Proceedings of the Symposium on Aritificial Intelligence in Military Logistics, Arlington, VA: American Defense Preparedness Assoc. pp. 177-182.

DISTRIBUTION PLANNING
An Integration of Constraint Satisfaction & Heuristic Search Techniques

Richard Camden*, Carolyn Dunmire
U. S. Army Human Engineering Laboratory (HEL)
(301) 278-5867, DSN 298-5867

Rajay Goyal, Neena Sathi, Bill Elm, Ivan Johnson*
Carnegie Group Inc.
(412) 642-6900

Dr. Mark Fox
The Robotics Institute, Carnegie Mellon University

BACKGROUND

The dynamics and complexity of logistics planning require decision support tools that do more than provide data base access or spreadsheet solutions. Techniques to aid the logistician in combat service support (CSS) planning must be capable of rapidly developing and evaluating alternate plans under battlefield conditions characterized by uncertainty, missing information, dynamic changes, too many variables, and lack of planning time. The Human Engineering Laboratory (HEL) is developing artificial intelligence (AI) technology for use in logistics *Planners* that will provide the basis for rapid logistics planning and flexible, easy-to-use "What if" capabilities in a computer-aided decision support environment including digitized terrain displays, graphical interfaces, and AI planning techniques.

OBJECTIVES

In combat, ammunition is a dominant factor in determining the outcome. The primary goal of the ammunition support structure is to provide the ammunition needed to support the plans of the tactical commander. Ammunition supply is based on a continuous refill system with stocks issued to the combat user replaced by stocks moved up from the rear. CSS planning at the corps includes the requirement to plan for the distribution of ammunition to support the tactical commander's plans and actions. The Distribution Planner described in this paper is being developed to establish the ability of AI to assist corps logisticians in developing and evaluating ammunition distribution plans over a normal corps planning horizon. The Distribution Planner seeks to satisfy stockage objectives at numerous supply points within a theater and generates transportation orders to distribute ammunition among these supply points over a user-defined planning horizon on the order of 5 to 15 days.

The purpose of the Distribution Planner is to assist the logistician in generating plans that simultaneously meet goals such as

- Meet Required Delivery Date & Quantity
- Fill highest priority needs
- Maintain stockage objectives
- Maintain safety stocks
- Making best equipment, transport, and supply route assignments

An important underlying reality is the need to plan against multiple goals, some of which may be conflicting. For example, the goal to maintain stockage objectives may conflict with the goal to reserve 20% of main supply route (MSR) capacity. Logistics resources are usually highly constrained. The objective is to provide information to the logistician that will help him guide the distribution planning process, and thus satisfy these multiple goals in as even-handed a manner as possible. The logistician and the Distribution Planner are key elements in the solution, working togethe iteratively to find a good, resilient solution the problem at hand. The Distribution Pla gives logisticians the capability - up ti not possible with only manual me explore many alternatives and vari basic plan - thus dramatically effectiveness and productivity.

KNOWLEDGE-BASED LOGISTICS PLANNING

An architecture has been developed (see Figure 1) that will support the logistician as he takes the "data" from the CSS standard systems, and the "information" from algorithmic programs, geographic information systems (GIS), and data base management systems, to effectively plan inventories and transportation capabilities. The knowledge-based planning architecture is built to address the problem of conventional ammunition support in a corps.

