
A Review of European Applications of Artificial Intelligence to Space


Mark Drummond and Helen Stewart, Eds., Ames Research Center, Moffett Field, California

October 1993



National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035-1000



LIST OF AUTHORS

David Atkinson
Richard Doyle
Mark Drummond
Bob Engelmores
Peter Friedland
Kathy Healey
Astrid Heard
Mel Montemerlo
Monte Zweben

PRECEDING PAGE BLANK NOT FILMED

CLASS ... INTERNATIONAL ... iii





TABLE OF CONTENTS

	Page
PREFACE	vii
THE NASA TEAM.....	xi
REPRESENTATIVES OF EUROPEAN ORGANIZATIONS VISITED	xiii
SITE REPORTS.....	1
MATRA Marconi Space	3
Arianespace	11
MBB Deutsche Aerospace	15
European Space Operations Center (ESOC).....	19
Vega Space Systems	25
AI Applications Institute (AIAI).....	29
Computer Resources International (CRI)	33
European Space Research & Technology Center (ESTEC)	39
BSO/Origin	43
SUMMARY REPORTS	47
European Space/AI Applications: A Summary	49
European Space/AI Technology: A Summary	55

PREVIOUS PAGE BLANK NOT ENRICHED



PREFACE

The purpose of this report is to describe the applications of Artificial Intelligence (AI) to the European Space program that are being developed or have been developed. This report describes the results of a study sponsored by the Artificial Intelligence Research and Development program of NASA's Office of Advanced Concepts and Technology (OACT). The report is divided into two sections. The first consists of site reports, which are descriptions of the AI applications we saw at each place we visited. The second section consists of two summaries which synthesize the information in the site reports by organizing this information in two different ways. The first organizes the material in terms of the type of application, e.g., data analysis, planning and scheduling, and procedure management. The second organizes the material in terms of the component technologies of Artificial Intelligence which the applications used, e.g., knowledge based systems, model based reasoning, procedural reasoning, etc.

This Preface provides the reader with the context in which the study was undertaken and carried out. NASA's AI R&D program is responsible for the development of AI technologies and for their application to NASA projects for the purpose of reducing the cost of operations and for increasing the return on investment in Space science projects. The former is done through AI applications in areas such as intelligent design assistance, automated fault diagnosis, planning, and scheduling. The latter is done through the use of intelligent tools for science data archiving, retrieval, analysis, and visualization. The program is responsible for fundamental research into new AI technology which will be useful to NASA, and for the infusion of existing AI technology into on-going and planned NASA projects in which that technology yields reduced cost, improved operational capability, or increased levels of scientific study. A brief look at the history of this program will provide the context as to why this review of applications of AI to European Space projects was undertaken.

In 1985, NASA initiated an Artificial Intelligence R&D program. This was done in response to a request from Congress that NASA develop and implement advanced automation technologies for use on Space Station and in the Space program. I was named the AI program manager at its inception, and have held that position ever since. Ames Research Center was named the Lead Center, and Dr. Peter Friedland was brought on to build and lead a team of AI researchers at Ames which could fill the Lead Center role. Ames Research Center and the Jet Propulsion Laboratory were charged with developing AI technology both through internal R&D groups and through funding industry and academia. The technology was to be applied at all of the NASA Centers, including the Kennedy Space Center, the Johnson Space Center, Marshall Space Flight Center, and Goddard Space Flight Center, as well as Ames and JPL.

An early planning effort was undertaken to determine the "best" areas in which to apply AI technology to the Space Station project. Dr. Friedland and a group of consultants including Brad Allen, Bruce Bullock, Jaime Carbonell, Robert Engelmores, David Mishevich, and Ben Wah, carried out the study. They analyzed the opportunities across all of Space Station and they came up with a set of recommended applications.

In 1990, I met Francois Allard, who manages an Artificial Intelligence applications program at ESTEC (European Space Research and Technology Center). Mr. Allard told me that he knew of the

results of the study that Peter Friedland had led, and that he thought their approach was very good. As a result he implemented a similar approach to determine an appropriate set of applications for his program to undertake. In 1991, I asked Mr. Allard if a team of NASA AI people could visit the sites of the research and development under his program in Europe. Mr. Allard agreed and helped arrange the visits. The itinerary of the trip was extended to include some projects which were not under the ESTEC program.

The NASA team was divided into two groups, one with four people, the other with five. Each group made a series of site visits. The first group consisted of:

- Peter Friedland, Ames Research Center
- Astrid Heard, Kennedy Space Center
- Richard Doyle, Jet Propulsion Laboratory
- Mark Drummond, Ames Research Center

This group visited:

- AI Applications Institute, Edinburgh, Scotland
- CRI, Borkerod, Denmark
- ESTEC, Noordwijk, the Netherlands
- BSO/Origin, Nieuwegein, the Netherlands
- Vega Space Systems, Harpenden, U.K.

The second group consisted of:

- Melvin Montemerlo, NASA Headquarters,
- Katherine Jurica, Johnson Space Center
- Robert Englemore, Stanford University
- Monte Zweben, Ames Research Center

This group visited:

- CNES, Toulouse, France
- Matra Marconi Space France, Toulouse, France
- Arianespace, Evry, France
- ESOC, Darmstadt, Germany
- MBB-Erno, Bremen, Germany

At each of the sites, the hosts were given an overview of the NASA program, and then the NASA team was given a briefing and demonstration of the applications being generated there. Site reports were written on each place visited, and then sent back to that site to be verified for correctness. The purpose of the study was to develop a written description of the European applications of artificial intelligence to Space program operations that could serve as a database for all those who are interested. It was not to compare the US and European programs or to evaluate the European projects.

I offered Mr. Allard the opportunity to have a European team review the NASA AI applications. The review was planned and scheduled, but other circumstances kept the European team from being able to come to the United States at that time. The review will be rescheduled.

There are a number of people who are responsible for the success of our review. First and foremost are Mr. Allard, who was instrumental in setting up the visits in Europe, and our hosts at the sites we visited. They were:

- Michel Maurette, CNES
- Jean-Michel Darroy, Matra Marconi
- Christian Parquet, Arianespace
- Albrecht Kellner, MBB-Erno
- Herwig Laue, ESOC
- Roger Thompson, Vega Space Systems
- Austin Tate, AI Applications Institute
- Mogens Nielsen, CRI
- Francois Allard, ESTEC
- Tim Grant, BSO/Origin

I would like to thank Mr. Allard for all of his help in setting up the review and to thank all of our hosts for their excellent presentations and discussions, and for their hospitality. Without their help, this report would not have been possible.

I would also like to thank Mark Drummond for doing an outstanding job of organizing this study, and to thank him and Helen Stewart for editing the report. Finally, I would like to thank the NASA team for their tireless work not only during a week in which travel was a nightly occurrence, but also during the ensuing weeks during which this report was written. I hope that the resulting report is as valuable to those who read it, both in Europe and in the United States, as it is to those of us who participated in the study.

Melvin D. Montemerlo
Manager, Artificial Intelligence R&D, Code CD
300 E Street, SW, NASA Headquarters, Washington, DC 20546



THE NASA TEAM

David Atkinson
Jet Propulsion Laboratory
4800 Oak Grove Drive
Mail Stop 525-3660
Pasadena, CA 91109
Tel: (818) 306-6131
Fax: (818) 306-6912
Email: datkinson@nasamail.nasa.gov

Richard Doyle
Jet Propulsion Laboratory
4800 Oak Grove Drive
Mail Stop 525-3660
Pasadena, CA 91109
Tel: (818) 306-6149
Fax: (818) 306-6912
Email: rdoyle@isd.jpl.nasa.gov

Mark Drummond
Recom Software
MS: 269-2
NASA Ames Research Center
Moffett Field, CA 94035-1000
Tel: (415) 604-4710
Fax: (415) 604-3594
Email: med@ptolemy.arc.nasa.gov

Bob Engelmores
Executive Director
Stanford University
Heuristic Programming Project
701 Welch Rd., Bldg. C (Rm. C-202)
Palo Alto, CA 94304
Tel: (415) 723-8444, (415) 723-5728
Fax: (415) 725-5850
Email: rse@hpp.stanford.edu

Peter Friedland
MS: 269-2
NASA Ames Research Center
Moffett Field, CA 94035-1000
Tel: (415) 604-4277
Fax: (415) 604-3594
Email: friedlan@ptolemy.arc.nasa.gov

Kathy Jurica
Lyndon B. Johnson Space Center
Mail Stop ER2
Houston, TX 77058
Tel: (713) 483-2042, (713) 483-0123
Fax: (713) 483-3204
Email: khealey@nasamail.nasa.gov

Astrid Heard
John F. Kennedy Space Center
Mail Stop DC-AST
Kennedy Space Center, FL 32899
Tel: (407) 867-2780
Fax: (407) 867-2050
Email: astrid_heard%pt_gate@kssib.ksc.nasa.gov

Mel Montemerlo
Manager, Artificial Intelligence Program
National Aeronautics & Space Administration,
Headquarters
Code CD
300 E Street, SW
Washington, DC 20546-3191
Tel: (202) 358-4664
Fax: (202) 358-3535
Email: m_montemerlo@rccola.hq.nasa.gov

Monte Zweben, President
Red Pepper Software Company
2929 Campus Drive
Suite 150
San Mateo, CA 94403
Tel: (415) 578-2644
Fax: (415) 341-8064
Email: ZWEBEN@PEPPER.COM

PRECEDING PAGE BLANK NOT FRAMED



REPRESENTATIVES OF EUROPEAN ORGANIZATIONS VISITED

Michel Maurette
CNES
18, Avenue Edourad Belin
31055 Toulouse Cedex
France
Tel: (33) 61 27 44 77
Fax: (33) 61 28 19 96

Jean-Michel Darroy
Matra Marconi Space France
31, Rue des Cosmonautes
ZI, du Palays
31077 Toulouse labege Cedex
France
Tel: (33) 61 39 67 35
Fax: (33) 62 24 77 44

Christian Parquet
Arianespace
Boulevard de l'Europe
BP 177
91006 Evry, France
Tel: (33) 1 60 87 64 30
Fax: (33) 1 60 87 64 03

Albrecht Kellner
MBB-ERNO, OT 11
Hunefeldstrasse 1-5
W-2800 Bremen 1
Germany
Tel: (49) 42 15 39 49 87
Fax: (49) 42 15 39 44 00

Herwig Laue
ESOC
Robert Bosche Strasse 5
61 Darmsdadt
Germany
Tel: (49) 61 51 90 23 66
Fax: (49) 61 51 90 34 02
Email: HLAUE%ESOC.BITNET@vm.gmd.de
Email: HLAUE@ESOC.ESA.DE

Roger Thompson
Vega Space Systems
Arden Grove
Harpenden, Herts AI5 4SJ
U.K.
Tel: (44) 58 24 61 678
Fax: (44) 58 24 61 77

Austin Tate
AI Applications Institute
80 South Bridge, Edinburgh EH1 1HN
U.K.
Tel: (44) 31 650 2732 (General AIAI)
Fax: (44) 31 650 6513
Email: bat@aiai.edinburgh.ac.uk

Mogens Nielsen
CRI
Bregnerodvej 144
DK-3460 Borkerod
Denmark
Tel: (45) 45822100
Fax: (45) 45821766

Francois Allard
ESTEC
Postbus 299
2200 AG Noordwijk
The Netherlands
Telex: 39098
Tel: (31) 1719-84856
Fax: (31) 1719-17400

Tim Grant
BSO/Origin
Bakenmonde 2
3430 BK Nieuwegein
The Netherlands
Tel: (31) 3402-88807
Fax: (31) 3402-60577



SITE REPORTS



MATRA MARCONI SPACE, FRANCE

Organization Overview

Matra Marconi Space is an international corporation with major facilities in England and France, which are partially or wholly owned subsidiaries throughout Europe and in the United States. Its major "products" are satellite Design, Development, Test and Engineering (DDT&E), satellite subsystems DDT&E and satellite operations. Matra operates both as a "prime contractor" (telecommunications, scientific and earth observation satellites), and as a "subcontractor" (Ariane electronics, onboard satellite instrumentation, software data processing and onboard satellite control systems). Major customers are CNES, ESA, and the French Defense Agency. The Toulouse facility, which we visited, is the primary Matra facility for artificial intelligence technology and application development. Most work in this area is done by the Software Technologies and Innovation Organization, a member of Matra's Technical Engineering Division. Investigations in artificial intelligence were initiated in 1984. In 1985 and 1986, the first contracts for AI applications were received from CNES and ESA. The first operational systems, ARIANEXPERT and PLAN-ERS, were delivered in 1990. The success of these initial systems marked the turning point with respect to the Matra internal management's support of AI technology. AI technology transfer is now strongly supported and encouraged by Matra's upper management and 30 percent of the funding for AI projects comes from IR&D.

Matra's development of products and prototypes utilizing AI technologies is distributed across multiple disciplines related to phases of spacecraft operations. These phases include monitoring and failure diagnosis, planning and scheduling, operations support and procedures management, and design. Research in the fields of knowledge acquisition, man/machine interfaces, task analysis and software development methodologies is also conducted. Major operational products and mature prototypes applicable to the phases of spacecraft operation, and research efforts are described below.

Specific Systems Discussed and Demonstrated

Monitoring and failure diagnosis-

ARIANEXPERT:

ARIANEXPERT is an operational knowledge-based system developed by Matra for Arianespace to assist in the labor intensive process of off-line post-flight analysis for Ariane launches. Throughout each Ariane launch, approximately 800 telemetry parameters are recorded at frequencies ranging from 1 to 400HZ. This telemetry data is analyzed post-flight in order to assess the performance of onboard systems and subsystems and to detect potential flight anomalies which could possibly prevent or delay the next Ariane launch. The first level of analysis, Level 0, is designed to detect such "launch stoppers" and must be performed in a matter of a few weeks after a launch. ARIANEXPERT assists human analysts in Level 0 analysis for one of twelve launch "domains", the propelled piloting phase. (This domain's parameters represent a subset of the total set of telemetry parameters.) ARIANEXPERT performs its analysis by selecting and executing suites of analysis

procedures (organized in decision trees) which compare flight data values for sets of parameters against expected or reference data values. (Reference data values are recorded from previous launches and are indexed via specific events such as "Phase 1 separation." Reference data is updated after each flight to include that flight.) ARIANEXPERT combines AI techniques, signal processing techniques, advanced graphical and statistical capabilities, and a highly sophisticated interactive human interface to provide a powerful tool to assist the human analysts. It also automatically generates formal post-flight-reports. The use of ARIANEXPERT has reduced the effort required to perform post-flight analysis for one domain from four man/days to one man/day. It was developed in KEE by a staff of 1.5 people over two years, and runs on a SUN workstation. It includes capabilities for explaining and justifying its reasoning procedures.

