

A New Technique Enables Dynamic Replanning and Rescheduling of Aeromedical Evacuation

Alexander Kott, Victor Saks, and Albert Mercer

■ We describe an application of a dynamic replanning technique in a highly dynamic and complex domain: the military aeromedical evacuation of patients to medical treatment facilities. U.S. Transportation Command (USTRANSCOM) is the U.S. Department of Defense (DoD) agency responsible for evacuating patients during wartime and peace. Doctrinally, patients requiring extended treatment must be evacuated by air to a suitable medical treatment facility. The Persian Gulf War was the first significant armed conflict in which this concept was put to a serious test. The results were far from satisfactory—about 60 percent of the patients ended up at the wrong destinations. In early 1993, the DoD tasked USTRANSCOM to consolidate the command and control of medical regulation and aeromedical evacuation operations. The ensuing analysis led to TRAC2ES (TRANSCOM regulating and command and control evacuation system), a decision support system for planning and scheduling medical evacuation operations. Probably the most challenging aspect of the problem has to do with the dynamics of a domain in which requirements and constraints continuously change over time. Continuous dynamic replanning is a key capability of TRAC2ES. This article describes the application and the AI approach we took in providing this capability.

U.S. Transportation Command (USTRANSCOM) is the U.S. Department of Defense (DoD) agency responsible for evacuating patients during wartime and peace. Doctrinally, patients requiring extended treatment must be evacuated by air to a suitable medical treatment facility (MTF). The process

of identifying an MTF that constitutes a suitable destination for a given patient (based on a match between the patient's medical condition and MTF's capability and on economics and transportation availability) is called *regulating*. The process of routing and scheduling the required aeromedical evacuation flights (missions) and assigning patients to suitable missions is *evacuation planning and execution*.

Problem Description

Other than in peace time, there has been limited experience with this approach to handling patients. The Persian Gulf War was the first significant armed conflict in which this concept was put to a serious test. The results were far from satisfactory—about 60 percent of the patients ended up at the wrong destinations and half in the wrong country (Endoso 1994).

In early 1993, the DoD tasked USTRANSCOM to consolidate the command and control of medical regulation and aeromedical evacuation operations during peace, war, and projected contingencies. The ensuing analysis led to the development of a decision support system—TRAC2ES (TRANSCOM regulating and command and control evacuation system).

The integrated medical regulation-evacuation problem requires the dynamic identification of appropriate MTFs for new patients and the planning-scheduling of aeromedical evacuation operations to transport these patients from current locations to selected MTFs. This is a large-scale, highly dynamic planning and scheduling

From the computational point of view, the problem of military aeromedical regulation and evacuation (ARE) of patients to MTFs is a complex extension of the well-known dial-a-ride problem

problem that can involve hundreds or even thousands of simultaneous patient movement requests. Each patient has one or several medical requirements that constrain the type of MTF to which he/she can be evacuated and a ready time prior to which evacuation cannot start. Additional constraints can include a maximum altitude above which the evacuation aircraft cannot take the patient and a maximum number of hours that a patient can spend in a flight before requiring an overnight rest. Planning-scheduling operations in this domain require the dynamic coordination and (re)allocation of a large number of resources subject to a wide variety of constraints. Key resources and associated constraints include aircraft and their different characteristics (for example, capacity, refueling requirements), air and medical crews and restrictions on the number of hours they can work in any given day, airports and their characteristics (for example, capacity, types of aircraft they can accommodate), and hospital beds at MTFs located all around the globe and the types of patients each MTF can accommodate.

Probably the most challenging aspect in planning and scheduling medical evacuation operations has to do with the dynamics of a domain in which requirements and constraints continuously change over time. New patient requests come in; others get canceled. Patient conditions change over time, possibly requiring the delay, acceleration, or cancellation of a patient's evacuation or a change in the patient's destination. Availability of key assets is also subject to unpredictable events (for example, aircraft maintenance problems, hospital beds not getting freed on time, airfield attrition). Weather conditions can affect evacuation, requiring that a mission be delayed, rerouted, or canceled.

In building and revising evacuation plans, a number of objectives and preferences need to be taken into account. The number of patients evacuated to adequate MTFs has to be maximized, with urgent patients given priority over routine ones. The time to pick up and deliver patients to their MTFs, especially urgent ones, should be as short as possible. The time each patient spends in a flight and the number of stops during his/her evacuation also have to be minimized. Other important considerations include minimizing the number of missions and maximizing aircraft capacity use.

