Overview: ARPA-Rome Laboratory Knowledge-Based Planning and Scheduling Initiative (ARPI)

Northrup Fowler III
Rome Laboratory
RL/C3C
525 Brooks Road
Rome, NY. 13441-4505
fowler@ai.rl.af.mil

Thomas D. Garvey
ARPA
Information Systems Office
3701 N. Fairfax Dr.
Arlington VA. 22203-1714
tgarvey@arpa.mil

Stephen E. Cross
Carnegie Mellon University
Dir. Information Tech. Cntr.
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213-3891
scross@cs.cmu.edu

Mark Hoffman
ISX Corp.
1165 Northchase Pkwy, Ste. 120
Marietta, GA. 30067
mhoffman@isx.com

Abstract

This dedicated volume of the AIPS '96 proceedings is devoted to a progress report on the ARPA / Rome Laboratory Knowledge-based Planning and Scheduling Initiative (or ARPI for short). The ARPI has been co-sponsored by the Advanced Research Projects Agency (ARPA) and the United States Air Force Rome Laboratory (RL) since 1989. The ARPI's main purpose is the development of the next generation of fundamental artificial intelligence planning, scheduling, and resource allocation technology. An equally important secondary goal is to demonstrate significant capability improvements over current planning trends in operational domains. This paper provides a historical background and the founding principles and visions for the ARPI which is now in its third phase.

BACKGROUND

In 1986, RL and the Air Force European Office of Aerospace Research and Development co-sponsored a workshop entitled the "Future of Expert Systems Workshop" [1] in Edinburgh, Scotland, that brought together leading U.S. and European researchers to discuss the state of the art in expert systems and to identify emerging technologies considered vital to future directions.

This article is based on material which was first published in IEEE Expert, vol. 10, No. 1, pp. 4-9, February, 1995 which is (c)Coryright 1995 IEEE. It is used with permission of IEEE.

Three discussion areas included very large knowledge base systems, real time problem solving, and architectures for large integrated expert systems. While the workshop was itself useful and interesting, side discussions between several of the participants (Paul Cohen, Nort Fowler, Austin Tate, James Allen, and Steve Cross) laid the groundwork for a major initiative in knowledge-based planning and scheduling.

During several late night brainstorming sessions, we discussed ways to improve AI research methodology, how AI research could be better focused on large scale problems, and how to improve technology transition. We conceived a framework for a new government / university / industry research and development partnership. Elements of this ideal program addressed the following issues and concerns.

- Planning in the real world is ubiquitous and hard.
- Real world planning requires a combination of approaches.
- Faster progress is possible via a shared, large scale problem
- Metrics-based evaluation is fundamental for sustained research progress.
- Success of the program is dependent on the early involvement of the end-user.
- Technology transition is every researcher's responsibility.

We made several attempts at refining these visions in the ensuing years. In 1987, we sponsored a Workshop at AAAI on Real-time Problem Solving. Later that year, the USAF (through the Rome Laboratory, the Wright Laboratory and the Air Force Office of Scientific Research) initiated a program in Real-time Problem Solving. [2] In parallel, the government members of the discussion group worked with the DARPA Strategic Computing Program which also sponsored a small knowledge-based planning component [3].

In 1990, several events transpired which provided us with the support required to implement the program we envisioned in 1986. Rome Laboratory and ARPA secured approval for the research program. Program approval was aided by a previous planning workshop that was held in Clearwater, Florida, in December 1989 [4]. The decision at this meeting to focus on military crisis planning, and in particular transportation planning, scheduling, and analysis, was most fortuitous. Technology assessments conducted by RL in early 1990 highlighted this need at both the Military Airlift Command (now the Air Mobility Command) and the US Transportation Command, both at Scott AFB Illinois. Transportation planners described the need for intelligent, interactive support for planning. They were becoming very anxious and frustrated by currently fielded systems and continued delays in promised improvements. The advent of Operation Desert Shield (ODS) in August 1990 provided the motivation to both address new ways to build software systems and insert new technology. ODS provided the opportunity for us to jump-start the program with a successful technology demonstration. This demonstration in turn resulted in several critical program features:

- the creation of a shared vision between end-users and technologists of what planning technology might achieve,
- a research methodology to transfer needs of an end-user community back into the research community,
- and a software development process to rapidly insert research results into demonstrable (and transferable) technology demonstrations.

A snapshot of the technical progress within the initiative through Feb. 1994, Phases I & II of the ARPI, which included 35 research groups, is described in [5]. This AIPS special volume reflects an update of that progress, carrying over into Phase III up through Feb. 1996.