ARIANEXPERT was a "big win" for Matra and is considered by the technologists to represent a turning point with respect to the level of management support provided for further technology development and application projects.

DIAMS2:

DIAMS2 is a near real-time, off-line fault diagnosis and recovery expert system developed by Matra for CNES to be used in the Spacecraft Control Center for the Telecom 2 satellite. (The system does not do fault detection.) DIAMS2 will be operational in 1993. DIAMS2 is also designed to provide a training capability and incremental knowledge refinement capability. It incorporates two "models" of the satellite subsystems for fault diagnosis, a behavioral model and a functional model. Decision trees are used as the representation structure for the behavioral model. The behavioral model provides the global method for diagnosis and is used to focus the diagnosis to particular functions or components. The functional model uses a quantitative object-oriented model representation and "Kate-like" diagnosis algorithms to provide the final isolation of the detected fault. DIAMS2 was developed over a period of three years.

X-ANALYST:

X-ANALYST is a generic tool for data analysis and intelligent task chaining which is derived from and based on techniques employed in the development of the ARIANEXPERT system. The tool was developed internally by Matra and has been used to develop off-line data analysis applications for the Telecom 2 and HISPASAT satellites and for a new version of ARIANEXPERT built for Ariane 4. These applications will be operational in the near future and provide all the capabilities available in the first ARIANEXPERT system. Applications built with the X-ANALYST tool are designed to allow users to add or modify knowledge by themselves.

Planning and scheduling--

PLAN -ERS:

Plan-ERS is a scheduling tool designed to generate schedules for use of instrument activity onboard the Earth Resources Satellites (ERS). ERS-1 is currently in orbit and ERS-2 will be launched in the near future. Plan-ERS is also designed to be used as a mission analysis tool by the satellites' schedulers. Instrument activity requests are input by the user. Activity attributes include

instrument class start/stop times, resource utilization, satellite attitude and environmental constraints, and orbit specification (changeable only by the user). Plan-ERS attempts to schedule as many instrument activity requests as possible, while optimizing both onboard and ground resource utilization. Plan-ERS utilizes object-oriented modeling of the spacecraft, resource modeling (constraints and utilization), and rule-based modeling of operational scheduling knowledge, (e.g., scheduling/rescheduling knowledge). It provides the following capabilities through an interactive user interface: schedule editing, schedule repair and patching, schedule decomposition and recombination, and “zooming”. Schedule patching is done automatically. Schedule repair is accomplished interactively with the user. Plan-ERS has been operationally used by ESA for the ERS-1 mission.

Optimum -AIV (Assembly, Integration and Verification):

Optimum-AIV is a project management tool for scheduling spacecraft assembly, integration and verification processes. It was developed by Matra under contract to ESTEC and was delivered in April of 1992. It has been adopted by the ARIANE 4 production team for equipment bay AIV planning scheduling. The scheduling problems addressed, and the technical approach utilized, by Optimum-AIV are similar to the problems addressed and approaches utilized by the Gerry scheduling tool. Optimum-AIV provides an extensive, highly interactive user interface. Inputs to Optimum-AIV include activity descriptions, resource descriptions, resource and project calendars, and resource and conflict resolution strategies. Types of constraints which may be defined include global, state, resource and temporal. Optimum-AIV verifies the logical consistency of the plan input by the user before generating a schedule. During scheduling, Optimum-AIV collects conflicts which cannot be automatically resolved and supports the user in solving them.

PADRE:

MMS has been recently awarded, by the French Ministry of research, a new project called “PADRE”. The goal of this project is “real-time” planning and resource allocation. Such systems are needed when planning must be done in parallel with the on-going process, and under strong time constraints. Typical applications of such techniques in the space domain are:

- scheduling/resource allocation for a telecommunication system (e.g., TDRSS or DRS);
- planning and re-planning of equipment performed on-board a space station.

Operations support and procedures management–

Matra has developed, or is developing, multiple procedure generation and procedures execution or management systems and tools. Included in the procedure generation set are the Procedures Operation Management (POM) tool, the Expert Operator’s Associate (EOA) prototype application and the PREWISE application. Included in the procedures execution or management set are the EOA and the Crew Support System (CSS) prototype system.

The general approach utilized in the procedures generation tools/applications is to provide the operations engineer with a library of procedure steps for constructing procedures. This promotes reusability, consistency, and maintainability. Procedure representations allow the definition of

pre-execution steps, execution constraints, actions, post-execution checks, post-execution constraints, and branching actions for error conditions.

In some cases procedures are defined by the user in a "natural language" format and the tool parses the input to produce the procedures' internal representation. In other cases, the user defines procedures via constrained editing. The procedure generation tools/applications also provide a procedure verification environment.

The general approach to procedures management or execution tools/applications is to provide an intelligent assistant to the human operator (ground-based or onboard) responsible for procedures execution. These tools/applications are being designed to support procedure search and browsing and autonomous procedure selection. Additional design goals include autonomous cooperation with diagnostic and planning tools/applications during procedure execution and a reactive architecture which enables "real-time" responses to incoming alarms and user inputs. The interpreter components of these tools are designed to allow procedure interruption/resumption and the synchronization of parallel procedures. Matra is developing two languages to support procedure execution, TL1 and ELISA.

POM:

POM is a hypermedia-based procedures generation tool which is in operational use both by Matra (Hispasat, SOHO, ...) and by CNES (TC2). It is being used to generate operation procedures for the TC2, HISPASAT, and SOHO satellites for which Matra is the prime contractor. There is no parsing provided by the POM. The user enters procedures via a highly constrained form close to that of the internal representation. POM includes a state simulator module for procedure verification.

PROCSAT:

PROCSAT is a tool which is very similar to PREVISE, but which addresses SPOT 4 earth observation satellite ground operations procedures. This contract is sponsored by CNES, and the system will begin to be operationally used by CNES operations in June 1993.

PREVISE:

PREVISE is a procedures generation and verification application currently under development. It is the first of 12 AI applications planned by the ESTEC's Expert System Demonstration Project. (The goal of ESTEC's Project is to demonstrate the benefits of AI applications. The domain of all 12 applications is "in-orbit infrastructure.") The PREVISE application is focused on the generation of crew procedures for IVA, EVA, Rendezvous and HERA robotics operations to be used onboard the HERMES. The procedures generation and verification architecture includes a procedures editor and a knowledge editor which provide inputs to a procedures compiler that produces the internal representation used by the procedures formatter and procedures executor. (The procedures editor will provide both a syntax driven approach and a natural language approach, which will allow the user to expand the language and syntax.) A procedures checker, via qualitative procedures simulation, utilizes verification knowledge (methods, constraints and resources) and state descriptions from a State editor/loader to perform local, temporal, quality and logical checks on the procedure internal

representation generated by the procedures compiler. Error messages and warnings are generated when "checks" are not satisfactorily completed. PREWISE is being developed in PROKAPPA and the initial prototype will be available in May 1993. Although focused on HERMES procedures, PREWISE is being designed to be adaptable to other domains.

EOA:

The EOA is both a procedures-generation and a procedures execution application. It is an early KEE-based prototype which was developed over five years by Matra and CRI under funding from ESA for ESOC. The objective of the EOA project was to determine whether knowledge-based techniques could improve the reliability and efficiency of complex procedure-based spacecraft control tasks. The telecommunications satellite MARECS was chosen as the application. Procedures are defined by the user via a menu-driven language and are editable. During procedure execution, procedures can be interrupted and resumed. EOA saves the current state at interruption and restores that state upon resumption. EOA can execute procedures in parallel. It also responds to incoming alarms and user input by selecting appropriate procedures. If procedures are aborted, EOA "cleans up" by satisfying predefined abort constraints (like PRS). The EOA prototype is under evaluation at ESOC and has demonstrated the feasibility of a knowledge-based approach. EOA has been successfully experimentally demonstrated at ESOC on the MARECS satellite, during eclipse operations.

Crew support system (CSS):

In 1988 Matra, under CNES sponsorship, initiated a R&T project to develop a procedures-execution application prototype to study the application of AI to the support of HERMES astronauts during rendezvous operations. Included in the objectives were the analysis and test of man-machine interaction concepts. The RVD EXPERT prototype, and an associated simulator, were completed in 1990 and demonstrated automated procedures execution and mission replanning, or procedures adaptation, based on automated anomaly diagnosis and vehicle configuration management. As a follow-on in 1991, ESA and CNES initiated the CSS, a feasibility study for implementation of a system to assist in procedures execution (as opposed to the automatic procedures execution demonstrated by RVD EXPERT). The study is currently focusing on a ground-based procedures execution assistant which could be applicable to multiple procedure "types", including rendezvous and payload operations. Efforts are being made to ensure compatibility between internal procedure representation formats produced by PREWISE and utilized by CSS.

Training-

ITS:

In the fall of 1992, MMS jointly with CISE (Italy) were awarded the second of the 12 AI applications planned by the ESTEC's Expert Systems Demonstration Project (of which PREWISE is the first).

The objective of ITS (Intelligent Tutoring System) is to develop software systems which can support self-training of personnel. Such tools can be used for different categories of space personnel: astronauts, satellite operators, mission controllers, integration engineers, In this context,

Artificial Intelligence Techniques are used to adapt the pedagogical approach (pedagogical methods, level of dialogue, ...) to the individual preferences and skills shown by the trainee.

Design-

Matra's advanced computer-aided design applications utilize AI techniques (e.g., search, geometrical reasoning, reasoning by analogy, simulated annealing) and OR techniques (multi-variable optimization). The software engineering approach is object oriented, (LISP or KEE). Three applications have been developed which solve difficult, but well-circumscribed design problems. Details on technical approaches were not available due to the confidential nature of the applications. All applications require specially developed model representations of system designs (i.e., do not use CAD databases).

SWITCHWORKS:

SWITCHWORKS is an operational system, deployed in 1991, which analyzes the reliability and efficiency of telecommunications satellite payload designs. It is considered by Matra and its customer to be an enabling application and has convinced the customer of the essential nature of AI technology in future applications. Reliability is assessed by analyzing all redundancy paths in the design to ensure that the failure of traveling wave tubes does not prevent any set of N selected channels among M to be routed to the remaining functional tubes. If the design is not reliable, an analysis is performed of all possible component failures which may result in the inability to restore full payload functionality in order to determine redesign requirements.

PAYLOAD EXPERT:

PAYLOAD EXPERT is an operational system which analyzes the reliability and efficiency of a multi-satellite telecommunications system. A global evaluation of the design space (all possible bindings of channels to satellites) is performed to ensure that the failure of satellites does not compromise the telecommunications mission.

ANTENNA EXPERT:

ANTENNA EXPERT is an antenna design optimization application which assists a designer in making design choices by simulating the system design to ensure satisfaction of constraints, such as structural and thermal, and by suggesting alternative techniques when constraints are violated.

Action for research and applications in man-system interactions-

ARAMIIHS:

ARAMIIHS is a multi-organizational European research effort which was initiated in 1988. Matra is the most significant "player" contributing 40 percent of the funding for the project. The scope of the research effort is broad and comprehensive, covering the production of documentation, knowledge acquisition, human factors, man-machine interaction, linguistic engineering, and technology-based training. AI efforts are focused on knowledge acquisition, design of

knowledge-based systems which are interactive and cooperative (with human operators), and validation of knowledge-based systems. Progress has been made in the development of a tool, MACAO, which assists a knowledge engineer in the acquisition and representation of expert knowledge and reasoning during knowledge engineering sessions. The final internal representation of this knowledge and reasoning is in the form of graphs (currently built manually). A fault diagnosis assistance application has been developed for ARIANE 4 equipment bay pre-launch tests, based on knowledge representation resulting from the use of the MACAO tool. This diagnostic system is PROKAPPA based and utilizes case-based reasoning. A prototype has been completed, and the first operational system will be deployed in 1993.

ARIANESPACE

Organization Overview

Our team visited Arianespace in Évry, France, on October 7, 1992. Our host was M. Christian Parquet, an Engineer in the Flight Evaluation Department. Also present at the meeting were: M. Jean-Pierre Dulout, Dept. Manager for Mission Analysis in the Industrial Directorate; M. Michel Bartolomey, Division Chief of Launch Operations and Chief of Operations "Ensemble de Lancement"; Mlle. Isabelle Rongier, a flight control expert from CNES in Évry.

General Observations

During the previous two days at Matra Marconi Space we heard a description and witnessed a demonstration of a system developed at MMS called Arianexpert (q.v. site visit report for MMS). At Arianespace we were able to get the user's perspective on this system.

Mr. Parquet led the presentation. His viewgraphs were so thorough and well-written that this report will be largely a transcription of that material.

Arianespace is required to perform a post-flight analysis (PFA) of every flight of their Ariane launcher. The purpose of this analysis is to detect any anomaly that might have occurred so that it can be corrected before the next launch. Using telemetry data that is transmitted, processed and stored at Toulouse during the week following a launch, the PFA is performed at Évry by some 12 groups of engineers and experts (approximately five per group), from the various companies involved in ARIANE's program. This analysis, which is called a "Level 0" analysis, is performed over a three-day period. The analyses of 12 different technical domains are then summarized and presented to a Director's Committee in order to authorize the next launch or to initiate corrective actions. Later on, a deeper analysis is made (called a Level 1 PFA) for each flight. The Level 1 analysis is a six-month task.

Several factors motivated the development of an expert system to assist in post-flight analysis:

- Level 0 PFA must deal with 500 to 800 parameters, some of which are transmitted at a rate as high as 400 values per second. (The data is assumed to be free of noise, and data loss is approximately 1% for all the technical domains. Errors in the data are expected during particular periods of flight. The interpretation of data quality is made by humans, not machines.)
- Arianespace may launch their rockets at a rate as high as one every 22 days, so there is a very short time to conduct an exhaustive analysis of each technical domain.
- Because the availability of experts varies, and engineers normally change jobs occasionally, the quality of analyses from one flight to another can vary.
- The PFA is largely a manual process, making it difficult to compare one analysis with another.