Features of the Aeromedical Regulation and Evacuation Problem

From the computational point of view, the problem of military aeromedical regulation

and evacuation (ARE) of patients to MTFs is a complex extension of the well-known dial-a-ride-problem (Sadeh and Kott 1996). The ARE problem exhibits a number of features commonly found in dynamic transportation problems that require dynamic replanning and rescheduling: (1) multiple demands to transport commodities or entities (in our example, patients) from or to origin or destination points (for example, airports and hospitals), (2) multiple resources (for example, planes, hospitals) that are to be routed and scheduled to meet the demands, (3) time-window constraints (a patient cannot be picked up at the origin point until he/she is prepared for departure and delivered to the airport), (4) capacity constraints (an aeromedical evacuation mission has a limited number of seats available), (5) multiple other constraints of varied nature (for example, constraints on duration of tours associated with vehicles and/or particular requests), (6) demands that can change dynamically while the schedule is executing (for example, new patients might need to be evacuated, or medical condition of a patient might change), (7) dynamically changing resources (for example, a mission can be delayed or canceled, or an airport can be closed because of the weather), and (8) dynamically determined destination and/or origin points (for example, patient's destination might be determined based on the available missions or available hospital beds).

Planning and Scheduling—Reactive and Predictive

Here we discuss some of the terminology we use in this article, particularly the meaning of *reactive*. In a planning problem, the solver is given a description of the current state of the problem world and a set of goals to achieve. The task is to find a plan—a sequence (or, more generally, a network) of activities that will lead to the desired goals. In a scheduling problem, the solver is given a plan of activities and a set of available resources. The task is to find a *schedule*, an assignment of the activities to the appropriate resources in suitable time windows.

In both problems, the solution must satisfy a given set of constraints, such as precedence constraints on activities or capacities and capabilities of resources. Usually, the solution should also attempt to minimize some cost or value measure, such as the cost of the resources used to accomplish the activities. In real-world problems, these two problems are often closely coupled; we treat these two problems as a combined planning and scheduling problem. We




Figure 1. The Context of Dynamic Replanning.

use the terms *plan* and *planning* to include also schedule and scheduling.

The *predictive* (also called static) formulation of this problem assumes that all activities are to take place in the future. In the *reactive* (dynamic or real-time) problem, some of the activities are occurring at the same time that the planning problem is being solved. The reactive problem typically occurs when the current plan has been disrupted either by unexpected events in the world or changes in goals (figure 1). This article is concerned with the more difficult reactive problem.

At first glance, the reactive problem can be reduced to a predictive problem by determining what the current state of the world is and then formulating a new predictive problem, where all activities are to take place in the future. The difficulty with this approach is that a new “clean-sheet” plan is likely to introduce major destabilizing disruptions in the plan-execution and plan-control process.

The key concern in reactive replanning is plan continuity—the new plan should not introduce unnecessary disruptions.

Continuity—The Key Issue of Reactive Replanning and Rescheduling

It is useful to note that the ARE problem can be seen as a constrained optimization problem. The example problem formulation can

briefly be outlined as follows:

Given: Current data of patients, missions, MTFs, airfields, airport staging facility ([ASF] for patients in transit), and so on

Find: Mission schedules and patient itineraries

Subject to constraint: Mission capacity, MTF capacity, and so on

Optimize or prefer: Minimum time to delivery, minimum use of resources, minimum deviation from the already existing plan-schedule

The last preference—minimum deviation from the existing plans—deserves special attention and presents a particularly difficult challenge. When searching for a solution to a dynamic rescheduling problem, one must attempt to minimize the extent of disruptions that the new plan introduces into current execution activities. Dynamic changes to a current plan violate some commitments and waste some investments made in preparing or executing the activities scheduled in the current plan.

For example, significant efforts are expended to prepare for a particular mission: Flight and medical crew are assembled and briefed, paper work is issued, flight controllers insert this mission into their plans, maintenance and




Figure 2. The Degree of Commitment to Plan Decisions, and the Associated Cost of Changing These Decisions, Is a Function of Time.

refueling resources are allocated, and so on. By changing a mission, we negate some of these investments and force the expenditure of additional resources to undo the effects of some of the activities.

Increased risk is another harmful effect. A plan change increases the risk of mishaps, miscommunications, erroneous data entries, and erroneous decision making. Thus, a mandatory feature of a dynamic replanning process is the ability to minimize the number of changes—violations of commitments made by the earlier plan.

