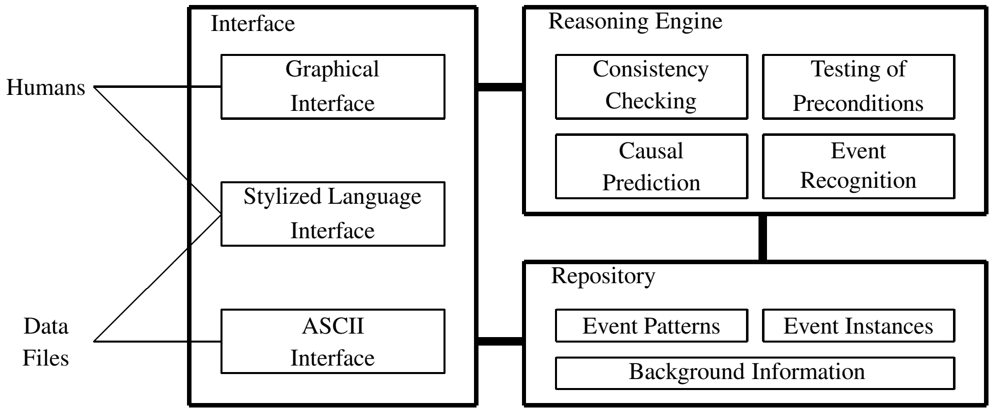

Figure 1: A block diagram of the IMPACT system, indicating interfaces to humans and data files, functional capabilities, and maintained information types.

Research Support:

"An Anchor Desk Associate for Mobile Command and Control" funded by DARPA under contract number F30602-94-C-0199 administered by Rome Laboratory.

References:

[1] G. C. Borchardt. *Thinking between the Lines: Computers and the Comprehension of Causal Descriptions.* MIT Press, 1994.

[2] T. M. Cook and R. A. Russell. *Introduction to Management Science.* Prentice Hall, 1993.

[3] B. Drabble (ed.). *Proc. Third International Conference on Artificial Intelligence Planning Systems.* AAAI Press, 1996.

[4] H. E. Shrobe (ed.). *Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence.* Morgan Kaufmann, 1988.