For these reasons, Arianespace asked Matra Marconi Space to develop a high performance tool, Arianexpert, to assist operators in making a more systematic analysis. It was agreed that the first system would assist in the analysis of one of the most important of the 12 technical domains; namely "Propelled Phase Piloting," which covers the period of the launch from lift-off through the separation of the last stage. A prototype was delivered at the end of 1990, tested on five previously analyzed flights, and then made operational for flight V45, which was launched on August 14, 1991. A full week was spent with the experts to validate the domain knowledge and to prove to them that the system was sufficiently flexible that they could modify the knowledge base by themselves.

Specific System Discussed and Demonstrated

Arianexpert-

What does Arianexpert do? It helps the operator to select the necessary data for the domain. It presents the data in a form that is natural to the operator and in a logical order. It provides a methodology for each analysis and systematizes the PFA for each domain, ensuring consistency and completeness. It explains the purpose of any analysis should the operator require such help. It can also process data, using a set of user-provided functions. It compares the results of the analysis with the specifications, or poses questions to the operator in order to qualify the analysis. It generates an analysis overview for the director, and a report of the flight results, both of which can be easily edited and updated.

The system has now reached a steady state. Since 1991 it is the only method of analysis used for the Propelled Phase Piloting domain, having completely replaced the traditional hand analysis. The system assists any trained operator in going through 35 analysis steps (unusually large for one domain), each having five to seven elementary units, in one day, as compared to two days by the traditional method. The users especially appreciate the automatic report facility, and the additional time available to them to make a deeper analysis if an anomaly arises. When anomalies arise, the experts are responsible for identifying them, and Arianexpert facilitates that identification. At least four engineers have become familiar with the system. They can modify the knowledge base themselves and have learned to go quickly through "routine units" and spend more time on the non-routine events.

The progressive introduction of Arianexpert was a success – users now trust the system's analysis and are satisfied with its functionality.

What's ahead? Some obvious improvements will be implemented, such as new processing functions or automated backup procedures. The main next step is to add additional PFA domains, which will require the time and willingness of the experts and the administrator (non-trivial requirements!). The choice of domains to add will depend to a large extent on the willingness of the users to adopt this new approach to analysis. The PFA for the inertial guidance domain, which heretofore has not been covered in the Level 0 analysis, is expected to be implemented by the end of this year. Further down the line they would like to implement a "technical memory" which will give a textual account of previous anomalies.

Arianexpert runs on a Sun 4 Sparcstation. The cost of the hardware was about \$100K. The cost of the software (so far) was \$300K to \$400K.

Our hosts also talked to us briefly about another application in the area of planning and scheduling. Using the tool PLANOPS, from Thomson CSF and modified to Arianespace specifications, Arianespace plans the assembly of each launcher. The system makes use of planning heuristics, such as “postpone the assembly as much as possible toward the launch date.” The scheduling is done daily, in real time.

General Comments

- Arianespace does not exploit the diagnostic capability of Arianexpert that Matra described to us. Users found that the fault tree representation was not very useful to them because known faults from past flights are fixed and don't recur. Moreover, the fault tree itself is difficult to generate. Thus, use of the system for diagnosis is not foreseen.
- Although use of Arianexpert reduces the analysis task from four man-days to one man-day, the primary objective was to increase the quality, consistency and thoroughness of the analysis.
- Despite the success of Arianexpert, upper management is not convinced that the introduction of AI is justified, mainly because it is not ARIANESPACE's job to manage software development projects. However, the ARIANESPACE Industrial Division is convinced that they have to initiate this kind of improvement.

MBB DEUTSCHE AEROSPACE

Organization Overview

MBB is a large corporation recently acquired by the German Daimler-Benz Corporation and integrated into Deutsche Aerospace. The facility visited is now called the ERNO Raumfahrttechnik and was previously called MBB-ERNO. The overall group performing the Artificial Intelligence activities is the System Development Organisation lead by Mr. Seibl. The host of the visit was Dr. Albrecht Kellner, head of Mission and Automation Technology organisation, under System Development. Dr. Kellner previously ran the Basic Research department of the Informatics organisation (also under System Development). This department is the leader in Artificial Intelligence research and applications, and Dr. Kellner initiated these activities. Informatics is managed by Mr. Fesche, and Mr. Norbert Schielow now leads the Basic Research department.

General Observations

One interesting note about the System Development organisation is that the Informatics department and the Mission and Automations Technology department are at the same level as the Operations department. This structure facilitates technology transfer and research relevance because the groups are co-located and appear to collaborate without any resistance. This should be contrasted with the relative resistance to introducing new technology that operational groups at NASA and the aerospace contractors exhibit.

Specific Systems Discussed and Demonstrated

The ERNO projects that were presented were all practically motivated. There was no equivalent of what we typically label "basic research." Every project had a mission bias—some for current missions and some for longer term missions still in the design phases. As a result and similar to the NASA AI program, ERNO specializes in fault detection, isolation, and recovery (FDIR) and mission planning and scheduling. The main projects are producing the CONNEX, SIMMEX, and MARS/NEPTUNE tools. Additionally, there are a number of application projects where these tools are being used in a variety of mission contexts.

CONNEX—

CONNEX: An experiential FDIR tool CONNEX was developed to remedy shortfalls that were observed with decision-tree representations of symptom-fault associations. The ERNO team feels that greedy decision-making is inappropriate for diagnosis because of the noise embodied in real-world applications. They view greedy systems as those systems that test symptoms, one by one, rejecting hypotheses when the expected symptoms are not observed. They observed that if a mistake is made near the root of the tree, then valid hypotheses are rejected. They argue that exact symptom fault pattern matching is inappropriate for real-world applications. They solve this problem in CONNEX (which stands for Connection Matrix Expert System) by forming a matrix of

114

symptom-fault entries and find the most likely fault by using what they call a "proximity measure." The proximity measure is a weighted ratio of observed "exceptions" (i.e., symptoms) to expected exceptions. The system sorts the candidate anomalies by the proximity measures and presents them to the operator on a graphical user interface. The interface is also used to create and edit connection matrices.

The ERNO team has widely applied CONNEX. The most significant application of CONNEX is the Columbus (SSF module) FDIR prototype. Since ERNO is the lead contractor for Columbus, they believe CONNEX will be operationally deployed. Below is a list of applications that use CONNEX.

COMPASS-

COMPASS: Computer-based Payload Operations Support System/**TIKON:** Technology for Intelligent Control/**SACV:** Satellite Autonomy Concept Validation

The system is developed in LISP and runs on Symbolics machines. A version for the SUN will be available soon.

COMPASS: Computer-based Payload Operations Support System

COMPASS is an application of CONNEX for the D2 Space Shuttle Spacelab mission. COMPASS will diagnose the telemetry from the HOLOP payload aboard the D2 mission. It will be a Macintosh application that will most likely be deployed at GSOC. COMPASS is an ESTEC project led by DLR.

SIMEX-

SIMEX: A simulation-based FDIR tool. The goal of the SIMEX project is to augment traditional and CONNEX-based FDIR systems with the use of simulations. The objectives are similar to the motivations of the model-based diagnosis community. The main insight is that a simulation of a device or system can be compared to the actual sensor observations of the system in operation, and when the two differ, a fault is detected. Then the model of the simulation is used to generate candidates for isolating the faulty component. One important design principle for SIMEX was that the simulator and modeling language chosen had to be one already accepted and used by designers. Consequently, design models could be used for multiple purposes. They chose the Core Simulation Software (CSS) that includes the Model Development Environment (MDE). These systems were already in use for Columbus. SIMEX does not use any well-known model-based diagnosis algorithms. Instead, when a discrepancy between the simulated behavior and the observed behavior is detected, the system launches a generate and test search for the faulty candidate. SIMEX makes a single-fault assumption. The model is used to prune candidates from the search process and also imposes a search bias. Candidates are explored in a breadth-first manner back from the observed discrepancies according to the connections in the model. The system is partially complete. A prototype of the hypothesis generator and tester has been demonstrated and the integration to the MDE has begun. The system is implemented in LISP using the CLOS object language and the CLIM GUI language.

MARS/NEPTUNE-

MARS/NEPTUNE: A heuristic scheduling system. The Mission Activities and Resource Scheduler (MARS) is a mature scheduling tool. The system is a complete architecture with an expressive task definition language, a sophisticated graphical user interface, a heuristic backtracking search mechanism, and a rule-based search control mechanism. The resource modeling in MARS allows for both reusable and consumable resources. No sophisticated reasoning is performed to establish when consumable resources should be replenished. Additionally, TRI-STATE constraints are supported which allow for mutual exclusion and is first step towards state constraints. MARS provides hardcopy reports in a variety of formats including Wordperfect and PostScript and also includes a database interface. The system also supports task hierarchies. MARS supports all of the Allen temporal relations and uses a temporal constraint network algorithm to ensure the satisfaction of these constraints. Due dates, duration intervals, and min-max delays between temporally related tasks are also easily modeled with MARS. MARS has been specifically designed to handle very large data sets.

The scheduling algorithm in MARS is straightforward. The system fires a set of task selection rules to choose the next task to schedule. Then the earliest time (that has not yet been tried) when all temporal constraints, due dates, and resource constraints are satisfied is then assigned. If the schedule can not be extended, a set of backtracking rules determines what assignments should be rejected and retried.

MARS also supports rescheduling in that a set of rescheduling rules decides what assignments should be rejected and retried in case of an external scheduling change.

MARS is implemented in LISP on Symbolics machines and is currently being ported to the SUN using CLOS and CLIM. Neptune is the next version of MARS that concentrates on distributed scheduling. Neptune allows multiple MARS processes to communicate with each other in order to break down a large-scale scheduling problem into manageable pieces.

MARS has been applied to a number of problems as studies. These include:

Analysis of HERMES operations	ESA	completed
ROSSA crew productivity analysis	ESA	completed
EURECA payload operations	MBB	completed
MARS/MIPS comparison	ESA	completed
EURECA maneuver planning	EURECA	completed
Satellite Autonomy Concept Validation	ESA	on-going
HERMES timeline engineering	ESA	completed
Crew Workstation Architecture	ESA	on-going
BIOLAB phase A	ESA	completed
SACV	ESA	completed
TIKON	ESOC	completed
HFCC Design Study	DARA	on-going
MARS and MDA	MBB	on-going
ARIANNE Production Planning	MBB	on-going

MARS for D2
MARS for Crew Training

MBB
MBB

on-going
proposed

TIKON:

Technology for Intelligent Control and SCVS: Integrating FDIR and scheduling for autonomous satellites

TIKON and SACV are architectures for full system autonomy. They integrate the CONNEX diagnosis systems and the MARS scheduling systems. A monitoring component observes the telemetry stream and uses CONNEX to diagnose failures. Upon a failure, a recovery procedure is added to the mission plan which MARS then deconflicts. These projects take a first step towards full autonomy. The TIKON project focuses on the ROSAT satellite. The SACV project actually demonstrated significant autonomy on the EURECA simulator for an extended period of time.

EUROPEAN SPACE OPERATIONS CENTER (ESOC)

Organization Overview

The European Space Operations Center (ESOC) is located in Darmstadt, Germany. Over the past 25 years ESOC has been responsible for the operation of 21 ESA/ESRO spacecraft missions as well as numerous satellites for more than 20 agencies. Spacecraft operations cover a wide variety of tasks carried out both before and after the launch of a spacecraft. These tasks include mission requirements analysis, orbit determination, ground segment preparation, tracking and control of spacecraft on orbit, and reception, processing and distribution of both spacecraft and payload data.

In support of its mission operations task, ESOC initiates studies of a variety of technologies, including artificial intelligence. The goal is to determine the potential impact of evolving technologies on future mission requirements and facilities investments. Studies in the area of artificial intelligence assess the applicability of AI with respect to current and future mission operations challenges for ESOC in the areas of unmanned experimental and scientific missions, unmanned operational missions, and manned in-orbit infrastructure. The challenges include spacecraft and mission complexity, life-time, availability, multiple spacecraft, and safety issues (see table 1).

Table 1: ESOC operational challenges for AI

	Complexity	Life Time	Availability	Mult. S/C	Safety
Unmanned Experimental & Scientific Missions	X				
Unmanned Operational Missions		X	X	X	
Manned In-Orbit Infrastructure	X	X	X		X

Specific demands on mission control systems include more payload modes and higher payload duty cycles, shorter mission planning cycles, combined payload operations at the limit of the resource envelope, and shorter turn-round times for payload data. Spacecraft complexity is increasing in such areas as subsystem redundancy, automatic and autonomous on-board functions, on-board software, and on-board data handling systems.

Knowledge-based systems are the technology of primary interest. The overall development philosophy is to conduct individual studies and develop stand-alone prototypes. Original R&D is generally not conducted in-house at ESOC, but is contracted out to specialists in multiple countries. Occasionally, studies and small prototypes are developed in-house by students or support contractor personnel. Selected prototypes undergo additional development to become an integrated mission-independent prototype which may be evaluated in an operational setting such as a control room. Ultimately, it is hoped that advanced prototypes may mature into operational systems or affect the

design of future operational systems. No operational systems using artificial intelligence are currently being developed, nor currently in use at ESOC.

Specific Systems Discussed and Demonstrated

Several applications of artificial intelligence at various stages of maturity are under study under the cognizance of ESOC or at ESOC itself. The area of automated tools to assist in operations procedures generation and procedures execution is a prime focus of attention at ESOC. Efforts underway include ESSOPE, ATOS, FOPSET, and EOA.

Expert Systems—

The ESSOPE system (Expert System for Spacecraft Operations Planning and Execution) is being developed by FIAR Space Division and Intecs Sistemi (Italy) for ESOC. The goal of this prototype system is to schedule the operations and maintenance of OLYMPUS satellite payloads. The system monitors schedule execution, and then initiates contingency procedures if anomalies are detected. The prototype was implemented in OPS83 on a Sun workstation, and was interfaced to the OLYMPUS simulator at ESOC for testing. ESSOPE uses a state-transition diagram to represent spacecraft states. A planner uses this diagram and knowledge about constraints to plan state transitions by given deadlines. The system is designed to issue commands directly to the spacecraft. However, an operator may elect to authorize each command if desired. Future plans for development of an operational system are uncertain since OLYMPUS suffered a serious malfunction which has invalidated the ESSOPE knowledge base and the simulator. In the interim, results are being used in the ATOS project.