Sliding Scale of Commitment

It is important to note that the less time that is left between the current moment and the time when the activity is planned to happen, the costlier the change is. For example, it is much more expensive to make a change to a mission that takes off in 20 minutes than a mission that is scheduled to take off in 20 hours.

Other factors in addition to time increase the cost of disruption. One such factor is the extent to which the decision has been communicated to the world; for example, if the MTF has already begun preparations to receive the patient, then the cost of disruption is greater.

These observations led to sliding scale of commitment (SSC), a concept proposed by Victor Saks and William Elm at Carnegie Group, Inc., circa 1993. The SSC concept states that there is a cost (the commitment cost) to be paid for rescheduling, even for actions yet to be executed; the cost is a function of time,

the closer—the costlier (figure 2); and the cost can also be dependent on domain-specific details of commitment, for example, who has been notified and what preparations are necessary for the actions.

In summary, we argue that the issue of continuity-stability is key to effective dynamic replanning. In fact, it is the dimension that distinguishes reactive replanning from repetitive application of predictive planning. The sliding scale of commitment adds yet another dimension of complexity to the issue of plan continuity. Prior research has not addressed this issue in a systematic and explicit manner.

Application Description

According to Endoso (1994), during the Persian Gulf War, the medical regulating and evacuation process produced less than satisfactory results—about 60 percent of the patients ended up at the wrong destination. In early 1993, the DoD tasked USTRANSCOM to develop a global command and control (c2) system to remedy deficiencies in medical regulating and evacuation. USTRANSCOM undertook a business reengineering study and identified two important concepts:

First, separate regulating and evacuation activities were combined into a single, “one-stop-shopping” process for patient movement. Merging separate responsibilities of the medical regulator and evacuation coordinator creates a single *evacuation broker* responsible for both regulating and evacuation. The second major concept was that of the “lift bed,” which relates medical capability to transportation capability. The study pointed to the need for automated support to implement these two major concepts and provide a global c2 system.

USTRANSCOM decided to explore possible automated information system solutions. This exploration led to the development of TRAC2ES, a large-scale system that involved in excess of 140 person-years of development and deployment effort.

TRAC2ES Functions

TRAC2ES is a command and control application that provides users with enhanced computer-aided capabilities to forecast, plan, coordinate, and execute the global process of regulating and evacuation of military patients, specifically: (1) provide entry and communication mechanisms that enable personnel of MTFs (for example, field hospitals) to request evacuation of a patient and define the medical requirements for the evacuation process; (2) perform automated or semiautomated plan-

ning of such new requests: matching of the patient with the appropriate hospital, aeromedical evacuation flights, intermediate rest facilities, and so on, possibly requiring the dynamic revision of prior plans; (3) request additional resources, for example, additional evacuation flights along suitable routes or hospital beds in bottleneck locations when current capacity is insufficient to meet the demands; (4) coordinate evacuation plans between different regional planning centers (theater patient movement requirements centers [TPMRCs]) and assure that there are no conflicts between their intended uses of shared resources; (5) store all plans in the databases that are globally accessible to all authorized parties (medical and transportation personnel) and issue notifications to hospitals (when and whom to prepare for evacuation or for arrival and treatment), transportation organizations (when, where and whom to pick up and deliver), and so on; (6) collect and provide near-real-time data regarding the current in-transit status (location, condition, future plans) of all patients and critical assets (medical transportation, equipment, crews); and (7) continuously monitor the execution of the plans, alert about disruptions of the plan, and provide dynamic replanning to accommodate the disruptions.

TRAC2ES Users

TRAC2ES users include (figure 3) (1) personnel at the origination hospitals (for example, military field hospitals), who use TRAC2ES to enter requests for evacuation of a patient and to receive notification of when and how this request will be satisfied; (2) personnel at the destination hospitals (for example, large hospitals in the United States), who use TRAC2ES to notify of their available bed capacity and receive notifications about arrival of new patients; (3) staff members of the TPMRCs or joint force patient movement requirements centers (JPMRCs), who use TRAC2ES to monitor all new requests or disruptions in the current plan (figure 4), generate a modified plan, and issue orders to the performing organizations; (4) staff members of the global patient movement requirements center ([GPMRC], part of the USTRANSCOM staff), who monitor the overall global process of evacuation, resolve conflicts over the use of shared resources, and procure and allocate additional resources; (5) personnel of the transportation organizations (both dedicated aeromedical transporters and military cargo transporters and civilian airlines who are also frequently pressed into aeromedical evacuation service), who use TRAC2ES to receive orders to transport a patient and report




Figure 3. The Evacuation Planning and Execution Process Involves Originating Medical Treatment Facility (MTF), Destination MTF, One or Several Theater Patient Movement Requirements Centers, Aeromedical Transport Organizations, and Global Patient Movement Requirements Center.

the actual success of executing the order and any disruptions; and (6) all authorized personnel of the U.S. government, who can use TRAC2ES to find the location and status of any TRAC2ES-monitored patient who is in transit.