ARPI's VISIONS

The initial thoughts on the program's management structure were discussed at a 1989 workshop between AI researchers, application builders, and potential military users.

> The AI planning community believes that it has many of the constituent theories in place, but what has yet to be demonstrated is what is important and what is not. What is significant about this initiative is that by providing the AI community with a problem rich in complexity, and with performance goals, it will be able to measure the efficacy of its approaches, and demonstrate to the application communities what can be done. The program will enable the determination of how planning techniques respond to issues such as scale, uncertainty, resource contention, optimization, learning, etc. Secondly, by creating a set of demonstration programs based on this domain, the field will demonstrate to the application community that the technology is relevant and worth pursuing. [4:54]

Vision #1: Integrated, well engineered planning modules in an application ready form.

A major goal of the initiative was to develop a well engineered, generic set of knowledge-based planning, scheduling, and resource allocation tools. Previous research, for example in classical (or generative planning), was concerned with reasoning about how actions affect the world. But resultant AI systems would address only a portion of the real world planning problem. For example, classical planning methods would generate a plausible plan but give little assurance that the plan could be implemented efficiently in the world. Constraint-based scheduling methods provided techniques to efficiently allocate and monitor the use of resources but assumed as input a complete plan. Simple, loose coupling approaches between existing planning and scheduling methods would not scale to large problems - a tightly integrated approach that would allow resource considerations to focus planning during plan generation was desired. In addition to resource allocation considerations, real world planning also involves managing multiple risks and maintaining plan feasibility from multiple perspectives. Encouraged by users, the ARPI recognized opportunities for progress in the early informing of the plan generation process of all these concerns in order to avoid the wasted effort of generating a classically plausible plan that was destined to be discarded for pragmatic reasons. Almost no previous work had addressed the system engineering and architecture (or framework) of an integrated planning system and the pragmatic issues of database integration.

Thus, the ARPI research program is extending and integrating several planning paradigms in artificial intelligence. Classical plan generation methods are being extended by incorporating reasoning under uncertainty, decision theory, and plan justification methods. This allows flexible operation plans to be created with precompiled options. Constraint directed scheduling techniques are being integrated into current transportation scheduling and analysis programs to provide execution monitoring and replanning support. Case-based planning methods are being demonstrated on plan reuse tasks, such as plan assembly and subplan reuse with force module libraries. These methods are being made available in an environment which provides uniform intelligent access to sources of planning data and allows simulations of plans to be composed from constituent models and executed to test and evaluate planning options. It is significant that each of these paradigms represents much more powerful capabilities than the more modest forward-chaining rule-based approaches that have been previously applied to real world planning tasks.

Vision #2: Apply TQM to Research Management.

In 1989, there was general agreement that AI research community was poised to break out of its preoccupation with small, clean abstract domains and that a focused community effort on a shared realistic problem of large scale would provide a new insights. It was also felt that there needed to be more emphasis on metrics-based evaluation - both as an integral part of an individual research methodology and as a means to gauge progress in the overall field. Many of the new areas for breakthroughs in AI were felt to lay at the boundaries of the current research areas, but incentives and infrastructure needed to be provided as catalysts to facilitate a community working on a shared problem.

Thus a second vision was to apply the principles of Total Quality Management (goal directed, measurable progress towards goals, "drive out fear", cooperation among stakeholders) to the achievement of the ARPI goals. We first recognized that the initiative contained many organizational specific goals that were representative of the groups in the research community (universities, corporate research laboratories, and federally funded research labs), systems integration community, government funding organizations and government laboratories, and end-user communities. Our management goal was to define an approach that allowed each participant to achieve their individual goals and at the same time encouraged support of consensus based community goals. For example, the government funders were concerned about demonstrating the utility of the supported research in domains of interest to potential end-users but were also sensitive to the "demo or die" mentality of previous multiple organization research program. The program structure was thus defined along three tiers: a research tier, a technology transfer tier, and a demonstration tier as shown in Figure 1.