ATOS:

The ATOS project (Advanced Technology Operations System) is seeking to develop a framework for application of AI and expert systems in mission operations. Specifically, the goal is to develop a system architecture which is compatible with external interfaces and the conventional spacecraft control system. The architecture must provide tools for storing and managing knowledge which is used, generated, or exchanged by multiple intelligent application modules. Common internal and human-computer interfaces are also of interest to the project. ATOS application modules which are planned include mission preparation, mission planning, operations execution (possible use of ESSOPE technology), and adaptive training. The ATOS concept has a requirement for a common repository of all system and mission knowledge. The repository, called the "Mission Information Base" consists of a shared database, user tools, internal management functions, and distributed access facilities. Portions of the MIB are being written in Ontolingua, and there is a desire to fold in the PACT architecture. The project is still in the first phase.

FOPSET:

The FOPSET project (Expert Tools for Flight Operation Plan Production) is prototyping a set of "basic functions" for editing and formatting the FOP (Flight Operation Plan), including automatic fetching of relevant data from databases in phase 1, and advanced functions for procedure validation

in phase 2. The current phase 1 prototype uses Interleaf TPS for FOP document preparation. In phase 2, the ProKappa tool will be used for procedure validation. The output of FOPSET is intended to be in a form which could be executed (i.e., spacecraft commands).

EOA:

The EOA system (Expert Operator's Associate), developed jointly by Matra Marconi Space, CRI, and ESOC, is a prototype expert system which assists an operator in controlling satellites. The prototype has been previously demonstrated using an existing real-time simulation model of the MARECS-B2 spacecraft, and is currently installed and awaiting operational test with MARECS during the eclipses season beginning 1993. The goal is to make the EOA technology operational for the next generation Spacecraft Control and Operations System (SCOS II) in about 1995. EOA was the first system for telemetry monitoring to be developed for ESOC. MARECS operations personnel stated that much more experience and testing is needed before they will be comfortable with the system.

Meteosat:

The Meteosat workstation prototype is an expert system developed by VEGA (U.K.) for ESOC. The "near operational" prototype is currently installed in the Meteosat operations center and is used in parallel with existing operations systems on an occasional basis. The system monitors approximately 350 channels of engineering data from each of three satellites, including one on loan to NOAA. The system provides capabilities for dynamic alarm limits, animated block diagram schematics of satellite systems (using DataViews), concurrent procedure execution, and rule-based diagnosis and command receipt validation (using ProKappa). Operations personnel are very enthusiastic about the system. However, there are no current plans to continue system development since the Meteosat operations center will be closed when operations are turned over to a new contractor and operations center in 1995.

Radiometry Expert System:

The Radiometry Expert System is also a near-operational prototype installed in the Meteosat operations center. The system monitors the state of two counters and multiple switches on the Meteosat satellites using real-time or playback data. The system diagnoses problems in the radiometer and recommends recovery procedures. It was noted that training on radiometer operations for recovering from failures is the most significant portion of operator training since a mistake could cause a loss of the mission. Many such failures, up to two or three per day, have occurred on pre-operational Meteosat satellites. The traditional approach would require an operator to analyze about 60 pages of printout from a mainframe computer. With the assistance of the Radiometry Expert System, the operator is able to recover within 1/2 hour of a failure episode. Although still a prototype, the system is constantly in operation, and it was reported that it has significantly increased the performance of the Meteosat mission. The first prototype system was developed in-house by graduate operations trainees, and the current system was developed by VEGA. Like the Meteosat workstation prototype, there are no plans for further development.

AMFESYS:

The AMFESYS (Automatic Mirror Furnace Expert System) is a prototype developed to monitor the status of the AMF instrument on-board the EURECA satellite. The goal of the system is to accurately model the state of the instrument in real time despite short periods of spacecraft access (visibility) by the controlling ground station. The prototype, which is demonstrated in a laboratory using a spacecraft simulator, uses model-based reasoning techniques to simulate the instrument. When access to real-time telemetry is obtained, the model is automatically synchronized with the AMF instrument. Differences between the model and the actual state of the instrument are analyzed, although deep model-based diagnosis is currently not performed. The technology was specifically cited as potentially very valuable in helping to improve prediction of AMF-state and reorientation of satellite operators after long periods of no contact with the instrument. The prototype was developed using the ProKappa tool.

GSDAS:

The GSDAS study is developing a diagnostic and advisory system prototype for Telemetry and Tele-Communications (TTC) ground stations. The eventual goal of the system is to improve ground station availability by using expert systems for malfunction detection and recovery advice. A prototype called DAS-ON is nearing completion in the first phase of this study. DAS-ON performs on-line diagnostics using shallow heuristic knowledge and model-based knowledge in a two-level architecture developed using the CLIPS 5.0 tool. A preliminary conclusion is that the approach is appropriate for system-level (e.g., TTC chain) diagnosis but not for individual subsystems. Subsystem diagnostics were found to be too slow using the CLIPS 5.0 system so the decision was made to have subsystems provide their own diagnostics. The real-time requirement at the subsystem level is to monitor 20,000 parameters every two seconds. The development of the DAS-OFF prototype will be initiated soon, with the goal of providing off-line diagnostics and recovery advice. The prototype will accept DAS-OFF diagnostics as input and will provide recommendations for TTC reconfiguration based on availability of redundant subsystems and scheduled mission support.

Other expert system projects in the area of mission operations which are currently being investigated include the Broadband System Configuration Management System, which has the goal of detecting performance and accessibility problems in the ESOC broadband LAN. A prototype is currently being implemented on a Sun 3 using Prolog and C. The Expert System for Network Management prototype has the goal of providing the operator of an integrated network management system with alarm correlation facilities. This prototype is currently being implemented on a Sun 4 using Prolog and C++.

In the area of automated tools for data analysis, several systems were discussed. The HERMES/COLUMBUS Mission Analysis Front-End, also called the Mission Analysis Assistant, was developed. The Mission Analysis Assistant prototype functions as an intelligent interface between a mission analyst and a library of mathematical software used for such tasks as trajectory optimization, orbit computation, thrust calculations, rendezvous and docking, re-entry, and landing simulation. The prototype was demonstrated on a specific HERMES/COLUMBUS rendezvous and docking scenario in 1991 and was described as a complete success. The system was developed using ART by Telefonica Sistemas and GMV of Madrid, Spain. However, the system will never be used

because both HERMES and the COLUMBUS Free-Flyer have been cancelled. The prototype is currently broken due to software incompatibilities, and there are no plans to maintain or upgrade it for other applications. The ART tool was cited as being very difficult to use for development and maintenance of applications.

Another prototype, called ESIOD, for computer assistance to flight control operations in the area of initial orbit determination, was discussed. The goal of the prototype was to proceed in the case of poor orbit determinations. The prototype was developed for application to GTO and NSO orbits only. The system was written in Prolog, and demonstrated in April 1991. The prototype failed to perform acceptably in the area of diagnosis and recovery recommendation. This is attributed to the problems in knowledge acquisition and knowledge representation since the knowledge engineers did not have an adequate knowledge of the flight dynamics domain. The human-computer interface to the system proved valuable, and further developments will focussed on that area alone. No continued use of AI in this area is contemplated.

The area of neural networks is just beginning to be studied. Several efforts are currently underway. A university is investigating the use of neural networks for solar flare prediction. Neural networks are also being studied for use in extracting wind vectors from ERS scatterometer data and for extraction of wind vectors and classification of cloud types in Meteosat data. No prototypes have been developed to date.

Other Discussion Topics

The Plan-ERS system, developed by Matra Marconi and others for ESOC (see elsewhere in this report), was mentioned briefly by ERS mission planning personnel. They stated that Plan-ERS never produced a spacecraft schedule, but in tests did perform limited resource allocation for the ERS payload only. It was reported that "plan-ERS does not work" and it is not used in ERS mission planning at ESOC for any function.

The need for validation of knowledge-based systems was expressed by the ESOC participants, who also said this was an extremely important concern with any system that is capable of issuing commands to a spacecraft. Discussion centered around the idea that expert systems can "get into situations" where the software has not been exhaustively tested. ESOC personnel noted that they are watching a current ESTEC study in this area with interest. Maintenance of knowledge bases is also an area of concern. Finally, ESOC personnel noted a need to keep up the operator's level of interest and proficiency while depending on automated tools as an issue for the future.

VEGA SPACE SYSTEMS

Organization Overview

Vega Space Systems is a spacecraft operations consulting company founded in 1978. They currently employ about 140 people in three locations (main office in Harpenden, England, a field office in Germany near ESOC, and other field personnel in the U.K.) They went public in the U.K. a few months ago; 1991 revenues were 6.2M pounds. Major customers (both directly or indirectly as subcontractors) are ESA (both ESOC and ESTEC), EUMARSAT, and EUMETSAT. They provide support for flight operations plans and procedures, launch systems, and post-flight analysis. Their work is about 50% software and about 50% systems engineering. Only a very small percent (10-15%) of their effort is military.

Vega's interest in AI comes from its role in total systems solutions; they do no true AI research, although they are heavily involved with tool development. Essentially all of their work is in unmanned satellite operations and support. Because of Britain's miniscule contribution to the ESA manned program (now only Columbus module for SSF since Hermes and MTF were cancelled during the end of September 1992), they are precluded from participating on Columbus.

General Observations

Comments on specific systems we were shown will follow, but first a few general points about their AI work.

1. Vega appears to view itself as one of three significant corporate AI software developers for ESA, the others being CRI and MBB-ERNO. The latter two are both frequent partners and competitors of Vega.
2. While Vega has been a heavy user of AI shells in the past, they feel somewhat burned by relying on small AI companies (Software A&E was mentioned specifically). They seem to be moving towards relying on themselves and using major general software products like C++ as a development environment.
3. They have a very heavy emphasis on procedures as a mode of operation with satellite control and maintenance. They see a strong resistance in Europe to realtime diagnostic systems--ESOC operations standards are to examine and validate all procedures in advance and not create solutions on an as-needed basis.
4. When asked about any operational "wins," they pointed to the MWS system for the METEOSAT. This is an interesting satellite control workstation (discussed in more detail below) that is currently running in shadow mode for the satellite.
5. ESOC and ESTEC are the biggest government customers for Vega. ESOC is responsible for satellite operations through orbital checkout (then the satellite owner takes over) and does its own

research on operations through companies like Vega. ESTEC is the research center for advanced technology on the satellites themselves. They do ground checkout of satellites among other things. This makes ESOC and ESTEC occasional friendly competitors on the boundary between those two functions.

Specific Systems Discussed and Demonstrated

SACV—Satellite autonomy concept validation

This was a very interesting design study and prototype done in cooperation with MBB-ERNO. The goal was to simulate a fully autonomous satellite capable of onboard FDIR, re-scheduling to reflect changes in onboard resources, expanding high-level macro commands, and communicating satellite status intelligently to ground controllers. The work also included a ground control system that could create and validate plans to be uplinked to the spacecraft, and could interact with all of the onboard functions when manual intervention was desired. The concept was implemented at ESOC using an existing simulator for the EURECA satellite (launched recently from the Shuttle and to be picked up by the Shuttle). The total effort to date has been about 4 man-years split between Vega and MBB-ERNO.

The OBMM (Onboard Mission Manager) encompasses the “smart” part of the SACV prototype. The onboard system uses MARS (from MBB-ERNO) for planning and CONNEX (also from MBB-ERNO) for FDIR. CONNEX is a strictly experiential, matrix-based diagnostic tool. The prototype actually had a much greater emphasis on assisting operations than diagnosing faults. Note that the kind of high-level commands OBMM was capable of understanding were strictly macros; it could deterministically expand macros but had no context sensitive capability to understand and instantiate true high-level goals.

The ground system for SACV consisted of three components:

1. COMPASS, a MARS-based command and planning assistant
2. BROWSER, a satellite status analysis system
3. RECONSTRUCTOR, a visualization tool to help ground controllers understand the current status of onboard satellite plans.

SACV gave a major demonstration in March 1992 when four scenarios were shown:

1. the updating of the onboard Master Schedule from goal oriented commands (GOCS) (which are really macros)
2. onboard autonomous rescheduling in response to three kinds of power system failures
3. recovery from an scientific instrument (satellite payload) failure onboard
4. autonomous changes to the onboard schedule in response to a serendipitous external event (a gamma ray burst of interest to an onboard instrument)

Vega considers the demonstration a major success and produced an excellent final report in August 1992. (This is available in the Appendices to this Study.) However, Vega believes that the

successful transition to onboard use of such autonomous systems may take as much as 10 years. This is partially due to technical reasons like the availability of enough onboard computing (Thompson CSF may be working on a rad-hardened SPARC chip). Vega believes that it is mostly due to political considerations within ESA and overall extreme conservatism in satellite operations.

Noticeboard—

This is a tool that facilitates communication among multiple dynamic processes. The system is what most would consider to be a relatively straightforward blackboard system that was motivated by the difficulties Vega had in building the MWS workstation for the METEOSAT satellite (discussed below). Vega wanted a system that could communicate between multiple different hardware workstations and software tools without users having to worry about specific interface protocols. From the demonstration we saw, the tool seems to work well and includes an easy to use demon facility. Vega considers Noticeboard a bit too slow for real-time use. It is not yet a commercial product and is in use only at Vega (on three projects) and at ESTEC (for four projects). It does not seem to have particularly sophisticated triggering or agenda mechanisms that some US blackboard tools do have.

MWS (MeteoSat WorkStation)—

MWS was built for the EUMETSAT organization that paid for and operates the MeteoSat meteorological satellite. Its construction began in 1988, and it is now in its third incarnation as MWS3. It runs on a Sparc, doing both subsystem diagnosis and operational procedure execution. However, the diagnosis portions are simple rule-based, off-the-shelf mechanisms; the major emphasis was on the procedural execution component. The demonstration we saw showed an integrated system for validating procedure uplinks, monitoring satellite status, and warning about limit problems (of several severities) in downlink telemetry. As such, it bore a great resemblance to SHARP. It did a nice job of guiding mission controllers through possibly complex procedures which could have multiple branch points. It did not have much in the way of features for online modification of procedures or noting of suggested changes to procedures; this illustrated a possible difference to many NASA mission control centers where the controllers are highly skilled engineers—Vega stated that European mission controllers tended to be relatively unskilled technicians. The display seemed a little busy at times. Vega may have given the customer exactly what he said he wanted rather than what he “really” wanted; they were frank about their desire to let conventional user modes of operation totally drive their systems.

MWS3 is in current use in “shadow” mode for the METEOSAT satellite. An excellent final report is available in the Appendices to this Study.

