A New Technique Enables Dynamic Replanning and Rescheduling of Aeromedical Evacuation

Alexander Kott, Victor Saks, Albert Mercer

Carnegie Group, Inc.
5 PPG Place
Pittsburgh, PA 15222
akott@cgi.com, saks@cgi.com, mercer@cgi.com

Abstract

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 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 Persian Gulf war was the first significant armed conflict in which this concept has been put to a serious test. The results were far from satisfactory -- about 60% of the patients ended up at the wrong destinations. In early 1993, the Department of Defense 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 paper describes the application and the AI approach we took in providing this capability.

Problem Description

U.S. Transportation Command (USTRANSCOM) is the 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 matching 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.

There has been very limited experience with this approach to handling patients other than in peace time. The Persian Gulf war was the first significant armed conflict in which this concept has been put to a serious test. The results were far from satisfactory -- about 60% of the patients ended up at the wrong destinations and half in the wrong country (Endoso1994).

In early 1993, the Department of Defense 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 Medical Treatment Facilities (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 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 or 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, a maximum number of hours that a patient can spend in a flight before requiring an overnight rest, etc. 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 assets/resources and associated constraints include aircrafts and their different characteristics (e.g. capacity, refueling requirements, etc.), air and medical crews and restrictions on the number of hours they can work in any given day, airports and their characteristics (e.g. capacity, types of aircraft they can accommodate, etc.), hospital beds at MTFs located all around the globe and the types of patients each MTF can accommodate, etc.

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 (e.g., aircraft maintenance problems, hospital beds not getting freed on time, airfield attrition, etc.). 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 pickup 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, maximizing aircraft capacity utilization, etc.

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 medical treatment facilities is a complex extension of the well-known Dial-A-Ride-Problem (DARP) (Sadeh and Kott 1996). The ARE problem exhibits a number of features commonly found in dynamic transportation problems that require dynamic replanning and rescheduling:

- Multiple demands to transport commodities or entities (in our example -- patients) from/to origin or destination points (e.g., airports and hospitals);
- Multiple resources (e.g., planes, hospitals) are to be routed and scheduled to meet the demands;
- Time window constraints, e.g., in our example a patient cannot be picked up at the origin point until he/she is prepared for departure and delivered to the airport;
- Capacity constraints, e.g., an aeromedical evacuation mission has a limited number of seats available;
- Multiple other constraints of varied nature, such as constraints on duration of tours associated with vehicles and/or with particular requests;
- Demands can change dynamically while the schedule is executing, e.g., new patients may need to be evacuated, or medical condition of a patient may change;
- Resources may also change dynamically, e.g., a mission can be delayed or canceled, or an airport may be closed due to the weather.
- Destination and/or origin points may have to be determined dynamically, e.g., patient's destination may be determined based on the available missions or available hospital beds.

Planning and Scheduling - Reactive and Predictive

Now a few words about some of the terminology we use in this paper, 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 employed 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 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 normally occurs when the current plan has been disrupted either by unexpected events in the world or by changes in goals (Fig. 1). This paper 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. Although this is indeed possible, the difficulty with this approach is that a new "clean sheet" plan is likely to introduce major destabilizing disruptions into the plan execution and control process.

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

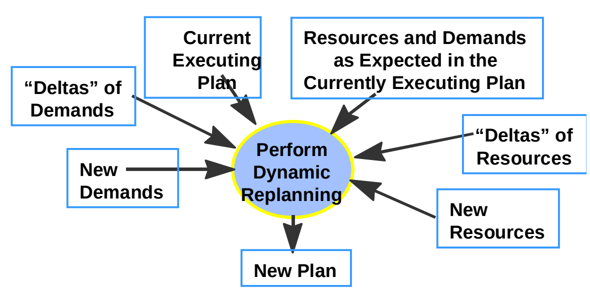

Figure 1. The context of dynamic replanning.