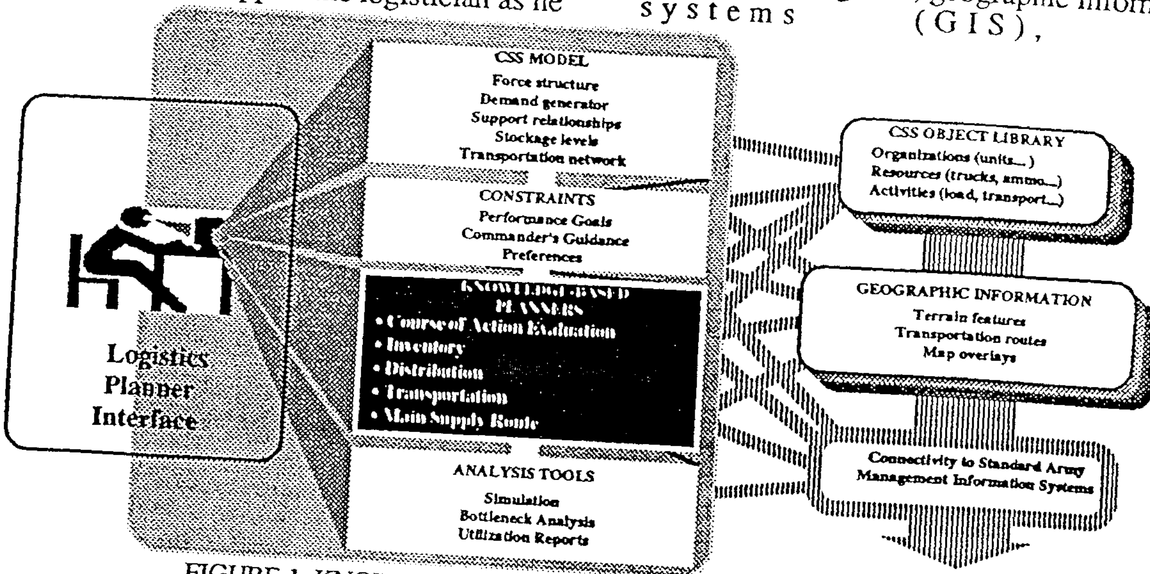

FIGURE 1 KNOWLEDGE-BASED LOGISTICS PLANNING SHELL

The Distribution Planner uses knowledge and constraint representations, data, analysis algorithms, and geographic information common to all logisticians. External data, such as inventory levels and combat requirements, will be accessed by the Planner. Model building interfaces support an interactive, graphical process of constructing specific models from a library of CSS objects. The Planner uses the knowledge represented in this CSS model to generate plans to accomplish user-specified goals. Interfaces allow the logistician to establish support relationships, specify supply site locations, specify planning goals, edit planning constraints, establish commander's guidance, and so forth. "What if" questions can be analyzed by varying the constraints, goals, or model objects to reflect varying battlefield conditions (e.g. - What if the main supply route capacity is reduced to 50%? What if the initial stockage is doubled?). Other types of plan analysis, including rule-based inspection and various statistical measures will allow the user to analyze resource utilization and identify bottlenecks and scarce resources, as well as analyze the impact of making changes to the plan.

AI IN LOGISTICS PLANNING

Planning is the selection and sequencing of actions to transition from an initial state to the desired goal state. Practical planning problems tend to be of a size that makes it impossible to completely enumerate all possible sequences of activities. Ammunition distribution plans, which consist of the ordered activities of loading, transporting, and unloading tens of thousands of tons of ammunition, using hundreds of trucks, several transport modes, over a planning horizon of 5 to 15 days, would create a combinatorially impossible problem to completely enumerate. Knowledge-based techniques make it possible to reason about the logistics concepts under consideration so that a plan is generated without completely enumerating all possibilities or mathematically solving a complete set of state equations.

When generating a plan for a real-world logistics system, it rapidly becomes apparent that the plan is really a tightly interwoven set of plans covering inventory, equipment, time, MSRs, and personnel while considering physical capacities, limitations, prerequisite conditions, and objectives distributed over time. The problem distills down to managing multiple perspectives and performing numerous trade-off analyses at the same time. Given the large number of issues to manage, performance then also depends heavily on minimizing the amount of backtracking (partial plan "unravelling") that must be performed, with the ideal being a viable plan on the first attempt.