TRAC2ES Operation

Let us consider an example. Suppose an American soldier is injured in Europe and receives initial care at an MTF in Germany. Data about the injured soldier are entered at the MTF and sent to the European TPMRC workstation. A few days later, the soldier is ready to be moved to an MTF near the soldier's home in the United States. Information about all flights of the military aircraft is located on a workstation at the GPMRC in the United States. Available beds and medical care capabilities at hospitals in the United States are stored on a TPMRC workstation located in Illinois. The TRAC2ES component at the European TPMRC will analyze mission data and bed data and assist planners in providing the soldier with a safe, speedy, and cost-effective trip home to a hospital having the appropriate personnel, equipment, and facilities. The journey can involve multiple connecting flights and can require use of overnight rest facilities, available at some airports. This planning process can, if necessary, involve replanning of movements of some other patients to open the bottlenecks in transportation or medical assets. TRAC2ES will monitor the movement of this soldier from the beginning of the soldier's journey in Europe to




Figure 4. Advanced Graphic User Interface for Monitoring the Dynamic Changes in the Evacuation Process Were Designed Using Cognitive Systems Engineering (Potter, Ball, and Elm 1996).

the destination hospital in the United States. The system will react along the way, if necessary, to problems such as airport closings, missed flights, and changes in hospital bed availability.

Even an example of a single patient involves a degree of complexity. TRAC2ES, however, is designed to handle many thousands of patients simultaneously, which leads to a combinatorial puzzle of missions, hospitals, airports, and other resources of astronomical size and complexity.

TRAC2ES Components

TRAC2ES includes MTF components, TPMRC components, fly-away components, and a GPMRC component.

MTF component: Multiple MTFs access TRAC2ES on the World Wide Web using commercially available web browsers. The TRAC2ES server provides a site that the MTF personnel can access to enter a request for evacuating a patient, the condition of the patient, and any preferences for the evacuation. The same site

allows them to find out about the status of the request, that is, if the patient has been scheduled for evacuation. MTFs also receive e-mail notification of the evacuation schedule, automatically generated using the TRAC2ES server. Hardware: PCs running WINDOWS. Technologies used: web server, web browser, HTML, JAVA SCRIPT.

TPMRC component: Each theater (military region) of operations includes an organization responsible for movement of patients to, from, and within the theater. This organization—TPMRC—uses a TPMRC-specific component of TRAC2ES, which receives requests for patient movements, plans and replans them, and issues orders to appropriate transportation and medical organizations to execute the planned movements. In particular, the dynamic replanning capability resides within the TPMRC component. When a TPMRC component is used in the context of a joint force operation, we call it a JPMRC component. Hardware: Sun workstations under SOLARIS. Technologies: Versant distributed object-oriented database management system, C++, MOTIF.

Fly-away component: The fly-away component is a version of the TPMRC component in a self-contained configuration that is used to rapidly set up a movement planning center in any new region.

GPMRC component: The GPMRC at USTRANSCOM utilizes this component (a modification of the TPMRC component) to monitor the movements of patients worldwide and perform what-if planning.

Use of AI Technology— Continuity-Guided Regeneration

A very large, distributed application with multiple types of user and function, TRAC2ES encompasses a number of novel technological aspects. In this article, we focus on one of these technological aspects, but it is worthwhile to mention others at least briefly and to point the reader to the corresponding references: (1) broad use of worldwide distributed object-oriented databases and rigorous object-oriented design and development methodology, (2) novel user interfaces developed using the methodology of cognitive system engineering (Potter, Ball, and Elm 1996), and (3) a multi-agent problem-solving process (Saks et al. 1997).

In the following discussion, we focus on one key function of TRAC2ES that is particularly relevant to the AI perspective—the planning and scheduling function—and the corresponding AI technology—continuity-guided regeneration (Kott and Saks 1996).