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 be briefly outlined as follows:

- Given: current data of patients, missions, MTFs, airfields, ASFs (Airport Staging Facility for patients in transit) ...
- Find: mission schedules and patient itineraries
- Subject to constraint: mission capacity, MTF capacity...
- 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 a 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 into 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 refueling resources are allocated, etc. By changing a mission, we negate some of these investments and force 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 is left between the current moment and the time when the activity is planned to happen, the costlier is the change. For example, it is much more expensive to make a change to a mission that takes off in 20 minutes than to 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. E.g., if the MTF has already begun preparations to receive the patient, then the cost goes up even more.

These observations led to Sliding Scale of Commitment (SSC) - a concept proposed by Dr. V. Saks and W. Elm at Carnegie Group, Inc., circa 1993. The observation is that the key to dynamic rescheduling is to minimize violations of commitment made in the current executing schedule.

Specific components of the SSC concept include:

- 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 (Fig. 2);
- the cost may also be dependent on domain-specific details of commitment, e.g., who has been notified, what preparations are necessary, etc.

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.

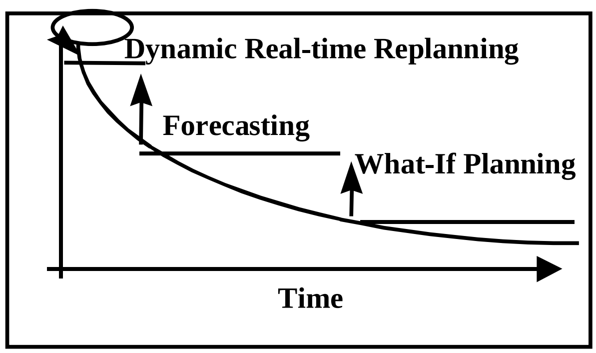

Figure 2. The degree of commitment to plan decisions, and the associated cost of changing these decisions, is a function of time.

Application Description

TRAC2ES Roots and Motivation

According to (Endoso 1994), during the Persian Gulf war medical regulating and evacuation process produced less than satisfactory results - about 60% of the patients ended up at the wrong destination. In early 1993, the Department of Defense (DOD), via DOD Directive 5154.6 and DOD Instruction 6000.11, tasked USTRANSCOM to develop a global command and control (C2) system to remedy deficiencies in medical regulating and evacuation. USTRANSCOM undertook a business re-engineering 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 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 (C2) application that provides USTRANSCOM with enhanced computer-aided capabilities to forecast, plan, coordinate, and execute the global process of regulating and evacuation of military patients. Although a very large and multi-functional system in itself, TRAC2ES is seen as a part of even broader framework - Global Transportation Network (GTN) which combines a number of databases and Decision Support Systems such as TRAC2ES and encompasses all aspects of U.S. military transportation. Within the GTN, TRAC2ES focuses on the medical transportation aspects, and shares will GTN many functions that also apply to non-medical transportation:

- provide entry and communication mechanisms that enable personnel of MTFs (e.g., field hospitals) request evacuation of a patient and define the medical requirements for the evacuation process;
- perform automated or semi-automated planning of such new requests: match the patient with appropriate hospital, aeromedical evacuation flights, intermediate rest facilities, etc.; this may require dynamic revision of prior plans;
- request additional resources, e.g., additional evacuation flights along suitable routes or hospital beds in bottleneck locations, when current capacity is insufficient to meet the demands;
- coordinate evacuation plans between different regional planning centers (TPMRC's) and assure that there is no conflicts between their intended uses of shared resources;
- 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), to transportation organizations (when, where and whom to pickup and deliver), etc.;
- 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);
- 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 (Fig. 3):