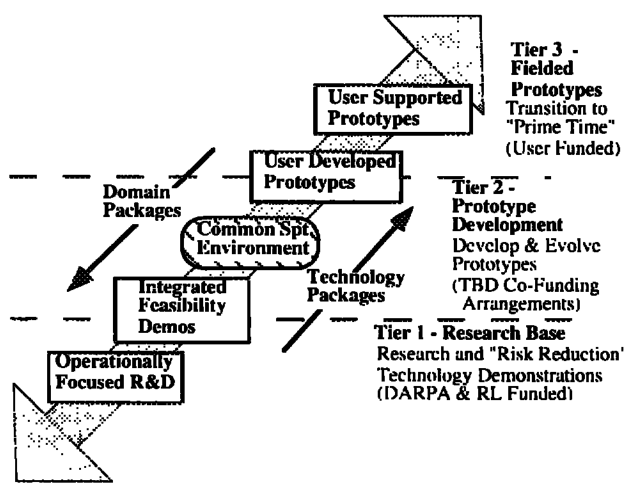

Figure 1: ARPI Program Philosophy

The ARPI research and development process is driven by a series of Integrated Feasibility Demonstrations (IFDs) and Technology Integration Experiments (TIEs) which assess technical progress and evaluate its operational impact. A Common Prototyping Environment (CPE) was developed during Phases I & II and an ACPT Testbed Environment during Phase III to support demonstration and testing of technologies; experimental system integration and evaluation activities; and re-use of databases, knowledge bases, software modules, and test scenarios.

Tier 1 includes a number of independent research projects that are oriented toward developing operationally focused knowledge-based reasoning technology that addresses critical problems in military planning and scheduling. The exit criteria for technology migration to *Tier 2* include successful demonstration of capabilities in research oriented TIEs. The *Tier 2* effort consists of the IFDs and CPE-supported activities which evaluate technical progress, merge the individual developments in Tier 1 into experimental systems and rapid prototypes, and integrate ARPI developed technologies with other components to address specific operational problems. An IFD shows the operational communities new planning and scheduling capabilities to obtain constructive feedback on their applicability to critical operational functionality measured against criteria for success defined by end users. *Tier 3* involves the user-guided insertion of ARPI technology and systems into user-supported operational prototypes.

Vision #3: A new paradigm in crisis management.

Planning to resolve a crisis generally is viewed as beginning with a stage in which the situation is assessed, high level mission objectives are determined, and possible constraints on means and methods for accomplishing those objectives are considered (see Figure 2).

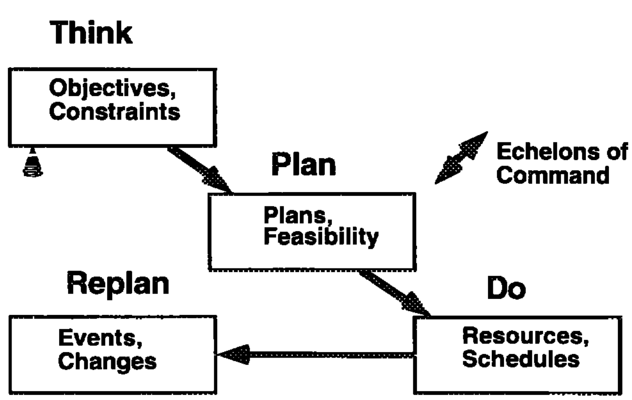

Figure 2: Serial planning process

The resulting mission objectives and planning constraints are passed to a plan generation stage, in which notional assets are allocated to one or more candidate courses of action and feasibility is analyzed from several perspectives, including transportation, resource utilization, and effectiveness. Once a feasible plan has resulted, actual forces are assigned, actual schedules and movement orders generated, and the execution of the plan can begin. As the plan is executed, events occur, and the situation changes. A replanning cycle begins; and the plan is modified by iterating the process described above.

The serial, sequential planning process and the hierarchical organization of details just described may conform to the way in which plans are represented and studied, but it does not accurately characterize the information processing and communication that occurs during planning. In military planning, as in ordinary business planning, simultaneous communication among levels of command and across areas of responsibility is commonplace, and many elements of the plan are developed and evaluated concurrently. Even when explicit communication and collaboration do not occur, planners actively model the potential decisions or data inputs of other participants in the planning process. (See Figure 3)