OES—a model based system for automatic creation of procedures—

OES was built as a demonstration of the automatic generation of operational procedures from system fault models. It works for the power and attitude control systems of the Olympus telecommunications satellite (although that seems merely the test domain, there was no attempt to currently get into operational use). The system uses simple functional models of major satellite components to a priori produce single fault scenarios for procedural remedy. The mechanism seems more

sophisticated than simple digraphs methods (e.g., FEAT), but considerably less sophisticated than state-of-the-art model based approaches employing qualitative physics.

FOPSET (Flight Operations Procedures System Engineers Tool)–

The final system discussed was work-in-progress to provide a tool to assist satellite engineers in producing operations procedures. It concentrated on closely modelling current methods: input was mainly textual, output in the form of structured tables. Several of us felt a more graphical approach to procedure generation would be better from a user interface point of view; the Vega personnel seemed to agree but said that they had met extreme user reluctance to depart from the current protocols for doing business. In other words, current methods were being automated rather than allowing the opportunities of automation to modify current methods.

AI APPLICATIONS INSTITUTE (AIAI)

Organization Overview

The Artificial Intelligence Applications Institute (AIAI) is part of the University of Edinburgh, and is co-located with the Department of Artificial Intelligence in Edinburgh, U.K. AIAI was established in 1984 to act as a technology transfer organization for artificial intelligence research taking place in the university. AIAI is primarily concerned with the practical delivery of solutions to their clients' problems through the use of AI techniques. While AIAI is "solutions oriented", it is fair to say that they are largely concerned with the application of AI technology. If a potential client presents a problem that does not appear solvable by the institute's AI technology set, then that client is unlikely to be taken on. AIAI achieves technology transfer by three mechanisms: consultancy, project development, and technical training. AIAI is self-funded and non-profit, and does roughly one million pounds worth of business each year. There are three main technical foci for the Institute: knowledge-based modelling and reasoning, knowledge engineering methods, and planning and scheduling. There are roughly equal numbers of people in each group, with around 24 technical members on the AIAI permanent staff. Present on behalf of AIAI were Austin Tate and Robert Rae.

General Observations

AIAI carries out a great deal of work that does not involve the European space industry, and this summary will not discuss that work at all. In terms of AI applied to space, the institute has carried out roughly eight projects. One of these was funded by the U.K. government, six were funded by parts of ESA, and one has been (and continues to be) funded by the U.S. Air Force. The bulk of the AIAI work that is directly relevant to space has been funded by ESTEC and ESOC; in this work, AIAI has often acted as a partner to other European organizations in the context of some larger project. AIAI's role is often to provide advice and consultancy regarding the more complex aspects of planning and scheduling. Specifically, AIAI have advised their collaborators on how to represent plans and schedules, and how to perform basic reasoning functions over those plans and schedules. In essence, AIAI has had the role of general system architect and helps to ensure that the various software and aerospace companies involved have an understanding of the relevant AI technology.

List of the Systems Discussed and Demonstrated

Funder	Project Name
SERC	T-SAT/T-Sched
ESTEC	Plan-ERS
ESTEC	Optimum-AIV
ESOC	Tools Evaluation
ESOC	Simulator in Pro-Kappa
USAF	O-Plan2

ESTEC Planning and scheduling consultancy
EUMETSAT Mission planning and control

Descriptions of specific *systems* developed in these projects follow.

Specific Systems Discussed and Demonstrated

T-SAT/T-Sched-

This work was done for the U.K. Science and Engineering Research Council (SERC), and involved working with a number of other U.K. organizations to design a leading-edge technology satellite. The AIAI contribution was a design and implementation of a mission sequence and on-board schedule execution system. The first implementation of the mission sequencer was done in the O-Plan architecture (described below). This implementation demonstrated some success, but there were problems that gave rise to further requirements on subsequent versions of the O-Plan system. The final implementation called T-Sched used a resource-centered scheduler.

Plan-ERS-

This work involved consultancy to a European consortium that included CRI (Denmark), Matra (France), and AIAI; the work was funded by ESTEC. The work was carried out from February 1987 to June 1988. The purpose of the project was to identify planning and scheduling problems in the space industry, and to analyze the feasibility of applying AI techniques to solve the identified problems. The project's first stage involved identifying the following possible planning problems: spacecraft assembly, integration and test plans (generation and execution); earth remote sensing spacecraft mission planning; data relay satellite mission planning; Hermes return phase planning; Columbus polar platform mission planning. The project's second stage selected the Earth Resources Satellite 1 as a test case, and built a prototype system for doing ERS-1 mission sequencing. The prototype system, Plan-ERS, was able to generate mission sequences for certain simplified demonstration problems.

Optimum-AIV-

This work was funded by an ESTEC contract awarded to a consortium that included CRI (Denmark), Matra (France), Progrespace (France), and AIAI. The contract dates ran from June 1990 through April 1992. The purpose of the project was to develop an adaptable kernel for supporting activity planning and the life cycle of spacecraft assembly, integration, and verification. AIAI's role in the project was to provide advice and consultancy on the underlying plan and resource representation. The AIAI design allowed for the following basic functionality: domain knowledge editing and plan specification; plan generation (using precedence links and configuration constraints); schedule generation (involving constraint satisfaction -- both temporal and resource usage); plan repair and plan execution monitoring.

The AIAI representation allowed domain objects to include activities, events, precedence links, resources, and calendars. Actions (and indeed entire plans) had a hierarchical structure, used to encode natural project structure. The constraints associated with a plan include predecessor/successor links, preconditions/effects, temporal constraints (target dates and durations), and resource constraints (requirements).

The system was demonstrated to ESTEC working on a satellite assembly, integration, and test problem. The final system used a sophisticated interface with the commercially available Artemis scheduling system.

O-Plan -the open planning architecture

O-Plan is a general architecture for the definition, manipulation, and execution of activity plans. This work extends from the now-classic AIAI work on NonLin, the first technically correct partial-order planner. The O-Plan architecture and plan representation have revealed themselves in both the Plan-ERS and Optimum-AIV projects. In a sense, both Plan-ERS and Optimum-AIV are applications of the basic O-Plan research. The overall architecture includes an agenda-based search mechanism, where a user can provide knowledge sources specific to a given problem or domain. While the architecture does attempt to address plan execution, not much work has been done in this area yet.

Eumetsat

New work taking place over the period 1992 to 1994 involves AIAI working with British Aerospace to produce the mission planning and scheduling support system for the EUMETSAT organization. This organization has responsibility for operation of the Meteosat range of weather satellites and for the distribution of meteorological products these satellites produce.

COMPUTER RESOURCES INTERNATIONAL (CRI)

Organization Overview

Computer Resources International is a private Danish company with its Headquarters office located in Birkerød, just outside of Copenhagen. The company is 50% owned by IBM with field offices also located in India and Luxembourg. Personnel are also scattered throughout Europe as software support personnel in groups of one to four people. CRI employs roughly 600 people mainly in the area of providing software support services to European and Danish customers. Out of six primary Divisions, one is devoted to Space Applications, and it is primarily in this division where AI applications are developed and explored as a part of a total solution to customer problems. The primary customers of the Space Division are ESA Directorates such as ESOC (European Space Operations Center) and ESTEC (European Space Research & Technology Center) with some support beginning this year to ESRIN, an Italian-based ESA directorate responsible for being a repository of information and science data. Another large source of project funding comes from the European Commission which funds the ESPRIT program, encompassing a billion dollar program devoted to information processing in general, but which attempts to stay away from the Space Applications arena.

Within the Space Division, two sections are devoted to supporting ESTEC operations by both "body shop" type support services and also by the specific development of satellite checkout systems. One section is devoted to operations support for ESOC where CRI is one of a consortium of five companies continuously competing for development of satellite control systems. A new section, Information Systems, has just won its first contract supporting ESRIN and hopes to build up to a similar support contract as that held with ESTEC for software support. The Support Systems section, headed up by Mogens Nielsen (our host), is where all of the projects to be discussed have been or are under development. It is a combination of Operational projects and new technology development. This appears to be an excellent organizational strategy for integrating AI applications within operational systems for satellite support. From developments in this section, technology can be transferred to operations projects in other sections also. But it must be stressed that AI is seen as one facet of many solutions to systems development problems and that there is no pressure to go from research prototypes to operational systems.

General Observations

CRI considers the deployment of AI applications as a byproduct of a total operational system. AI system development and deployment is not a primary goal, nor does it appear to be a corporate focus. Functionality, not AI, is the product.

CRI is generally not interested in marketing commercial software projects but rather the development of systems uniquely tailored to customer needs. Software reusability comes from their ability to reuse the developments from project to project. A "win" would be defined as a technology development that becomes reused in several customer contracts.

PRECEDING PAGE BLANK NOT FILMED

CRI provides advanced technology projects to ESTEC who has the responsibility to transfer this technology to the operations sector, but does not always have a customer defined at the time of system procurement.

There appear to be some technology transfer problems, mostly due to politics, between ESTEC and ESOC, with contractors like CRI caught in the middle. Contractor frustration was evidenced by a discussion of the Plan-ERS system which was considered ready for use by ESOC, but which remains with ESTEC running in “shadow” mode to evaluate the quality of plans versus the actual ESOC planning system products.

The model of selecting five contractors to continuously compete for satellite control systems projects breeds familiarity with the established system platform and requirements, as well as promoting excellence in technology advancement and cost cutting measures for competition. The manpower rates were fixed up-front when the five consortium members were selected.

CRI views the AI market in Europe as just out of the “feasibility” stage and now in the prototype development stage. They are trying to show some benefits whereby AI technique use can proceed at a faster rate and gain acceptability.

CRI views competitors in the AI market primarily as Logica, Vega, MATRA, and BSO.

General Comment (not related to CRI)

The European Commission strategy of providing 50% funding, as well as the ESA practice of providing only partial funding on some projects, is akin to the US SBIR program. The goal here is to empower European companies with the technology development to allow them to excel in all markets utilizing software tools.

Once satellites, or space programs, go beyond the experimental stage, they are turned over (through contracts) to the private sector for Operations and Maintenance. Hence, AI (or any technology) “wins” are when private companies utilize it within the context of their satellite support contracts.

List of the Systems Discussed and Demonstrated

Funder	Project Name	Presenter/Lead Engineer
ESTEC	Optimum-AIV	Mads Aarup
ESOC	GMPT Study	Mads Aarup
ESPRIT	ITSIE	Ebbe Pedersen
ESPRIT	VALID	Inga With
ESPRIT	SIMPR	Klaus Heje Munch

Specific Systems Discussed and Demonstrated

OPTIMUM-AIV-

OPTIMUM-AIV: (Open Planning Tool Improving Use of Methodology for Assembly, Integration and Verification)

This system will provide support to planning activities for Assembly, Integration and Verification of spacecraft through all phases (life cycle) of spacecraft development. The system is a derivative of the Plan-ERS system, and CRI was part of a consortium of four companies which developed this product; CRI, MATRA-ESPACE (France), Progespace (France) and AIAI, UK. MATRA is actively pursuing marketing of this product as a general planning tool, while CRI sees its role as primarily tailoring the kernel of the planning tool to specific satellite support systems contracts.

This system appeared to be very robust in planning support and serves as a planning assistant, rather than an automatic planning system. The system runs on Sun-4/SPARC computers and is coded in Lisp(CL/CLOS). During the planning cycle, it will return to the user for decisions and conflict resolutions. The system first solves temporal constraints then resource usage. The repair capability attempts to resolve resource conflicts with "simple activity shifts." However, if the problem cannot be quickly resolved, the user will be queried for input. There are no plans to implement an automatic repair capability, and the system will remain interactive. The system interacts with Artemis for output and therefore provides commonly understood reports for the user, as well as a common scheduling database (although some information is lost when transferred to Artemis for output).

Development within this project is considered development only and there are no plans to operationally deploy this system, except as the concepts apply to projects such as GMPT, discussed next.

GMPT-

The GMPT, Generic Mission Planning Tool: An application of Optimum-AIV project, is an ESOC sponsored study which will define a generic mission planning facility in support of operations planning and scheduling for spacecraft operations. CRI is the prime contractor, and Science Systems Ltd. is a subcontractor. There are two phases to the project. The first phase will provide a survey of current mission planning approaches, analyze future ESA requirements for mission planning systems, define generic mission planning user requirements and related interface standards. The second phase will define baseline software requirements and overall architectural design resulting in a preliminary prototype system.

Phase I is complete as of September 1992 and resulted in the generation of three documents, Mission Planning Approaches, User requirements Document, and Interface Standards. Phase II is now in progress and was expected to be finished in November 1992. The aim of phase II is to elaborate GMPT concepts, i.e., producing a software requirements document, and a prototype demonstrating some of the elaborated concepts. The prototype will be based on OPTIMUM-AIV and will thus be written in Common Lisp. The prototype is not supposed to be used operationally, since this would require porting the entire system to C++, the language used for the new SCOS-II (Spacecraft Control and Operations System) infrastructure.

Out of all the systems demonstrated at CRI, this system stands the highest chance of making and operational impact at ESOC. This is mainly due to the fact that it is sponsored by ESOC rather than ESTEC, which seems to have some problems in the technology transfer arena. Nevertheless, the GMPT project is only tasked with producing a prototype for a demonstration in connection with the final presentation, and there is no driver or guarantee that ESOC will provide further funding for turning GMPT into an operational system.

ITSIE-

ITSIE: Intelligent Training Systems in Industrial Environments. This is an ambitious training project sponsored by ESPRIT totalling 3.5 yrs and 60 man years. CRI is again part of a consortium of companies working on the project. This training system provides the methodology for building intelligent training systems, domain models for industrial applications and some instructional strategies. In general, it bears the same problems as other intelligent training systems, in that the domain model is the largest consumer of programming resources, and as such makes it infeasible for most companies to field.

ITSIE offers three training strategies. First, the Tutoring strategy which plans the entire session, monitors the trainee behavior, provides error diagnosis and detailed explanations of errors. Second, is the Coaching strategy with the planning, monitoring and diagnosis of Tutoring, but which focuses on practice and tests, offering advice only in cases of serious misconceptions or upon request. The third is the self-teaching strategy, which allows input by the trainee or instructor on training objectives. The monitor and diagnosis is only to support the self-evaluation of the trainee and serious errors are noted to the trainee. There are also three training methods, the problem-driven, example-driven and test-driven methods.