- personnel at the origination hospitals (e.g., military field hospitals): they use TRAC2ES to enter requests for evacuation of a patient and to receive notification of when and how this request will be satisfied;
- personnel at the destination hospitals (e.g., large hospitals in U.S.): they use TRAC2ES to notify of their available bed capacity, to receive notifications about arrival of new patients;
- staff of the regional patient movement centers (TPMRCs or JPMRCs): they use TRAC2ES to monitor all new requests or disruptions in the current plan (Fig. 4), to generate a modified plan and to issue orders to the performing organizations;
- staff of the global patient movement center (GPMRC, part of the USTRANSCOM staff): they monitor the overall global process of evacuation, resolve conflicts over the use of shared resources, procure and allocate additional resources;
- personnel of the transportation organizations (both dedicated aeromedical transporters, as well as military cargo transporters and civilian airlines who are also frequently pressed into aeromedical evacuation service): they use TRAC2ES to receive orders to transport a patient, and to report actual success of executing the order and any disruptions;
- all authorized personnel of the U.S. Government - they can use TRAC2ES to find the location and status of any TRAC2ES-monitored patient while in-transit.

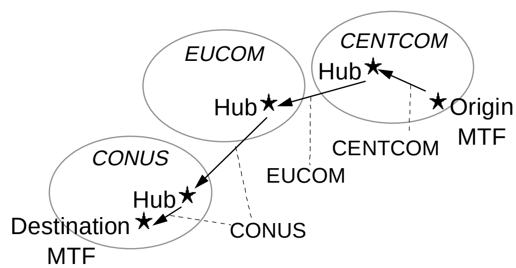

Figure 3. The evacuation planning and execution process involves originating MTF, destination MTF, one or several theater patient movement centers, aeromedical transport organizations, and global patient movement center.

TRAC2ES Operation

Let us consider an example (Braidic 1996). Suppose an American soldier is injured in Europe and receives initial care at a Medical Treatment Facility (MTF) in Germany. Data about the injured soldier is entered at the MTF and sent to the European Theater Patient Movement Requirements Center (TPMRC) workstation. A few days later, the soldier is ready to be moved to an MTF near the soldier's home in the U.S. Information about all flights of the military aircraft is located on a workstation at the Global Patient Movement Requirements Center (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. TRAC2ES component at 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 may involve multiple connecting flights, and may require use of overnight rest facilities, available at some airports. This planning process may, 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 the destination hospital in the U.S. 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 simultaneously many thousands of patients. This leads to a combinatorial puzzle of missions, hospitals, airports and other resources of astronomical size and complexity.

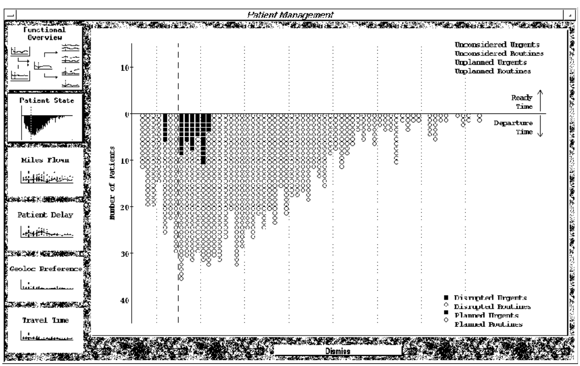

Figure 4: Advanced GUI for monitoring the dynamic changes in the evacuation process were designed using Cognitive Systems Engineering (Potter et al. 1996).

TRAC2ES Components

MTF component: multiple MTFs access TRAC2ES via the World Wide Web (WWW) using commercially available web browsers. TRAC2ES server provides a site which 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, i.e., if the patient has been scheduled for evacuation. MTFs also receive e-mail notification of the evacuation schedule automatically generated via the TRAC2ES server. Hardware: PC's running Windows. Technologies used: web server, web browser, HTML, JavaScript.

TMPRC 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/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. Hardware: Sun workstations under Solaris. Technologies: Versant distributed object-oriented DBMS, C++, Motif.

Fly Away component: a version of the TPMRC component in a self-contained configuration, used to set up a movement planning center rapidly in any new region.

GPMRC component: global patient movement center at USTRANSCOM uses this component (a modification of the TPMRC component) to monitor the movements of patients world-wide, and to perform what-if planning.