Continuity-Guided Regeneration in Planning and Scheduling

We saw the issue of continuity as the key challenge in the dynamic replanning and scheduling function of TRAC2ES. In response to this challenge, we developed the technique of continuity-guided regeneration (CGR). CGR is an extension of the ideas originally developed at CGI for solving redesign problems, circa 1990, within the SPEX design and configuration shell (Berry and Kott 1992). The key idea is to regenerate the plan while we use the currently executing plan as a constraint on the solution. One significance of this idea is that it assures plan continuity without having to answer the difficult question of how much of the current plan to undo. We extended the CGR idea of SPEX to make it capable of accounting for the sliding scale of commitment.

In TRAC2ES, we applied it to modify a particular scheduling approach, microopportunistic

search, an instantiation of constrained heuristic search (Sadeh 1991). The main steps of the microopportunistic search procedure can be summarized as follows: (1) a set of candidate plans is created for each requirement; (2) reliance (dependencies) of the requirement on each of the available resources is computed; (3) each requirement distributes its demand to the resources and time line; (4) overall contention for each resource is computed by time bucket; (5) the most contended-for resource is selected; (6) the requirement that is most reliant (dependent) on the selected resource is assigned to the resource; and (7) go to 2.

We modified this approach using the idea of commitment constraint. We elected to use the reliance computation as the area of the search procedure where the commitment constraints enter the picture. For each requirement (in our specific case, a patient movement request), (1) retrieve the resource that has been committed to this requirement in the existing plan, (2) compute a measure of importance of preserving the same commitment (IPC) based on how far in the future the use of this resource is to take place and what the other domain-dependent commitment measures are, (3) increase the reliance measure associated with this requirement-resource pair using the IPC, and (4) use the assignment mechanism to assign a resource to the requirement. The higher reliance value leads to a better chance that the requirement will be assigned to the same resource as in the existing plan if such an assignment is feasible.

We believe that the CGR paradigm can be used not only as an add-on to the microopportunistic search but also with other planning and scheduling approaches. A key advantage of the CGR approach is that it eliminates the key dilemma of the rescheduling problem: how much of the prior schedule to undo.

Applying the CGR Approach to the TRAC2ES Dynamic Replanning Problem

The solution process (figure 5) begins with the arrival of disruptive events, such as the closure of an airport. Each event is decomposed into one or more of so-called world deltas. Each *world delta* is a description of a difference between the expected state of an entity (for example, airport is open) and its actual state (for example, airport is closed).

The deltas are sorted, using domain-specific heuristics, in the order of significance and downstream impact. Depending on the nature of each delta, the algorithm then can create new deltas, update attributes and associations of existing resources and demands, or reconfig-




Figure 5. Overall Flow of the Dynamic Replanning Process Using the Continuity-Guided Regeneration Approach.

ure existing resources (missions). For example, the following operations can be performed on missions: (1) update modified missions with a complete route-schedule given, (2) repair missions that have been diverted for some reason and for which continuation of the mission is not given, (3) reroute missions, if necessary, to provide a suitable route for new urgent patients.

All patient itineraries that have been in any way affected by any of the deltas are marked as invalid. This undoing of itineraries in the original plan is done liberally; that is, we prefer to err on side of undoing more itineraries than necessary given that the CGR provides continuity.

At this point, the process of assigning patients to resources (for example, missions and hospitals) is performed. It involves reassigning patients to resources (missions, hospitals, and so on) using the CGR approach. An assignment of a patient to the same resources as in the original plan is given a higher score. The increase in score is computed in inverse proportion to how far in the future the use of

the resources begins (to account for the SSC).

Finally, the system generates a recommendation that describes the impact of the world deltas and the set of suggested modifications to the plan.

Experience with TRAC2ES

The most significant use of TRAC2ES to date has been at military exercises—the most demanding tests of military capabilities short of real war. It should be pointed out that benefits of TRAC2ES capabilities are more difficult to demonstrate in peace time—the flow rate of patients is fairly low; resources are not overconstrained; and a handful of human planners can do a good job planning, executing, and monitoring the evacuation process in an essentially manual process. However, the situation is entirely different at the time of war or its closest possible approximation—military exercise. The flow of patients is large, resources are overconstrained, disruptions are unrelenting, and human planners are unable to cope with the volume of information and the complexity of the problem. That's where TRAC2ES

delivers a capability that simply cannot be delivered by any other means.

Consider one of the exercises—Unified Endeavor 98-1 (UE 98-1)—where TRAC2ES has actively been used and where it demonstrated its benefits. UE 98-1 was conducted in October–November 1997. It was a computer-assisted exercise that involved components from the United States Army, Marines, Navy, and Air Force and the military of the United Kingdom. In the scenario, the operation responded to an aggressive northern nation in the Middle East that threatened a smaller southern nation's sovereignty. U.S. Central Command (USCENTCOM) tasked U.S. Atlantic Command to provide a combined joint task force to deter the aggressor while USCENTCOM deployed additional forces to the region. United Kingdom forces participated in UE 98-1 through a collateral exercise, UK PURPLE LINK 97; their joint rapid deployment forces consisted of a reinforced brigade with maritime, air, special forces, and logistics elements in coalition with U.S. forces.