The field of artificial intelligence includes a variety of knowledge-based approaches. Most widely known are Expert Systems, that are commonly (incorrectly) perceived as rule-based systems which emulate the performance of human experts. The planning system discussed here is very different from a rule-based expert system. Its unique characteristics are:

- The knowledge base is stored in a hierarchical inheritance network (in schemata or frames), allowing explicit representation of key concepts, relationships between concepts and entities, and abstractions of those concepts.
- The "execution" portion of the system is not represented in rules, but rather in the capabilities encoded into the objects themselves (the methods of the object-oriented approach).
- The real execution power of the system lies in the search through the knowledge structures, not in a chain of rules firing to a conclusion.
- The separation of generic concepts, specific instances of those concepts, and the problem solution techniques are explicit (and fiercely guarded).
- The interaction between major portions of the system is managed by a "blackboard architecture" where all results are freely "posted", and the total set is available for all other modules to use.

SEARCH AS A PLANNING TECHNIQUE

It is not intuitively obvious how Search can be used in generating a plan. Search is a general-purpose AI tool, (or "weak method") which has the advantage of wide applicability and the disadvantage of requiring extensive hand-tooling to apply it to a particular problem.

The underlying approach on which the Distribution Planner is built is called Constrained Heuristic Search (CHS). CHS can effectively manage the size of the problem representation, and the complexities of multiple interacting goals to allow generation of a successful plan, in a reasonable time. This is essential for a practical, fieldable system. The "constraint" concept proves to be versatile, representing knowledge about the physical world, the users' goals, and so forth.

CHS refers to the software-based problem solver, or "algorithm", which is guided toward good solutions by a number of different constraints that represent physical objects, limited resources (people, time, materials, etc.) and operating objectives, policies, and preferences. A key part of this emerging technology is knowledge representation, which provides a clear, straightforward means for representing the constraints and objectives (or "modeling" the environment or problem domain) in a way that's easy to define, modify, and extend.

The CHS approach cycles repetitively through the problem domain, gradually building up a plan across a user-specified time horizon. The CHS looks at the whole problem all the time, successfully turning its attention to those elements most in need of resolution before dealing with less constrained problem elements. Hence, it is data driven and responds "opportunistically" to different problem scenarios quire flexibly, that is., without the need for the user to intervene when problem circumstances change from time to time.

CHS is a powerful alternative to so-called "conventional" approaches to solving complex planning, scheduling, and logistics problems. Techniques including linear programming (LP) and various forms of integer programming (IP) have been used for many years in attempts to achieve optimal solutions to these very complex problems. CHS provides the means to provide excellent (sub-optimal) solutions to problems that cannot be solved quickly (or at all) with LP and IP techniques.

At one level, CHS can be used to solve problems that cannot be formulated in LP terms at all or that would require an inordinate amount of time to compute a solution; by the time the solution is available, it no longer applies since circumstances have changed in the meantime. CHS formulation also can make it easier for the user to change the problem definition without relying on software engineers to re-code the problem.

CHS provides a straightforward means to represent constraints and objectives that are different in kind from each other (some objectives are in dollars, e.g, others in time) and that often conflict with each other. CHS lets the problem solver work through the solution "opportunistically" looking at the data at hand and relying much less on prescribed definition of the problem and the structure of its solution, as in more conventional approaches.

PROBLEM-SOLVER

The Distribution Planner takes supply demands over time as input and determines how to allocate critical items and set stockage objectives for storage locations. The planning process is constrained by the available MHE lift capacity at each supply point and by user-specified "policies" (e.g., ATPs should be stocked directly from the CSA) which can be "relaxed" by the planner, if needed, to produce a better plan. The problem solving methodology combines constraint satisfaction and heuristic search techniques.

The DP constructs a constraint network that defines the interactions among a variety of physical entities and operating objectives, represented by variables and constraints. This constraint network represents time-varying demands for ammunition rounds at each supply point. This network is extended by including physical constraints (such as transport and material handling capacity limitations), preferential constraints (e.g., maintaining safety stock levels), and causal (or pre-condition) constraints. The internal problem-solver cyclically determines the appropriate node(s) to focus on (temporally and physically) to make (the next) planning decision(s). In effect, the problem-solver looks at the whole problem across the entire time frame under consideration on each decision cycle, chooses those elements most in need of assignment, updates the plan being generated, and returns to re-evaluate the plan and choose the next decision point. The planner cycles through the problem until orders and/or resources are exhausted. The resources being managed include

- Anticipated demand over time by ammunition type,
- Ammunition consumption rate,
- Locations within the area of operations,
- Source and destination locations,
- Current inventory status,
- Material Handling Equipment capabilities,
- Controlled and required supply rates,
- Ammo type and combat unit priorities, and
- Inter-supply point transportation assets capacities.