The first version of this system is complete, and two applications have been written. One application covers an electric metering unit, and the second began as an attempt to model a power plant but was scaled down when the domain model generation was determined to be too difficult and time intensive. Various tools that have been developed in Lisp are in various stages of development and serve as demonstration/prototype but would need to be converted to C for actual operational use. CRI has no intention to market this tool, however, a proposal has been submitted to ESTEC to apply ITSIE to Astronaut training. Chances of possible win becoming an operational training tool are minimal.

SIMPR-

The SIMPR, Structured information management: Processing and Retrieval, project, ran from January 1989 to June 1992 and was conducted by a Consortium of nine members ranging from privately-owned companies to various universities and Technical Institutes. The system provides authoring and editing software support for creation and management of very large information data bases resulting from items such as technical reports and manuals. The system is based on syntactical analysis of textual information for classification purposes. Statistical analysis is not used. Essentially, it is a knowledge-based system whereby linguistic knowledge is combined with domain knowledge of experts for classification and retrieval. The indexing system is called MIDAS, and the

retrieval system is called PROSPECT. Our group saw the MIDAS portion of the project demonstrated.

Once again, the purpose of this project was not to develop an operational system but to perform the research to demonstrate the available new technologies and enable companies to develop more efficient information retrieval systems.

VALID-

VALID, a validation tool for knowledge base consistency, is another ESPRIT project which CRI performed as part of a consortium of four companies including Cognitech from France, and CEAB and UPM from Spain. The project went from January 1989 to January 1992 at \$5.9 M covering 30 person years. The goal of this project was to develop methodologies for validating and checking the consistency of Expert Systems and their knowledge bases. To achieve these objectives, a Common Conceptual Representation had to be achieved and a Meta-language to convert knowledge bases into this representation had to be developed. In addition, a validation environment had to be established for the user. Then, various sets of validation tools had to be built for use dependent on the validation or verification goals applicable to the specific life cycle in which the expert system existed.

The system is developed in Common Lisp on a Sun SPARC and has translators for three expert system development shells, KOOL, Milord and Ops83 (done by CRI). Seven types of validation tools are provided with capabilities such as the ability to check the control knowledge of a knowledge base and verification of knowledge base constructs and contents.

Once again, the purpose of this project was not to develop an operational system but to perform the research to demonstrate the available new technologies and enable companies to develop uniform knowledge base validation tools.

EUROPEAN SPACE RESEARCH & TECHNOLOGY CENTER (ESTEC)

Organization Overview

ESTEC has the primary responsibility for technology development within the European Space Agency. ESTEC personnel oversee technology demonstration projects relevant to ESA missions which are conducted by consortia whose members are drawn from the European aerospace and software industries and the European research community. They also develop testbeds for the evaluation of technology.

ESTEC is one of three ESA centers established to date, the others being ESOC, located in Darmstadt, Germany and ESRIN, located in Rome, Italy. ESOC is the principal space operations center while ESRIN is the central data archiving and management center.

General Observations

Operating within an ESA-wide research and development budget of \$42 million per year, ESTEC draws on a workforce whose composition strictly follows how the overall ESA budget is partitioned along national lines. France has a plurality of about 30%. There are both civil servants and contractors.

ESTEC is organized as a matrix, with each individual having a role in the monitoring of ESA projects as well as having a place in a hierarchical technical directorate which acts as support to other directorates at ESTEC, i.e., the space project directorates.

There are about 10 persons at ESTEC involved in projects which involve AI work. Among these are two AI-trained persons and one of them, Francois Allard, is clearly an AI champion. ESA projects involving AI work are scattered throughout the technical directorate, although most AI work is conducted out of the Automation and Informatics Department (The hierarchy at ESTEC is as follows: directorate, department, division, section.).

ESTEC personnel issue ITTs, much like RFPs or NRAs, review proposals, and make recommendations about merit and funding to the Industrial Policy Committee of ESA. To evaluate the AI aspects of these proposals, ESTEC personnel draw on the expertise of external consultants. These consultants form an ESA AI Steering Group. This group consists of prominent European AI researchers such as Guy Boy of EURISCO in Toulouse, France, Yves Kodratoff of the University of Paris, France, Robert Wielinga, University of Amsterdam, Holland, and Austin Tate of the AI Applications Institute in Edinburgh, U.K.

Responses to ITTs are made almost exclusively by consortia of aerospace companies, software companies, and research institutions. The reason given for collective proposals being the default is that the expertise needed to carry out ESA projects is usually distributed among several sources. In the US, this would not in itself be a sufficient reason for forming a consortium, for an aerospace company in responding to an RFP typically would acquire additional expertise either through

subcontracting or recruiting. A deeper reason for the prominence of consortia in Europe may be that they simply represent a more natural way of doing business in Europe's multinational environment.

ESTEC awards may include university contracts, but they do not include university grants. Grants for more basic research are made directly through blanket research projects funded by the European Commission, such as ESPRIT, which are funded from a higher level in the European Community.

The principal criterion for success of an ESA project is that the user, typically a member of industry, should continue funding the work after the completion of the project. Another, less objective criterion is that there should be a clear demonstration of benefit.

When asked to identify the project with the greatest success to date, Allard pointed to the MARS scheduling system, developed both with ESTEC and with MBB/ERNO funding. This system is being applied to Columbus payload operations scheduling.

The Expert Systems Demonstration Programme

ESTEC launched the expert systems demonstration programme in September 1991, with the goal of demonstrating the application of knowledge-based systems technology in space. Called the Expert Systems Demonstration Programme (ESPD), the program will allocate 3.4MAU to 16 projects over the next four years. This more focused program will extend the ESTEC investment in projects involving AI technology, which has been 5MAU to date.

The ESPD is divided into two broad application areas: payload engineering, and crew and operations. The capabilities being targeted include design assistance for payloads, experiments, and associated testing, reactive planning, automated analysis of experiment data, procedure generation, selection, and verification, computer-assisted training, sensor placement, automated monitoring and control, and knowledge capture.

Specific Systems Discussed and Demonstrated

CaeCALX-

CaeCALX is a geometric resource allocation system which assigns cargo items to storage racks onboard space platforms. The system was developed by two Spanish teams from the Grupo de Mecanica del Vuelo and the Instituto de Investigacion Tecnologica. Their system is a successful combination of straightforward mathematical analysis and well-known AI search techniques.

The degrees of freedom which enlarge the search space are: 1) storage racks may be subdivided into different size drawers, 2) cargo items may be assigned to any drawer, and 3) the position and orientation of each cargo item is unspecified.

The constraints which reduce the search space are: 1) items which contribute to a common purpose must be placed in the same drawer, 2) items expected to be used routinely must be placed near the front of a drawer, and 3) fragile items must not be placed at the "bottom" of a drawer relative to acceleration during launch.

Global constraints to be satisfied are: 1) maximum overall volume efficiency, and 2) placement of the overall rack center of mass. The investigators found that there was sufficient constraint available in the problem to employ a best-first beam search and avoid backtracking. In their approach, several incomplete solutions (partially filled drawers) are expanded in parallel and evaluated against efficiency of volume use and the other constraints. Final solutions utilizing 85% of available volume were achieved on average, far exceeding the performance of humans manually interacting with a CAD system.

PREVISE-

PREVISE is a system for partial automation of crew procedure generation and verification. This project is of two-year duration and began in mid-1991. The consortium which is executing this project consists of Matra Marconi Space and Dassault Aviation of France, Spacebel of Belgium, and Hernandez Engineering of The Netherlands.

The strength of this project is the amount of effort devoted to tools for encoding the background knowledge needed to define and reason about procedures in a space domain. PREVISE includes a knowledge editor which enables users to define the set of objects, the vocabulary and the set of constraints appropriate to their domain. There is a procedure editor for defining procedures from the encoded domain knowledge and to instantiate constraints to be used for procedure verification.

A procedure checker performs several forms of verification: logical errors, time and resource conflicts, missing steps, violations of global constraints, and poor optimization. A procedure executor enables a user to interactively simulate the execution of a procedure to perform manual verification.

PREVISE will run on a SPARCstation under Unix/Open Windows. It is a primary example of the uniformly strong European emphasis we witnessed on software tools for defining, manipulating, and reasoning about procedures.

MARS-

The Mission Activity and Resource Scheduler (MARS) is a mission scheduling tool developed at MBB/ERNO partially with funding from ESTEC. MARS runs in Symbolics CommonLisp on a SPARCstation with a LISP coprocessor board. A complete port to SUN/Unix/XWindows is planned. MARS has been evaluated for Columbus payload operations.

The feature of MARS most highly valued by its users is its highly flexible language for defining tasks, resources, constraints, and scheduling strategies.

Once a user has specified a scheduling problem, MARS detects resource conflicts and shifts tasks forward in time to resolve conflicts. An interface permits the user to manually edit the schedule. MARS appears to be similar in capability to the COMPASS system developed by Barry Fox of McDonnell-Douglas.

MARS-generated scenarios produced efficient use of European resource shares on Freedom/Columbus, as follows: 92% of crew resources, 89% of power resources, and 100% of upload capacity. MARS is considered to be the greatest success story to date among ESTEC-funded projects.

The next generation MARS system is the New Expert Planning Tool for Users in a Networked Environment (NEPTUNE). NEPTUNE has been developed and delivered already to support a form of distributed scheduling where a single schedule is worked at different levels of abstraction by different planning agencies within ESA. A C version is under development.

Huygens DMS ES-

The last briefing and demonstration we received at ESTEC described a prototype for an onboard data management expert system for the Huygens Titan Probe (part of the Cassini orbiter mission at Saturn). The prototype was developed in the KEE knowledge-based system environment with additions in CommonLisp.

The ambitious concept involves encoding onboard models of Titan, its atmosphere, and probe behavior, and to perform reasoning to control the probe to maximize science return during the brief mission. Unresolved issues include: knowledge representation for multiple types of models, testing and validation of the expert system, and real-time performance.

The scope of the concept is reminiscent of Mars Rover mission designs where the mobile robot performs all of: navigation, environment model updating, sample collection, and reasoning about scientific goals. The Huygens DMS ES project has not been approved by ESA for further investigation at this time.

BSO/ORIGIN

Organization Overview

BSO is a systems house based in the Netherlands that has a worldwide network of BSO-owned companies. Inside the Netherlands, the organization is known as “BSO”; outside the Netherlands it is known as “Origin”. The overall BSO/Origin group is incredibly decentralized, where each component company is run as an autonomous profit-and-loss unit. The head office for BSO/Origin is in Utrecht. BSO/Origin offer four basic classes of service: information systems, automation technology, organizational advice, and management support. The technologies utilized in these services include AI, multimedia, optical storage, image compression, virtual reality, instructional technology, and language technology. The markets to which BSO/Origin addresses itself include aerospace and defense, telecommunications, logistics, and general government. The group we met with were from BSO/AeS, an aerospace and defense company that is part of the BSO network. There is a BSO company that specifically works on AI, but we did not meet with them since they have not done any space-related research or applications. To the extent that BSO/AeS does research, it is applied, focusing on specific client needs.

General Observations

BSO/Origin is an interesting company. The network of component companies that together comprise BSO/Origin appears to be extremely loosely coupled, allowing a great deal of freedom of organization. BSO/AeS has had space-related contracts with ESTEC, ESOC, MBB-ERNO, DLR, Fokker Space & Systems, and the Dutch Government. We met with BSO/AeS at ESTEC, as they offered to save us the drive over to their office in Nieuwegin. Present on behalf of BSO were Piet Veenstra (managing director), Tim Grant, and Bas Kooter. We received presentations from Tim Grant and Bas Kooter.

List of the Systems Discussed and Demonstrated

There were three main systems discussed during our meeting, listed below.

1. Activity Scheduling System
2. High Performance Capillary Electrophoresis Payload Simulator
3. Cabin Atmospheric Pressurisation Subsystem Expert System

These systems are discussed in more detail below.

Specific Systems Discussed and Demonstrated

Activity Scheduling System (ASS), for the Dutch Utilisation Center (T. J. Grant)

The activity scheduling system (with the unfortunate acronym “ASS”) is a prototype of an integrated system intended to eventually support the design of payloads and experiments, the planning and scheduling of experiments, the control of nominal and non-nominal payload operations, and the post-mission comparison of achieved versus planned activities. Traditional approaches to these different phases of activity result in disparate software systems that cannot be easily integrated and share results. ASS is intended to act as a proof of concept for an integrated approach that combines all of these phases of activity. This work is funded by the Dutch Government, and is a single element in a bigger project, the goal of which is to build what is called the Dutch Utilisation Center. Functional goals of the system are being designed by consultancy with various organizations, including ESOC. Essentially, ASS allows a user to build a state-machine that describes the device to be modeled, and provides a mechanism for enumerating the exponentially-many behaviors that could be produced by that machine. This approach has been pursued with some success in previous work in AI Planning (the ERE project at NASA/Ames, for instance), and in Discrete Event Control theory.

High Performance Capillary Electrophoresis Payload Simulator (HPCE-PLS) (T. J. Grant and M. H. Wigmans)

The HPCE-PLS is a high-fidelity software simulation of a commercially-available HPCE instrument. The project is funded by the Dutch Government as part of the overall Dutch Utilisation Center project. HPCE is a general experimental method for separating molecules; the method works by moving charged particles through a fluid under the influence of an electric field. The HPCE-PLS simulates the operation of an HPCE in order to determine the performance of the instrument under space conditions, primarily to detect possible instrument faults and procedural problems without involving expensive hardware. The HPCE-PLS does not incorporate any AI techniques of interest beyond basic object-oriented programming. There are plans to add various AI-based capabilities, however, including the automatic generation of an architectural design document by combining the HPCE-PLS and the activity scheduling system, ASS, described above.

Cabin Atmospheric Pressurisation Subsystem (CAPS) Expert System (T. J. Grant and B. M. Kooter)

The CAPS expert system was developed for ESA; the CAPS prototype is currently installed in the Crew Work Station Test-Bed (CWSTB) that is located at ESTEC. The test-bed is intended to provide a reasonably high-fidelity environment in which to study human-related uses of computer workstations. The CAPS expert system is inherently human-related, as its goal is to monitor and control the partial pressure of various critical component gasses in the astronauts’ overall breathing mixture. The CAPS system was originally called the “N2O2 Expert System”, and its original goal was to investigate the value of combining expert systems and Graphical User Interface (GUI) techniques in support of spacecraft operations. The goal for the CAPS system has evolved from this starting point, where the emphasis has been on using the system as a vehicle for various human-computer interaction studies. BSO feels that the final system is best viewed as a GUI with an under-

lying expert system, rather than as an expert system that happens to have a GUI. The AI techniques used in CAPS are reasonably traditional for an expert-system; however, much like Georgeff's work on the Procedural Reasoning System, the authors of CAPS put a great deal of work into the representation of monitoring and diagnosis procedures.