The exercise was supported by the TRAC2ES FLYAWAY 1.3.1 that provided a single-theater patient movement capability consistent with the UE 98-1 concept of operation. Wounded-in-action casualties were generated from several sources, including the Master Scenario Event List and a battle-simulation model. The exercise controllers periodically injected additional casualty situations into exercise play. There were several MTF nodes distributed over the United States. Each MTF was submitting patient movement requests to the JPMRC component of TRAC2ES. Users and observers stated that the exercises demonstrated a number of value-added capabilities provided by TRAC2ES (compared to conventional near-manual operation). These capabilities included a much more robust method of communicating patient movement requests from MTFs to the JPMRC, the ability of MTF to see if a patient has been scheduled for movement and how, and the ability of MTF personnel to view the entire planned itinerary of the patient. In addition, users at a variety of locations (including United Kingdom components) were able to obtain information about patient's planned movements. Finally, the dynamic replanning capability deserves particular mention. Numerous extremely disruptive events were injected into the exercise: closure or movement of various facilities, cancellation or rescheduling of flights; and changes in requested patient movements. It was common for nearly all patients previously scheduled to be disrupted because of a single-event injec-

tion; many occurred while patients were already en route. It was clear to the observers that without TRAC2ES, JPMRC officers would be unable to deal with the resulting chaos and remedy disrupted movement schedules. The dynamic replanning capability of TRAC2ES helped users to quickly and easily create solutions to problems caused by disruptive events. Helping the JPMRC staff to solve the hardest, most time-consuming parts of the patient movement scheduling process, TRAC2ES freed them up for other tasks.

Development and Deployment Process

TRAC2ES automated information system development started with a proof-of-concept (POC) prototype, followed by two operational prototype releases that tested key functions and deployment to TRAC2ES owner issues. The POC, successfully completed in early 1994, featured tests of preliminary algorithms that used constraint-based reasoning for development of lift beds and rudimentary airlift schedules.

Operational prototype 1, delivered in October 1994, began prototyping TRAC2ES hardware with network communications. Operational prototype 1 software, although introducing some limited features, was primarily a demonstration capability used to illustrate future directions and application capabilities that TRAC2ES might contain. Operational prototype 2, delivered in June 1995, was based on a newly designed prototype architecture that includes a distributed object-oriented database. The operational prototype 2 graphic user interface (GUI) and algorithm provided enhanced functions, including coordinated planning by one global and several theater-based centers. In the fall of 1995, operational prototype 3 was released, featuring a number of enhancements to the GUI and the planning algorithm. In 1996 and 1997, the work focused on extending the operational prototype 3 system into a self-contained, readily deployable TPMRC system that could be used for rapid installation in exercises and real-life contingencies. This resulting system was named TRAC2ES FLYAWAY. It is this version of TRAC2ES that was used in the exercise UE 98-1 we discussed earlier.

From its inception, TRAC2ES has involved users in shaping its "to be" vision using business-process reengineering and the incorporation of multidisciplinary user feedback. USTRANSCOM sponsored numerous corporate information management workshops, each of which focused on reengineering a portion of

From its inception, TRAC2ES has involved users in shaping its "to be" vision using business-process reengineering and the incorporation of multi-disciplinary user feedback.

the patient regulation and evacuation process into a seamless whole. One of the key challenges in the TRAC2ES development and deployment process was the fact that both the system and the business process were evolving in parallel and heavily influenced each other. The users continued to work on modifying and optimizing the business process as they learned more about the feasibility and capabilities of TRAC2ES; at the same time, TRAC2ES was redesigned and modified to meet the evolution of the business process.

From the software methodology perspective, TRAC2ES development has two particularly salient features: (1) a prominent role played by the object-oriented computer-aided software engineering (CASE) tool with automated code-generation capability and (2) the extensive use of the object-oriented DBMS.

Maintenance

Developers of applications involving AI techniques and domain-specific knowledge bases often envision that the maintenance of the knowledge base will eventually be undertaken by the knowledge engineers who are part of the end user organization or possibly by the end users themselves. Other common visions of application maintenance involve self-learning of the system.