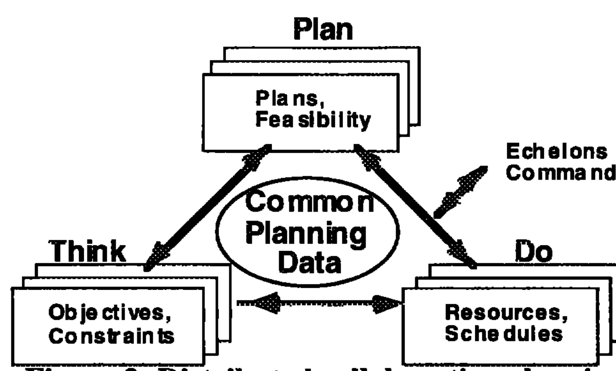

Figure 3: Distributed collaborative planning process

Instead of a sequence of discrete stages, with intermediate stored products and local control mechanisms, planning is seen as a continuous process with predictive control between stages and no intermediate products. Quick communication and a visible overall rationale are shared among stages to provide rapid correction or modification of the process as a whole in order to maintain overall efficiency. In practice, this improved interprocess connectivity and overall process visibility reduce intermediate error, increase overall efficiency, and enhance the ability of the process as a whole to adapt to new conditions. The collaborative planning model thus encourages optimization of individual components of the process in the context of optimizing the overall process.

While the goal of providing technical solutions to individual tasks involved in the military planning process remains a primary consideration in establishing technical requirements for new technology, the identification of the task requirements on which to focus is strongly influenced by overall process considerations: one requirement may contribute more strongly to overall productivity or plan quality than another. Thus, we worked with military end-users to define visionary planning systems that supported human specialists and computer based planning aids to work in a distributive, collaborative infrastructure to support the rapid creation, analysis, execution, and evaluation of a plan. Critical to the success of the program were the definition of functional requirements that were dependent on the ARPI development of "missing technology." Some specific correlations between operational planning requirements and knowledge-based scheduling and planning technology are shown in Figure 4.

AI/Planning Research World	Joint Planning and Operations World
Generative Planning	Commander's Objectives, Concept of Operations, Force/resource selection and re-use, objectives/task decomposition
Constraint-Based Planning	Resource Constraint Analysis, Feasibility Analysis, time-phasing , etc.
Case-Based Planning	Force Analysis, Planning module Library Development, Failure Analysis, Plan Revision Techniques
Intelligent & Object-Oriented Databases	Distributed, Heterogeneous Intelligence and Situation Assessment Databases
Interactive Graphics and Editing of timelines, schedules, resources, maps, graphs, representations, etc.	Manual Data Analysis, Plan Refinement, and Briefing Production

Figure 4. Specific Connections Between Researchers and Planners.

Vision 4: Accelerate delivery to end-users in a supportable form

Integrated feasibility demonstrations (IFDs) were the product of the program. However, the success of the IFDs provided the incentive for the end-users to advocate for early insertion and support of fieldable prototypes based on the IFDs. Based on our IFD experience, we created an intelligent systems engineering process, characterized by direct, early, and frequent user involvement through rapid prototyping and short development cycles, which we call "user-centered software development." The elements of this process are:

- Conduct a scenario-based domain analysis
- Define the intelligent systems architecture
- Prototype intelligent systems components
- Maximize re-use and re-engineering
- Conduct frequent user-centered testing
- Deliver it in an evolvable form with support tools

CURRENT PROGRAM STATUS

In response to the 1990 Broad Agency Announcement, more than 20 organizations were selected to participate in this program. Approximately a dozen more were added after a second Broad Agency Announcement was circulated in 1992. In 1994, a third BAA was circulated with a strong focus on mixed initiative planning and scheduling technology development. Again, nearly 30 organizations were selected to participate in phase III.

The individual Technology Developers shown in Figure 5 are pursuing operationally focused basic research. The research community is self-organized into "clusters" of activity or focus.

- Collaborative Planning Working Group
- ACP Scheduling Cluster
- Plan Evaluation, Explanation and Guidance
- Simulation Cluster
- Case-Based Planning Cluster
- Uncertainty Cluster
- Transportation Cluster
- Ontology Cluster

Membership in more than one cluster is encouraged, as are cross cluster activities and collaboration.

AIAI - Tate/Drabble	**NWU - Birnbaum/Edwards**
Brown - Dean/Kaelbling	**ORA - Calistri-Yeh/Segre**
CMU - Navin-chandra/Sycara	**Oregon - Ginsberg/Crawford**
CMU - Smith/Veloso	**Oregon - Ginsberg**
GE - Stillman/Bonissone	**Pitt - Pollack/Znati**
ISI - Swartout/Gil	**Rochester - Allen**
ISX - Edwards/Hoffman	**Rockwell - Darwiche/Goldszmidt**
Kestrel - Smith/Kambhampati	**SRI - Bienkowski**
Klein - Miller/Militello	**SRI - desJardins**
Loral - McKay/Finin	**SRI - Myers**
Maryland - Hendler/Kambhampati	**SRI - Wilkins**
Massachusetts - Cohen	**Stanford - McCarthy**
MITRE - Mulvehill	**Wash - Etzioni/Hanks/Weld**

Figure 5: ARPI Phase III Participants

The ACPT Testbed Environment provides a World Wide Web (WWW) based set of tools, resources, and domain knowledge assembled and developed by ISX Corp., the integration contractor. The ACPT Testbed Encironment is shared among research participants and functions in the capacity of both a domain familiarization tool as well as an infrastructure for demonstrating technical and operational progress through TIEs and IFDs. To date three Integrated Feasibility Demonstrations have been conducted and a fourth is currently under development.