SUMMARY REPORTS

EUROPEAN SPACE/AI APPLICATIONS: A SUMMARY

Introduction

The European AI market is one that has developed as a result of searching for the optimal solutions for development and improvement of spacecraft design and operations, both on the ground and in-orbit. Most of the applications are a direct result of contract awards which stipulate the problem to be solved or process to be improved, leaving the solution methodology and design to the contractors. This approach in turn, has led both the commercial sector and the government centers to investigate AI technology as a part of an overall systems solution. AI system development and deployment is not the primary goal, nor does it appear to be a corporate focus. Functionality, not AI, is the end goal. The product of the technology program is primarily in the form of demonstration prototypes, not operable systems. AI projects emphasize utilization of existing capabilities, rather than stressing basic research, for the purpose of producing demonstration prototypes and generically applicable systems which are able to span a multitude of applications.

The European Space Agency sponsors such demonstrations in hopes that they will improve overall contractor capability and influence internal company decisions. This approach causes their criteria for defining success to be quite different from the U.S. criteria, where the system must be in full operational use to be labelled as a "success". The effectiveness of this approach is demonstrated by the example at MATRA of the Arianexpert application. Prior to this application, only 2% of the funding for artificial intelligence research and development came from internal sources. As a result of the successful deployment of Arianexpert, Matra Marconi Espace now supplies 30% of the funding for internal AI research and development.

European AI applications, mainly in the form of prototypes, span multiple disciplines related to all-phases of spacecraft operations. These phases include design, monitoring and failure diagnosis, planning and scheduling, operations support and procedures management. The current primary focus appeared to be in the areas of planning and scheduling as well as procedures automation and management. These are seen as two areas where large cost reductions and efficiency improvement can be achieved.

As a result of the highly focussed European space program, most applications, in use or under development, are targeted for deployment in ground atmospheric satellite operations or for use with the upcoming Columbus mission.

Technology Development and Transfer

The European model of technology development and transfer to the space market is somewhat different from the U.S. where the government is the primary custodian of the space program, and hence both the developer and recipient of technology. The European Space Agency (ESA) develops space systems primarily through commercial contracts, and once developed, turns them over to the commercial market for operations. ESA orchestrates rather than implements the European space

program, and therefore has a role different from NASA's in the development and transfer of technology.

The AI technology projects sponsored through ESA are required to be multi-national and often are multi-site implementations. Management of such activities creates a cumbersome and difficult environment for technology transfer and operational use of prototypes. There is little pressure by ESA to show operational impact, but rather an effort to invest in the European space infrastructure supported by the commercial sector. The bottom line seems to be educating the industrial base through AI system development rather than implementing applications for efficiency or safety reasons.

ESA has divided its tasks primarily between two independent suborganizations (directorates), ESTEC and ESOC. ESTEC is the European Space Research & Technology Center and is responsible for the space program's technology development. ESOC is the European Space Operations Center, responsible for oversight of ground operations support for European missions and satellites. Nevertheless, technology development is sponsored by both directorates, sometimes in parallel, and there appear to be barriers to technology transfer between these directorates. While the quality of applications developed under each sponsorship is quite good, most "successes", per the American definition, which have come to operational use were developed and sponsored by ESOC, the operations directorate.

But within the European space market, one must also consider successes to be those applications whose commercially developed core technology has been reused or has the potential for reuse in other systems which the company might develop. This seems to be the approach taken by many of the European companies who look to in-house development of reusable technologies rather than stand-alone prototype applications.

In general, by studying the European model of technology development and transfer, we must be careful not to judge based on the U.S. model of a government space program and corresponding technology development. ESA is a development agency only, and operations of spacecraft fall in the commercial sector. Hence, true technology transfer is not achieved by inserting the technology within ESA programs, but rather by each company using the technology it developed thru ESA (or other) funds to enhance their satellite control systems or spacecraft operations of the future.

Current Applications

The site reports included in this document provide extensive detail on each application and its uses, and we will not attempt to duplicate this information here. In this Summary, we summarize the types, maturity levels, uses and general reception of the applications.

There were many examples of impressive applied technical work at the sites we visited in Europe. Some of the locations which stand out in both volume and quality of applications development work were MATRA Marconi Space, Computer Resources International (CRI), MBB-ERNO, and a smaller company called Vega Space Systems. Users of their technology included organizations like ArianeSpace and ESOC itself.

The demand for applications appears to fall into the three primary categories of data analysis, planning and scheduling, and procedures management. Figure 1 classifies many of these applications into these categories, and assigns a maturity level of 1 (lowest) to 5 (highest), indicating readiness for operational use. The Operational Use category indicates the projected use of the system in an operational environment.

Application	Data Analysis	Planning & Scheduling	Procedure Management	Maturity Level	Operational Use
Optimum-AIV		X		3	NO
Generic Mission Planning Tool (GMPT)		X		5	High Prob.
MARS		X		5	YES
Browser	X			5	NO
Reconstructor	X			5	NO
COMPASS		X		4	NO
OES			X	4	NO
FOPSET			X	4	High Prob.
MeteoSat Workstation	X		X	5	Shadow Mode
PREVISE			X	5	High Prob.
CaeCALX		X		3	NO
Expert Operator's Associate (EOA)	X			5	High Prob.
Radiometry Expert System	X			5	YES
Automatic Mirror Furnace Expert Sys	X			3	NO
GSDAS	X			4	NO
ESSOPE	X	X		4	NO
Activity Scheduling System		X		3	NO
HPCE-PLS	X			3	NO
CAPS Expert System	X			4	YES
Arianexpert	X			5	YES
Plan-ERS		X		4	NO
DIAMS2	X			4	YES
X-Analyst	X			3	High Prob.
POM			X	5	YES
Crew Support System			X	3	NO

Figure 1

The largest and most mature category is in the area of data analysis for purposes of system monitoring and failure diagnosis. However, the largest demand for applications appears to be in the areas of planning and scheduling as well as procedures management, since both these categories support the daily operations common to the sustaining engineering functions for satellites, the primary operational nature of the European space community.

There is some limited work in the areas of computer aided design, data management (both acquisition and maintenance), human factors, and training. In general, such applications were not targeted for operational use and remained internal to the organization developing them. These applications are not included in Figure 1 since they are in the infancy stages of maturity where actual use in the European space community is still far into the future.

Future Plans and AI Impact

It appears that the impact of AI on European space operations was first achieved by analytical tools like the Radiometry Expert System, Arianexpert, and the Meteosat Workstation. Later impacts in the planning and scheduling area came from systems like MARS, while procedural management systems like POM are just recently coming to fruition.

As an example of the effort to derive "generic" capabilities for specific types of applications based on previously developed special-purpose systems, there is X-Analyst which is derived from Arianexpert. This system also demonstrates the trend towards integrated systems as there are plans to integrate it with diagnostic capabilities derived from the DIAMS2 diagnostic system for applications in the areas of spacecraft integration and launch site testing analysis. The DIAMS2 system is also an example of the trend towards "hybrid" reasoning methods which combine techniques from disparate research areas, in this case qualitative reasoning and associative reasoning in diagnosis. Although DIAMS2 is an engineering prototype, it also highlights the awareness of the European research and development community to issues at the frontier of AI research.

Another noticeable trend is the emphasis on maintaining a human operator in the loop on procedures and analysis which could be opportunities for full automation. This trend may continue to hinder development of more highly automated systems. On the other hand, there are efforts at integrating AI with other technologies and evidence that AI research results and application requirements are influencing development in other fields. The SIMEX system at ERNO/Deutsche Aerospace is a case in point. The SIMEX project on model-based diagnosis is working to influence the development of new simulation models to provide data which is usable by other model based reasoning techniques. ERNO particularly highlighted their interest towards integrating the research efforts of AI groups with the research in other disciplines, such as simulation and CAD.

As stated previously, the demand seems to be in the procedural management and planning areas. Hence, it is in these application categories that we will most likely see operational AI systems in the future. Future trends are in the areas of highly integrated systems, as represented by projects like the Satellite Autonomy Concept Validation, which integrates both ground and spacecraft systems for highly autonomous operations. This is quite similar to the renewed interest of the American space program in vehicle health management systems. Both approaches are limited by spacecraft design and available computing power, so that implementation will have to wait for sophisticated spacecraft which can accommodate the technology. It therefore seems that while European space technology may be currently lagging behind U.S. capabilities, Europe displays excellent potential and the trends and general directions of space technology, especially in AI, are largely the same.

Observations and Conclusions

The major technology development funding sources for AI technology development are the European Commonwealth and ESA divisions with each having a similar approach to project funding much akin to the US Small Business Innovative Research (SBIR) Program to provide incentives to commercial industry towards technology development. The European Commonwealth only provides

50% funding and ESA often provides only partial funding which seems to provide the motivation for improving corporate infrastructure and capabilities in a technology area such as AI.

A consortium approach is enforced requiring contract awards to projects which obtain the cooperation of multiple companies. The granularity of definition at the contract award stage is much coarser and generally groups multiple functional tasks under one large umbrella task. While this approach provides the positive feature of enforcing inter-nation and inter-site cooperation, implementation and operational use of technology is sometimes hindered by such diversity of participation. There also seem to be some technology transfer barriers caused by political problems between ESOC and ESTEC which catch contractors in the middle and frustrates deployment of maturing applications, such as Plan-ERS.

An interesting model for technology development is in use by ESOC. They have selected 5 contractors to continuously compete for satellite control system development with established system platforms and requirements. This model seems to promote excellence in technology advancement and provides cost-cutting measures in the smaller scoped competition. Such approaches should be studied further for determination of their effectiveness.

It is also true that in Europe, as in the U.S., the role of the "AI Champion" cannot be underestimated. There were several examples of development efforts which were initiated, thrived, and had success due to the strong efforts of individuals in both the R & D and operational communities. Sometimes these efforts succeeded despite the reservations of upper management. However, it should be noted that among the visited aerospace contractors, support of senior company management in AI development appears strong and influential. It was also clear that the operational community still considers AI as substantially outside the realm of conventional methodologies, and therefore AI technology still awaits wide acceptance.

In general, by studying the European model of technology development and transfer, we must be careful not to judge progress based on the U.S. model of a government space program and corresponding technology development. European government agencies are overseers and caretakers of technology development while operations of spacecraft fall in the commercial sector. Hence, true technology transfer is not achieved by inserting the technology within government space programs, but rather by each company using the technology it developed through ESA (or other) funds to enhance their satellite control systems or spacecraft operations of the future. Europe seems to be attempting the development of a corporate infrastructure versed in latest state-of-the-art technologies which can respond to future envelopes of opportunity that will inevitably appear in space exploration as it responds to an increasing world market. If successful, such development could prepare the European industrial sector to be a driving force in future space markets.

EUROPEAN SPACE/AI TECHNOLOGY: A SUMMARY

Introduction

The purpose of this summary is to describe the AI technologies used in advanced software-based automation projects sponsored by the European Space Agency (ESA). These technologies include planning, scheduling, model-based reasoning, machine learning, procedural reasoning, knowledge-based systems and AI programming techniques.

General Observations

ESA activities are conducted at three centers, ESTEC, ESOC, and ESRIN. ESTEC, located in Noordwijk, The Netherlands, has primary responsibility for technology development within the European Space Agency. ESOC, located in Darmstadt, Germany, is the European Space Operations Center, but also sponsors and performs technology development activities. ESRIN, located in Rome, Italy, is the central data and information archiving and management center. This report is based on visits to ESTEC and ESOC, and to sites where ESTEC- and ESOC-sponsored work is conducted. The authors of this report did not visit ESRIN, so the report is necessarily incomplete in this area.

A few words about technology transfer models are in order. Within ESA, technology development projects are taken to the point of operational prototype demonstrations. A sign of success is when the user picks up further development and completes the delivery of an operational system. In contrast, technology development projects within NASA are taken further, to the point of operational system delivery.

One reason for the difference in technology transfer models is that ESA views technology transfer, at least partly, as education of the aerospace and software industrial base. From this viewpoint, ESA technology development projects serve the dual purpose of establishing the relevance of certain technologies to ESA problem areas and providing industry with experience in those technologies.

As part of this model, there is a clear emphasis on mature techniques. ESA does not fund in-house research, but does draw on a European AI Steering Committee to keep them informed of new or improved approaches which have reached a mature stage. This steering committee is comprised of some of the finest European AI researchers (Boy, Kodratoff, Wielinga, Tate).

This ESA model of the role of research in technology development is in contrast with NASA's, where successes have been achieved by pursuing projects where the resolution of basic research issues had been on the path towards what later became a completed operational delivery.

AI Technology Areas

Planning-

The emphasis of the systems that we viewed was primarily on scheduling. None of those systems was purely a planning tool. Two of the systems viewed by the NASA delegation, Optimum-AIV and Plan-ERS, did appear to have planning elements, but there was no clear definition and demonstration of a core planning capability. The same can be said of much US "planning" work. Most planning problems, when studied in greater detail, turn out to be scheduling problems.

It is normally agreed that scheduling is the problem of sequencing a pre-defined set of tasks subject to some temporal and resource constraints. The planning problem is that of determining the set of tasks to be scheduled. The traditional approach to combined planning and scheduling problems is to first plan, synthesizing a set of tasks, and then schedule, sequencing that set of tasks. Historically, people have been rather good at the planning part of the problem, but have required assistance with the picayune details of task scheduling.

Most automated systems thus first address the problem of scheduling, and provide small "hooks" for later planning functionality. In addition, most scheduling systems are in fact scheduling decision-support systems. The distinction is important: a scheduling decision-support system provides an electronic "spread-sheet" that keeps track of the details of time and resource assignment, something that people are not very good at. Such a decision-support system does not have to search a space of possible schedules, thus, it does not have to address the problem of search control. The human user whose decisions are being supported can provide all the guidance regarding which particular scheduling modification to make; the decision-support system simply tracks the decisions that are made and computes the detailed ramifications of each individual decision.

Thus, when one is presented with an "automated planning and scheduling system", one is usually really presented with a semi-automatic scheduling decision-support system. Planning is not really addressed, and the scheduling that is addressed is not fully automated. Nonetheless, such scheduling decision-support systems provide tremendously useful functionality, and indeed, they do actually address a large proportion of real problems.