We did not feel that these concepts were practical in the context of TRAC2ES. Although several key techniques used in TRAC2ES belong to the field of AI, TRAC2ES does not have a large and prominent knowledge base. TRAC2ES relies in its problem-solving approach on a relatively small collection of constraints and preferences, which are fairly stable and do not change frequently. The life-critical nature of the application also argues against any maintenance process that is not extremely rigorously controlled.

The maintenance process of TRAC2ES is structured around a rigorous tracking, analysis, and prioritization of request for bug fixes and enhancements; a continuous development process involving rigorous and formal cycles of requirements maintenance; design modifications using an object-oriented CASE tool; extensive reviews; an implementation that involves a large percentage of automatically generated code; verification and validation; quality assurance; and extensive testing.

Conclusion

TRAC2ES is a large system that serves a broad range of users worldwide and involved several hundred person-years of development effort.

Although providing a range of functions and benefits and relying on a number of advanced technologies, TRAC2ES is critically dependent on a particular capability—continuous dynamic replanning. One might argue that the entire doctrine of large-scale aeromedical evacuation (and the associated business process) could be considered infeasible without such a capability.

We found that plan continuity is a key concern in devising a dynamic replanning technique. It is the key differentiator between predictive and dynamic planning. Our prior work in design and redesign led to the novel CGR technique. The key idea is to regenerate the plan while the currently executing plan is used as a constraint on the solution. One significance of this idea is that it assures the plan continuity without having to answer the difficult question of how much of the current plan to undo.

Our implementation of this approach in application to aeromedical evacuation demonstrated that this approach addresses the continuity issue, enables the SSC, eliminates the “extent of undoing” question, and provides an effective mechanism to control the balance between the demands of continuity and optimality.

TRAC2ES has demonstrated capabilities of dynamic replanning, not feasible using current near-manual methods, in the demanding environment of modern military exercises.

Acknowledgments

This work was supported by USTRANSCOM. Leadership was provided by William Elm and Albert Mercer (Carnegie Group, Inc.), Richard Simpson (MITRE), Jack Simpson, and Correy Kirschner (USTRANSCOM).

References

- Berry, F., and Kott, A. 1992. The SPEX Shell: An Approach to Automating Application and Sales Engineering Processes. In Proceedings of the 1992 ASME Computers in Engineering Conference, 136–143. New York: American Society of Mechanical Engineers.
- Endoso, J. 1994. TRANSCOM Wants Casualties to Get Where They Belong. *Government Computer News*, May 2.
- Kott, A., and Saks, V. 1996. Continuity-Guided Regeneration: An Approach to Reactive Replanning and Rescheduling. Paper presented at the Florida Artificial Intelligence Research Symposium, 20–22 May, Key West, Florida.
- Potter, S.; Ball, R.; and Elm, W. 1996. Supporting Aeromedical Evacuation Planning through Information Visualization. Paper presented at the Third Annual Symposium on Human Interactions with Complex Systems, 25–28 August, Dayton, Ohio.
- Sadeh, N. 1991. Look-Ahead Techniques for Micro-

Opportunistic Job-Shop Scheduling. Ph.D. thesis, School of Computer Science, Carnegie Mellon University.

Sadeh, N. M., and Kott, A. 1996. Models and Techniques of Dynamic Demand-Responsive Transportation Planning. Technical Report, CMU-RI-TR-96-09, The Robotics Institute, Carnegie Mellon University.

Saks, V.; Braidic, G.; Kott, A.; and Kirschner, C. 1997. Distributed Medical Evacuation Planning: Which Problem Should Each Agent Solve. Paper presented at the Fourteenth National Conference on Artificial Intelligence, Workshop on Constraints and Agents, 27 July, Providence, Rhode Island.



Alexander Kott is the director of Research and Development at the Advanced Decision Support Unit of Logica Carnegie Group. During the TRAC2ES project, he led the development of the GPMRC component and the CGR dynamic replanning algorithm. He received his Ph.D. at the University of Pitts-

burgh in 1989. He worked on research and advanced application projects in search algorithms, intelligent design automation, requirements engineering, distributed problem solving, and planning and scheduling in dynamic and uncertain environments. His e-mail address is akott@cgi.com.



Victor Saks served as principal engineer at Carnegie Group, Inc., while this research was conducted and was the key designer of TRAC2ES algorithms. He is currently employed by Transarc Corporation. He received a Ph.D. in mathematics from Wesleyan University in 1972 and an M.S. in computer

science from the State University of New York at Buffalo in 1985. His research interests include dynamic replanning and rescheduling and distributed computing. His e-mail address is tsvi@transarc.com.