Use of AI Technology - Continuity-Guided Regeneration

A very large, distributed application with multiple types of users and functions, TRAC2ES encompasses a number of novel technological aspects. In this paper 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:

- broad use of world-wide distributed Object-Oriented Databases and rigorous object-oriented design and development methodology (Braidic 1996);
- novel user interfaces developed using the methodology of Cognitive System Engineering (Potter et al. 1996);
- 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 of this conference – the planning and scheduling function – and the corresponding AI technologies – 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 re-design problems, circa 1990, within the SPEx design and configuration shell (Berry and Kott 1992). The key idea is to regenerate the plan while using 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 Continuity-Guided Regeneration 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, micro-opportunistic search, an instantiation of Constrained Heuristic Search (Sadeh 1991). The main steps of the micro-opportunistic 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 are 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.
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 which 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 taking into account other domain-dependent commitment measures;
3. increase reliance measure associated with this requirement-resource pair using the IPC;
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 Continuity-Guided Regeneration paradigm can be used not only as an add-on to the micro-opportunistic search but also with other planning and scheduling approaches. A key advantage of the Continuity-Guided Regeneration approach is that it eliminates the key dilemma of the re-scheduling problem: how much of the prior schedule to undo.

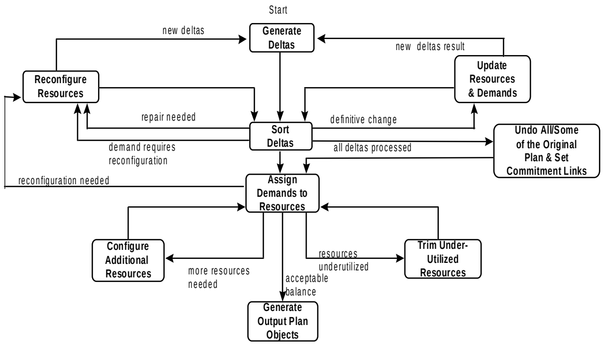

Figure 5: Overall flow of the dynamic replanning process using the Continuity-Guided Regeneration approach.

Applying the CGR Approach to the TRAC2ES Dynamic Replanning Problem

The solution process (Fig. 5) begins with the arrival of disruptive events, such as 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 (e.g., airport is open) and its actual state (e.g., 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 may create new deltas, update attributes and associations of existing resources and demands, or reconfigure existing resources (missions). For example, the following operations can be performed on missions:

- update modified missions with a complete route/schedule given;
- "repair" missions which have been diverted for some reason and for which continuation of the mission is not given;
- re-route missions, if necessary, to provide 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, i.e., we prefer to err on side of undoing more itineraries than necessary given that the Continuity-Guided Regeneration provides continuity.

At this point, the process of assigning patients to resources (e.g., missions and hospitals) is performed. It involves re-assigning Patients to resources (missions, hospitals, etc.) using the Continuity-Guided Regeneration 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 begin (to account for the Sliding Scale of Commitment).

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

Current Use of TRAC2ES

The most significant use of TRAC2ES to date has been at military exercises – 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 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 -- where TRAC2ES has been actively used and where it demonstrated its benefits. Unified Endeavor 98-1 was conducted in October/November 1997. It was a computer-assisted exercise that involved components from U.S. Army, Marine, Navy, Air Force, and military of 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. US Central Command (USCENTCOM) tasked USACOM to provide a Combined Joint Task Force (CJTF) 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 (JRDF) consisted of a reinforced brigade with maritime, air, special forces, and logistics elements in coalition with US forces.

The exercise was supported by the TRAC2ES FlyAway 1.3.1 version which supported 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 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 U.S. Each MTF was submitting patient movement requests the JPMRC component of TRAC2ES. Users and observers stated that the exercises demonstrated a number of "value added" capabilities provided by TRAC2ES (as compared to conventional near-manual operation), including:

- much more robust method of communicating patients movement requests from MTFs to the JPMRC; the ability of MTF to see if patient has been scheduled for movement and how; the ability of MTF personnel to view the entire planned itinerary of the patient;
- users at a variety of locations (including UK components) were able to obtain information about patient's information; planned movements, etc.,
- dynamic replanning capability deserved particular mention. Numerous extremely disruptive events were injected into the exercise: closure or movement of various facilities; cancellation or rescheduling of flights; changes in requested patient movements. It was common for nearly all patients previously scheduled to be disrupted due to a single event injection; many occurred while patients were already enroute. It was clear to the observers, that without TRAC2ES JPMRC officers would be unable to deal with the resulting chaos and to 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 functionality and deployment to TRAC2ES owners 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.