DISTRIBUTION PLANNER ARCHITECTURE

There are five major modules in the Distribution Planner:

- CSS Model
  - Corps Task Organization
  - Support Relationships (Network)
  - Activities Relationships (Network)
  - Weapons Systems Definition
- Stocking Planner
- Sourcing Planner
- Analysis Sub-System
- User Interfaces and Displays

CSS MODEL

The CSS model contains the underlying knowledge and data used by the Distribution Planner. This model is constructed from a set of "objects" that represent specific entities in the problem domain: storage locations, trucks, convoy definitions, corps hierarchy definition, point-to-point transportation routes, material handling equipment, and so forth. "The Library" contains the following classes of objects:

- Army Organization Hierarchy: corps, division, brigade, company, battalion
- Resources: MHE, trucks, people, ammunition, and so forth.

- Activities: load, ship, unload, inventory consumption

STOCKING PLANNER

The Stocking Planner computes the stockage objectives of all components at all supply points over the user-specified time horizon. Currently, this is based on the commander's guidance regarding how much ammunition (by type) should be stored at supply points and ammunition demand (by type) over time. A Demand Generator (DG) was created to generate demand by ammunition type over time. (The DG was implemented to provide the means to create demand data quickly and conveniently for testing and stressing the Distribution Planner.) The Stocking Planner uses demand factors and the commander's guidance inputs to translate these DG-provided demand profiles into stockage objectives at the defined supply points over planning time frame.

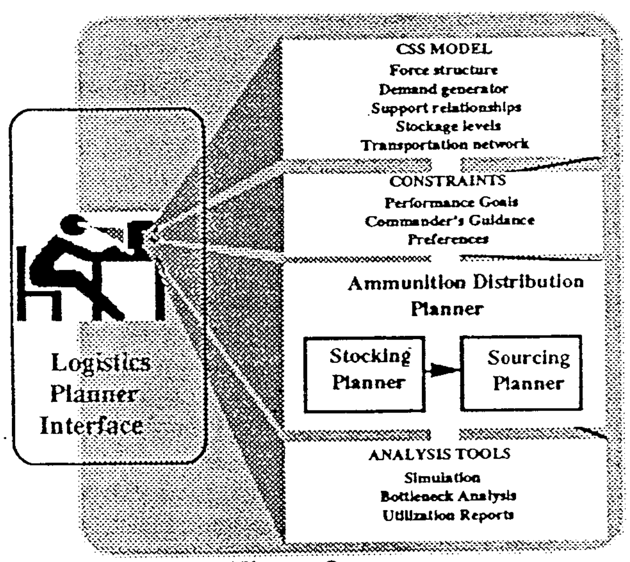

Figure 2
Distribution Planner Architecture

SOURCING PLANNER

The Sourcing Planner, the heart of the Distribution Planner, has been implemented as a blackboard architecture. Resources are modeled as knowledge sources, or "objects". The constraint-directed planning process cyclically evaluates the net demand of orders for the key resources modeled, and makes those resource assignments that are most in need of being dealt with on this decision cycle. This can also be thought of as a "Least Commitment" approach; the planner makes resource order assignments only for those elements that are "tightest"; those orders with more flexibility or multiple options are deferred until later decision cycles.

Knowledge sources "vote" for each order, in effect making a statement from their respective points of view how important it is to deal with each order on this decision cycle. In the process, a contention matrix is developed, which in turn is further elaborated to include an order due date and priority factor; this results in a "goodness matrix." The order with the highest "goodness" (urgency, importance) is selected; timebounds are then established for all activities dealing with this order.