Albert Mercer is the director of operations at the Advanced Decision Support Unit of Logica Carnegie Group. He experienced the real-world challenges of military medical evacuation while he served as a military plans and operations officer during his 22-year career in the U.S. Department of Defense. He received his Master's degree from Webster University in Health Systems Management in 1981. His interests include advanced applications for decision makers in health care and military software-engineering methodologies and software project-management processes. From 1995 to 1997, he was the manager of the TRAC2ES project. His e-mail address is mercer@cgi.com.

Call for Papers

The Thirteenth International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems (IEA/AIE-2000)

New Orleans, Louisiana, USA • June 19-22, 2000

*Sponsored by the International Society of Applied Intelligence
Organized in Cooperation with AAAI, ACM/SIGART, CSCSI, IEE, INNS, JSAI, SWT, and USL*

IEA/AIE-2000 will focus on methodological as well as practical aspects in the development of KBSs, knowledge modeling, reasoning, evolutionary computing, distributed and multi-agent systems, hybrid techniques that integrate the symbolic and connectionist perspectives in the industrial application of AI, and application of intelligent systems' technology to solve real life problems. Practical experiences with state-of-the-art AI methodologies are also acceptable when they reflect lessons of unique value to the conference attendees. Accepted papers will be published at full length in the proceedings, and authors of selected papers will be invited to write an in-depth paper for publication in the *International Journal of Applied Intelligence*.

Authors are invited to submit no later than November 17, 1999: 1) a single hard copy of the paper to Dr. Raja Loganathanaraj, IEA/AIE-2000 Program Chair, Center for Advanced Computer Studies, #2 Rex Street CC 451, University of Southwestern Louisiana, Lafayette, LA 70504 USA and 2) e-mail a list of key words and a URL address containing a printable version of the paper to ieaaie2000@cacs.usl.edu. Submissions must be written in English, having up to 10 single-spaced pages presenting the results of original research or innovative practical applications relevant to the scope of the conference. Shorter works, up to 6 pages, to be presented in 10 minutes may be submitted as short papers representing work in progress or suggesting possible directions. Please indicate "short paper" in the submission letter in this case.

Also, the conference will contain a series of special track sessions. If an author desires to include his/her paper in a special track, please indicate that as well. Researchers are invited to submit a workshop, a tutorial, or a special track session proposal preferably by email before July 20, 1999 according to the guideline on the conference web page at <http://www.cacs.usl.edu/~ieaaie2000/>.

General conference information can be sought from the General Chair at the following address:

Dr. Moonis Ali, General Chair of IEA/AIE-2000
Southwest Texas State University, Department of Computer Science
601 University Drive
San Marcos TX 78666-4616 USA
Email: ma04@swt.edu



Android Epistemology

Edited by **Kenneth M. Ford, Clark Glymour, & Patrick J. Hayes**

Epistemology has traditionally been the study of human knowledge and rational change of human belief. *Android epistemology* is the exploration of the space of possible machines and their capacities for knowledge, beliefs, attitudes, desires, and action in accord with their mental states. From the perspective of android epistemology, artificial intelligence and computational cognitive psychology form a unified endeavor: artificial intelligence explores any possible way of engineering machines with intelligent features, while cognitive psychology focuses on reverse engineering the most intelligent system we know, us. The editors argue that contemporary android epistemology is the fruition of a long tradition in philosophical theories of knowledge and mind.

The sixteen essays by computer scientists and philosophers collected in this volume include substantial contributions to android epistemology as well as examinations, defenses, elaborations, and challenges to the very idea.

Contributors include Kalyan Basu, Margaret Boden, Selmer Bringsjord, Ronald Chrisley, Paul Churchland, Cary deBessonet, Ken Ford, James Gips, Clark Glymour, Antoni Gomila, Pat Hayes, Umar Khan, Henry Kyburg, Marvin Minsky, Anatol Rapoport, Herbert Simon, Christian Stry, and Lynn Stein.

ISBN 0-262-06184-8 336 pp., index. \$25.00 hardcover

The AAAI Press • Distributed by The MIT Press

Massachusetts Institute of Technology, 5 Cambridge Center, Cambridge, Massachusetts 02142

To order, call toll free: (800) 356-0343 or (617) 625-8569. MasterCard and VISA accepted.

Illustration: "Electronic Man," an aquatint etching by Max Papart; reproduced courtesy of Nahan Galleries, New York, New York