IFD-1: DART

The Dynamic Analysis and Replanning Tool (DART), set a baseline that led to the initiative vision and that demonstrated an ability to support transportation planning and analysis for the deployment portion of military planning. The joint command that participated in the development was the U.S. Transportation Command (USTRANSCOM), and DART is now deployed to a variety of joint sites, such as the US Pacific Command (USPACOM) and the US European Command (EUCOM).

IFD-2: SOCAP

The System for Operations Crisis Action Planning (SOCAP), moved toward the initiative vision of interoperability by demonstrating (1) the first steps towards an explicit, reusable computer representation of a course of action (COA), (2) automated assistance for the generation of the COA, and (3) the first explicit reasoning thread from the COA to the transportation data used by TRANSCOM. This thread enables immediate transportation analysis and feedback to the COA generation process. The participating operational joint command was U.S. Central Command (CENTCOM).

IFD-3: TARGET

IFD-3 moved further toward the initiative vision of distributed planning by demonstrating new technology to help resolve resource conflicts arising in a multiple deployment scenario and to further assist in employment COA generation. The demonstrated system, called Theater-level Analysis, Replanning and Graphic Execution Toolbox (TARGET), was designed to provide support for the U.S. Commander-in-Chief (CINC), Pacific Command (CINCPAC) Joint Task Force (JTF) concept of operations.

IFD-4: Extended ACPT

Phase III of the ARPI has several additional foci of attention beyond that of phases I & II (which concentrated on distributed, collaborative planning). Current areas of focus for Phase III include mixed-initiative planning with both human and computer agents, development of an operational common plan representation, tighter coupling of planning, execution, and replanning, and visualization tools to aid in each of these areas.

Phase III further reflects a change in domain from that of transportation planning and course of action (COA) development (pursued in earlier phases) to that of air campaign planning.

The precursor to IFD-4 is the recently developed plan authoring tool that supports air campaign planning (the Air Campaign Planning Tool, ACPT). This tool helps an air campaign planner, at the strategic level, analyze and capture all the required information needed to generate a prioritized master target list for a variety of scenarios. This target list can then be analyzed by a legacy software system for its effectiveness and resource utilization.

IFD-4 is being designed at the time of this article but will likely incorporate technology contributions from SRI (generative planning), ISI (plan critics), and GE (temporal constraint management and visualization) to the ACPT developed by ISX. IFD-4 is due to complete in the summer of 1996.

CONCLUSIONS

As we hope this special volume of the AIPS proceedings will show, the ARPI has established a new vision of success in each of the many diverse communities it embraces. It has provided a common framework wherein all technology R&D stakeholders share expectations and responsibilities, renewed enthusiasm and success. It has initiated a critical process review that defines both the terms and directions across a spectrum of future activities. It has, by example, enriched the dialogue and space of opportunities of other programs that are adopting the ARPI spirit of teamwork and accomplishment. It has been repeatedly singled out for recognition of its bold practice of uniting people of talent and vision in a common quest for excellence that translates directly into significant improvement in the process and procedures of its ultimate customers, the end users.

REFERENCES

[1] Milne, R. (Ed.) 1986. Report on the Future of Expert Systems, Workshop Report, Intelligent Applications LTD, Edinburgh Scotland.

[2] Erman, Lee D. (Ed.) 1990. Intelligent Real-Time Problem Solving (IRTPS), Technical Report, TTR-ISE-90-101, Published for the Air Force by Cimflex Teknowledge Corp.

[3] Dunlap, R. (Ed) DARPA Knowledge-based Planning Workshop, Austin Texas, Dec 1987.

[4] Fox, M., and Cross, S. (Ed) 1990. unpublished technical report.

[5] Burnstein, M (Ed) 1994. ARPA / Rome Laboratory Knowledge-based Planning and Scheduling Initiative, Workshop Proceedings, San Francisco, CA.: Morgan Kaufmann Publishers.