Hence, the problem-solving methodology for the Sourcing Planner is a combination of heuristic search and constraint satisfaction. A constraint network is constructed that defines the interactions among a variety of variables and constraints. The basic constraint network is constructed from temporal demand and resource availability relative to stocking and shipping ammunition components. The goal is to identify where (in temporal and physical space) to make the next planning decision and how successive decisions affect each other.

The problem-solver consists of two inter-connected networks: Resource Network and Activity Network. The Resource Network (RN) models each battle unit, storage facility, and transportation route as nodes. A node has a number of parameters defining resources (inventories, trucks, MHE, etc.) Arcs connect transportation routes to storage facilities and combat units. The Activity Network (AN) models inventory moves and dynamic demand for inventory. Each inventory fill request is modeled by a single activity, with properties defining the inventory requested, quantity, source and/or destination, and due date. This activity is instantiated into an and/or graph of activity instances representing alternate means of satisfying the request by alternate sources and routes. Each activity instance is elaborated into three sub-activities corresponding to load, transport and unload activities.

The Activity Network is connected to the Resource Network by the demand each activity places on resources. In particular, a load activity requires available inventory at a storage facility, MHE at the facility, and sufficient transport capacity.

In summary, the problem-solving steps are

(1) Propagate demand for ammunition at ATPs and other first level ammunition sources;

(2) Instantiate inventory move activities; elaborate the instantiated activities into load, transport, and unload sub-activities; propagate demand for resources onto the resource network;

(3) Compute Resource Network "contention" values, due to net demand versus net resource availability;

(4) Compute Activity Network "goodness" values, based on activity's priority and the degree of contention provoked by the resources it demands;

(5) Select the best activity to schedule; identify the time window during which it is to be performed; and

(6) Propagate the impact of the scheduling decision throughout the networks; update inventories and demands; go to step 2 (if new demand arises) or to step 3.

USER INTERFACES AND DISPLAYS

The principal output display is a symbolic map of storage locations and combat units with links representing the CSS support relationships. Division boundaries, main supply routes, and other features are also presented. The information available (selectable on user command) can be a snapshot of the state of the distribution plan at a user-selected time, or a summary of the overall plan. A time scale is available that enables the user to select the time tick for which the snapshot information is displayed.

The results of the Distribution Planner are reports and analyses for use by the logistician in evaluating supportability of operations and in conducting "What if" analyses:

- Stockage Profiles versus Objectives (by type and quantity) at each supply point during the planning time frame.
- A series of transportation orders, consisting of
  - Point and Date of Departure
  - Point and Date of Arrival
  - Ammunition type & quantity (components, by ton)
  - Use of Aggregate Truck and Roads Capacity (ton-hours)
- Summary and analysis of planning results (percent safety stock achieved, truck and roads utilization, MHE utilization, timeliness of ammo availability at front lines, etc.)

PERFORMANCE

The following table provides a summary of performance of the Distribution Planner for different planning horizons:

Planning Horizon	Constraints	CPU-TIME (SUN-4 Work-Station)
2 days	3,000	65 secs
5 days	7,000	388 secs
7 days	10,000	571 secs.

These results compare very favorably with the planning cycle currently experienced with manual approaches. The typical planning cycle can take from 1 to 2 days of intense effort to develop just one reasonable corps distribution plan. Using the Ammunition Distribution Planner as an on-line interactive decision aid, battle planners can explore dozens of alternatives in the time that currently only one or two may be developed and evaluated.

SUMMARY

Knowledge-based planning techniques are AI's unique contribution to the solution of problems that cannot be solved by deterministic algorithms. They are appropriate to problems that have numerous and conflicting goals, a high volume of data, activity bottlenecks and resource shortages. The authors believe that it is in these critical areas that AI can be most judiciously applied by the Army.

For more information, contact Richard Camden, US Army Human Engineering Lab, AV 298-5867, 301-278-5867, net address rcamden@hel4.brl.mil.