OP1 (Operational Prototype 1), delivered in October 1994, began prototyping TRAC2ES hardware with network communications. OP1 software, while introducing some limited features, was primarily a demonstration capability used to illustrate future directions and applications capabilities TRAC2ES may contain. OP2, delivered in June 1995, was based on a newly designed prototype architecture which includes a distributed, object-oriented database. The OP2 GUI and algorithm provided enhanced functionality including coordinated planning by one global and several theater-based centers. In fall of 1995, the next operational prototype OP3 was released, featuring a number of enhancements to the GUI (Graphic User Interface) and the planning algorithm. In 1996 and 1997, the work focused on extending the OP3 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 Unified Endeavor 98-1 we discussed in section 4 above.

From its inception, TRAC^{2}ES has involved users in shaping its "to be" vision via national award winning business process re-engineering and the incorporation of multi-disciplinary user feedback. USTRANSCOM sponsored numerous Corporate Information Management (CIM) workshops, each of which focused on re-engineering a portion of the patient regulating 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: on one hand, the user community continued to work on modifying and optimizing the business process as they learnt more about feasibility and capabilities of TRAC2ES; on the other hand, 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: a very prominent role played by the object-oriented CASE tool with automated code generation capability, and the extensive use of Object-Oriented DBMS (Braidic 1996).

Maintenance

Developers of applications involving AI techniques and domain-specific knowledge-bases often envision that the maintenance of the knowledge base will be eventually 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 are practical in the context of TRAC2ES. Although several key techniques utilized 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 indicates 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; continuous development process involving rigorous and formal cycles of requirements maintenance, design modifications via a object-oriented CASE tool, extensive reviews, implementation that involves large percentage of automatically generated code, verification and validation, quality assurance and extensive testing.

Conclusion

TRAC2ES is a very large system that serves a broad range of users worldwide, and involved several hundred person-years of development effort. While providing a range of functions and benefits, and relying on a number of advanced technology, 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 technique of Continuity-Guided Regeneration. The key idea is to regenerate the plan while using the currently executing plan 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 fully addresses the continuity issue, enables sliding scale of commitment, 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 via current near-manual methods, in the demanding environment of modern military exercises.

Acknowledgements

This work was supported by the United States Transportation Command. Leadership was provided by W. Elm and A. Mercer (Carnegie Group, Inc.), R. Simpson (MITRE), J. Simpson and C. Kirschner (USTRANSCOM).

References

Berry, F. and Kott, A. The SPEx Shell: an Approach to Automating Application and Sales Engineering Processes. In *Proceedings of the 1992 ASME Computers In Engineering Conference*. ASME.

Braidic, G. 1996. Mission-critical Patient Care. *Object Magazine,* February.

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, In *Proceedings of the Florida Artificial Intelligence Research Symposium*, Key West, FL

Potter, S., Ball, R. and Elm, W. 1996. Supporting Aeromedical Evacuation Planning Through Information Visualization, In *Proceedings of the 3rd Annual Symposium on Human Interactions with Complex Systems.*

Sadeh, N. 1991. Look-ahead Techniques for Micro-opportunistic Job Shop Scheduling. Doctoral dissertation, 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, Pittsburgh, PA.

Saks, V., Braidic, G., Kott, A. and Kirschner, C. 1997. Distributed Medical Evacuation Planning: Which Problem Should Each Agent Solve. 14th National Conference on Artificial Intelligence, Workshop on Constraints and Agents.

Saks, V., Kepner, A. and Johnson, I. 1992. Knowledge Based Distribution Planning. In American Defense Preparedness Association Conference.