With these points in mind, there is an element of planning technology to be found in both Optimum-AIV and Plan-ERS. Both of these systems have had the benefit of significant design advice from the AI Applications Institute of the University of Edinburgh. AIAI has extensive experience with state-of-the-art planning systems, and have had significant influence on the design of Optimum-AIV and Plan-ERS. Thus, it is likely that these two systems could in fact be used as true planning systems, autonomously synthesizing a set of tasks to be scheduled. However, such a capability was not demonstrated during our visits, and the real planning capability of these systems is not clear.

Scheduling–

The three most impressive scheduling systems are Mars/Neptune, Plan-ERS, and Optimum-AIV. Each provides an excellent user interface, and each interactively supports a user in making better scheduling decisions. Each system also provides some sort of automatic scheduling facility.

Mars (see MBB and ESTEC reports) provides an extremely flexible and general language for defining tasks, resources, constraints, and in fact, this language also allows for the definition of scheduling strategies. It has an excellent and well-defined formal basis, since it uses the standard “Allen Logic” of relations between intervals, and well-developed temporal propagation algorithms to ensure consistency of the underlying interval constraint networks. The automatic scheduling approach taken by Mars is essentially that of forward-simulation: times and resources are assigned forward in time, with the temporally earliest decisions made first. If backtracking is needed, it occurs according to a user-defined strategy. Neptune is a next-generation version of Mars designed expressly for distributed scheduling. We did not get sufficient detail on Neptune to be able to explain the basic approach taken.

Plan-ERS (see Matra, ESOC, and AIAI reports) is a scheduling system based, essentially, on AIAI hierarchical partial order planning technology. This technology is essentially state-of-the-art for AI planning systems, but the use of hierarchical partial orders is still uncommon in scheduling systems. The advantage of using an explicitly hierarchical representation is that scheduling problems can be solved at an “abstract” level, and then more detailed versions of the problem can be worked out lower down in the hierarchy. While this basic intuition seems sound, more work is required to formally and empirically demonstrate that it actually works well in practice. (It isn’t clear that the independence assumptions inherent in the hierarchy are actually valid in many real problems.) In Plan-ERS, the underlying AIAI hierarchical partial order representation has been augmented with resource and metric time constraints. Incremental algorithms are used to maintain time window consistency at each level in the hierarchy.

Optimum-AIV (see the Matra, CRI, and AIAI reports) is once again based on AIAI hierarchical partial order planning technology. The focus of this system has been scheduling for project management applications. Again, resources and metric time constraints have been added to the underlying representation, and incremental propagation algorithms are used to enforce consistency. One interesting aspect of systems based on classical planning representations is that they fairly easily and directly represent so-called “state” constraints. A state constraint is an ordering between two activities that is grounded in a cause-and-effect relationship. For instance, one task might open a given door, and a necessarily subsequent task might involve someone walking through that door. If another task intervenes between these two, and if this intervening task closes the door, then there’s obviously a problem. A scheduler that explicitly represents the cause-and-effect relationship between the tasks can notice the possible conflict, and can even suggest new orderings to resolve it. The underlying representation of Optimum-AIV allows it to perform this sort of reasoning.

In general, the scheduling systems exhibit a number of common characteristics. First, there is a definite focus on interactive schedule editing, with extensive graphical features provided for supporting a user’s scheduling decisions. There are also languages provided for allowing a user to define schedule transformations (and sometimes, scheduling strategies) directly. A system that provides

such a capability is, in some sense, highly parameterizable, and it should be possible to use the system in a large number of different scheduling application domains. As mentioned above, there are also excellent representations being used, and these are able to capture various sorts of state constraints. Finally, there is common use of incremental time bound propagation algorithms; such algorithms are used to maintain metric temporal consistency in a schedule.

There are a few areas in which further technology development would seem to help in getting AI-based scheduling applications more widely used in Europe. These include: new methods of search control, automatic schedule execution, and mechanisms for measuring and achieving schedule robustness. The schedulers that we saw emphasized the methodology of constraint satisfaction, which does not allow for maximization under a user-provided objective function. This can be achieved by viewing the problem as one of constrained optimization. These comments emphasize the fact that the field of scheduling is still developing and that much work remains to be done, both in Europe and in the U.S..

Model-based reasoning—

The key insight behind model-based reasoning is that a system representation which describes both structure (the pathways along which interactions take place) and behavior (the details of those interactions) implicitly encodes information about the possible causal relationships among events in a system, including the appearance of a fault in one location and its manifestation elsewhere. Because symptom-fault relationships are not actually enumerated, a greater degree of completeness can be achieved. Moreover, and most importantly, previously undescribed faults can be diagnosed.

The degree of sophistication of a model-based reasoning approach may be appraised at a high level according to whether both structural and behavioral knowledge is utilized.

Several of the ESA-developed systems which target monitoring and diagnosis assistance use traditional monitoring and associational diagnostic reasoning techniques, and have not yet begun to incorporate model-based reasoning techniques. Among these are the Meteosat WorkStation (MWS), developed by Vega Space Systems, the Arianexpert system, developed by Matra Marconi Space, and the CONNEX system developed by MBB.

MWS uses limit sensing to detect anomalies, and rules to capture symptom-fault associations. In many ways MWS resembles the SHARP system developed by JPL. The Arianexpert system performs discrepancy detection from a set of reference values.

The CONNEX system adds a form of uncertainty analysis to associational diagnostic reasoning. Symptom-fault relationships are captured in a matrix representation. The most likely faults are determined by taking weighted sums of the differences between observed data and known symptoms. The approach reduces sensitivity to noise and to the order in which parameters are checked.

Some other ESA-developed systems have begun to incorporate model-based reasoning concepts. The OES system, developed by Vega, generates fault recovery procedures from simple behavioral models of system faults. The models describe how single faults propagate. The procedures generated from these models direct an operator on how to recover from faults which may produce multiple

symptoms or fault cascades. The behavioral representation used in OES models is less expressive than that being used in the SADDAM task at JPL and JSC, which captures the quantitative and temporal aspects of behavior in order to support high-fidelity model-based simulation.

The SIMEX system, an extension to the CONNEX system, incorporates some of the insights of model-based diagnosis. Discrepancy detection is driven by an available model and simulation capability. Candidate generation is based on tracing structural connections in the model. However, SIMEX does not appear to employ behavioral knowledge, which is typically used for two purposes in model-based diagnosis: to further prune candidates, and to attempt to validate candidates by simulating known faults associated with candidates at the component level.

The DIAMS2 system, developed by Matra Marconi Space, was the most sophisticated diagnosis system viewed by the NASA delegation: DIAMS2 incorporates a more complete approach to model-based diagnosis, using both a component-connection structural representation, a quantitative behavioral representation, and a diagnosis algorithm reminiscent of KATE. DIAMS2 is also a hybrid system, using decision-tree based symptom-fault associational reasoning to initiate diagnosis, and model-based techniques to complete diagnostic reasoning.

Another hybrid system is the GSDAS system, developed by ESOC, which uses both rule-based and model-based diagnosis. ESA may be slightly ahead of NASA in exploring hybrid technology diagnosis systems. Such systems are likely to be developed and evaluated at the JSC Control Center Complex in the next few years.

The AMFESYS system, developed by ESOC, uses model-based simulation for a different purpose: to generate high-fidelity predicts to fill telemetry gaps for the AMF instrument on the EURECA satellite. This satellite is in view of ground stations only infrequently. The model-based simulation capability may be vital in reducing reorientation overhead when communication with the satellite is reacquired, especially if combined with model-based diagnosis. NASA may not have exactly the same scenario, but the concept appears to be an intriguing application of model-based techniques for efficiently and effectively recovering from communications downtime.

Procedural reasoning—

The term “procedural reasoning” typically refers to the explicit representation and manipulation of procedural knowledge. A procedure can be any time-oriented task description that one chooses. The idea of explicitly representing procedures in a computer appears to have originated with Mike Georgeff (of the Australian AI Institute); his system, called PRS (for the Procedural Reasoning System), provided the first real implementation of this idea. PRS has since been commercialized by a French company: the product is written in C, and uses the X windowing system to manage its graphical input and output.

What sort of reasoning is actually provided by a procedural reasoning system? While there is significant value associated with explicitly representing procedures so that people may examine, edit, and execute them manually, the real potential win seems to come from the computer automatically producing some useful inferences by examining the explicit representations. The demonstrations of procedural reasoning viewed by the NASA delegation focused both on underlying automatic

inference and the interactive creation, editing, execution of procedures. The number of variations on these concepts were impressive: there were systems for generating recovery procedures from traces of faulted behavior; entire architecture for encoding procedural primitives, assembling procedures, checking constraints, and validating procedures via simulation; workstations for flexible execution monitoring of procedures, including interrupt handling, parallel execution, contingent execution, and “clean-up” activities after aborts; interactive tools for on-line procedure modification.

Within the European technical community, the explicit representation of procedural knowledge is an extremely important topic. Far more work is being done on this area in Europe than in the U.S. The sorts of automatic reasoning that are to be supported by such representations are just beginning to be defined. The next few years will present an opportunity for the European community. With the number of procedural knowledge representation projects already large and increasing, there are sure to be some excellent insights regarding what sorts of inference one might provide with an automated system.

Knowledge-based systems—

A number of expert or knowledge-based systems were seen by or described to the NASA delegation: SACV, MWS, Arianexpert, CAPS ES, EOA, Radiometry ES, Mission Analysis Assistant, ESIOD, etc. Rather than examine each of these in turn to gain insight, two only will be examined: Arianexpert and SACV.

Arianexpert is a mature application to support post-flight analysis which has been tested, become operational, and is designed to be extensible by its users. The list of capabilities in Arianexpert reads like a textbook inventory of expert system functionality: an explanation facility, a report generator, capabilities for recording anomalies and experience, along with the knowledge-based core of analysis support functions for data selection, discrepancy detection and interactive procedure execution. Arianexpert is an expert system for expert system aficionados.

The Satellite Autonomy Concept Validation (SACV) project was a forward-looking study and prototyping effort to investigate the concept of interactive intelligent systems to achieve full onboard automation. Scenarios for onboard fault detection, diagnosis, recovery, and rescheduling were examined. Among the component prototype systems used in the study were the CONNEX diagnosis system and the MARS scheduling system. Although follow-on work to this study has not been initiated, the concept corresponds to one which has also been identified by NASA as a long-term automation goal, but which is not being actively addressed within NASA at the current time.

AI programming techniques—

Our hosts at CRI remarked how AI practitioners typically do not get proper credit for certain programming techniques invented within the field: object-oriented programming, blackboard architecture, procedural attachment, even search techniques. This is a variation on the more frustrating Catch-22 whereby AI practitioners are denied success as they succeed: Once the information processing details of an intelligent system are explained, the behavior exhibited by that system, considered “intelligent” formerly, does not retain the status of “intelligence,” apparently and merely by virtue of being explainable.

Nonetheless, there have been successful, readily acknowledged applications of AI programming techniques within ESA: The Noticeboard tool developed by Vega is a blackboard architecture with a demon facility. The CaeCALX application, developed by two Spanish engineering teams for ESTEC, uses best-first beam search and a careful understanding of constraints and search to avoid backtracking and solve a geometric resource allocation problem in storing cargo items for launch.

ESA also has made extensive use of expert system development shells. Among those mentioned in our visit are ART, CLIPS, KEE, KOOL, Milord, Ops83, and ProKAPPA.

Other AI technology areas—

There are a few other AI technology areas which we did not see a great deal of work in, but which merit some remarks.

At ESOC, there were some early studies being conducted on the applicability of neural networks to regression and data classification problems in scientific data analysis, such as wind vector determination and cloud classification from earth observing satellite data.

Matra has three efforts in design automation: SWITCHWORKS, Payload Expert, and System A. These tools are targeted for assisting designers in tasks such as redundancy analysis, critical component analysis, constraint satisfaction, and exploration of design alternatives. The general issue of system model reuse from the design stage through the system life cycle is also receiving attention.

The area of knowledge acquisition, validation, and maintenance is of considerable concern within ESA, particularly at ESOC where extensive expert system development has taken place. Matra is developing an interview tool called MACAO which assists the knowledge engineer. CRI is developing a tool called VALID for knowledge validation. Part of the approach is a generic knowledge representation for which consistency checking tools are being developed, both for base-level and control knowledge.

The area of Intelligent Computer-Assisted Training (ICAT) is being explored at CRI in a project called ITSIE (funded from ESPRIT). The developers are aware of concepts for both flexible teaching and testing strategies and for user modeling. At this time, the overhead for constructing a domain model detailed enough to be useful is considered prohibitive.

Natural language processing concepts are being applied in a project at CRI called SIMPR. The goal is to automatically create indexes for textual documents. The techniques being investigated are purely syntactic, meant to create indexes at the phrase level, not just at the keyword level. It appears that somewhat arbitrary and non-useful indexes may result without the application of some semantic information.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE October 1993	3. REPORT TYPE AND DATES COVERED Technical Memorandum	
4. TITLE AND SUBTITLE A Review of European Applications of Artificial Intelligence to Space		5. FUNDING NUMBERS 586-02-21	
6. AUTHOR(S) Mark Drummond and Helen Stewart, Eds.			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Ames Research Center Moffett Field, CA 94035-1000		8. PERFORMING ORGANIZATION REPORT NUMBER A-93127	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA TM-108789	
11. SUPPLEMENTARY NOTES Point of Contact: Mark Drummond, Ames Research Center, MS 269-2, Moffett Field, CA 94035-1000; (415) 604-4710			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified – Unlimited Subject Category 61		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The purpose of this report is to describe the applications of Artificial Intelligence (AI) to the European Space program that are being developed or have been developed. This report describes the results of a study sponsored by the Artificial Intelligence Research and Development program of NASA's Office of Advanced Concepts and Technology (OACT). The report is divided into two sections. The first consists of site reports, which are descriptions of the AI applications we saw at each place we visited. The second section consists of two summaries which synthesize the information in the site reports by organizing this information in two different ways. The first organizes the material in terms of the type of application, e.g., data analysis, planning and scheduling, and procedure management. The second organizes the material in terms of the component technologies of Artificial Intelligence which the applications used, e.g., knowledge based systems, model based reasoning, procedural reasoning, etc.			
14. SUBJECT TERMS Artificial intelligence, European space operations, Spacecraft automation		15. NUMBER OF PAGES 76	
		16. PRICE CODE A05	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT