SELECTED POPLOG APPLICATIONS

Application	Market Sector	Technologies
Design for Testability	Electronics	Design Rules Object Oriented Representation
ES for Road Improvement	Transport	Pattern Recognition Data Analysis
EXCAP	Engineering	Process Planning Prolog
RESCU	Research	Process Control Real-time KBS
COGSYS	Manufacturing	Process Control Real-time KBS
ACKT CASE Tool	Software	CASE
Interactive Vision Environment	Reseach	Image Processing
SPADE	Software	Theorem Proving Prolog
Traffic Information Collator	Transport	Natural Language Processing
SPRIDES - Spring design ES	Aerospace	POP-11
Design to Product	Manufacturing	Object Oriented Representation
Speech Synthesis	Research	Natural Language Processing
Tracking Polyhedral Objects	Research	Image Processing
EDWIN	Finance	Pattern Recognition
Currency Movement Analyst	Finance	Pattern Recognition
FAUST Fault Analysis Using Simulation and Testing	Electrical Utility	Object Oriented Representation Diagnosis
TASKER	Oil Exploration	Scheduling
CLORIS	Training	Intelligent Tutoring Interactive Video
Identifying Weld Defects	Engineering/Manufacturing	Intelligent Tutoring Interactive Video
Aircraft Maintenance Scheduling	Aerospace	Scheduling

Application	Market Sector	Technologies
Classification of Plankton in Neural Networks	Marine Biology	Pattern Recognition
GRIPE A Graphical Interface for Protein Exploration	Molecular Biology	Search
HACM Helicopter Air Combat Modelling	Defence	Simulation
RAP Resource Allocation Prototype	Defence	Command and Control
PANTECHNICON	Intelligent Documentation	POP-11
Automating Electrical FMEA	Car Manufacturing	Mixed Languages
Polytechnic Timetabling System	Education	Scheduling Prolog Genetic Algorithms
Ada Compiler Register Allocation	Software	Prolog
A Generalised Approach to Updating Scanner Strategies	Medicine	Expert Systems Prolog
An Expertext System for Building Standards	Regulatory Authority	Expertext
Interactive Materials Requirement Planning	Manufacturing	Planning
Land Segmentation	Aerospace	Image Processing
Information Fusion on the Battlefield	Defence	Object Oriented Representation
Optimising the Placement Sequence of Surface Mounted Devices on Printed Circuit Boards	Manufacturing	Rule Induction
The Writer's Assistant	General	Intelligent Authoring
PLINTH: An Expertext Shell	Regulatory Authority	Expertext
Model-based Understanding of Traffic Scenes	Transport	Image Processing

The Applications described in this booklet are examples of work done with Poplog. They are not necessarily commercially available products.

These applications synopses are based on published material supplemented in some cases by discussion with the authors. Applications range from systems in daily use to small scale research investigations or concept-proving studies. ISL is grateful to the Poplog users who have supplied this material. Responsibility for any errors in this document rests with ISL.

Integral Solutions Limited, Berk House, Basing View, Basingstoke, Hampshire, UK
Tel: 01256 55899 Fax: 01256 63467 Email: isl@isl.co.uk
Registered in England No. 2355560

SELECTED POPLOG APPLICATIONS

DTA: Design for Testability

ELECTRONICS
DESIGN

Design for Testability Advisor (DTA) is a program which gives advice on the testability of electronic equipment designs. RAF technicians observed that many printed circuit boards were difficult and hence expensive to test while others designs were inherently testable. DTA was developed to reduce the servicing costs of electronic equipment, by providing design engineers with pro-active advice about testability requirements. DTA works by extracting data from Racal-Redac's VISULA CAD system as a board is being designed, and building in Poplog an object oriented model of the circuit. Testability rules based on knowledge elicited from expert RAF test engineers are used to analyse the model. DTA then states where there are potential testing problems in equipment designs, gives explanations, and recommends solutions, with preferences. Although methods had already been found of producing more easily testable designs, their application depended on the awareness of the designer. Poplog made it possible to integrate expert systems with existing CAD tools, and provide design engineers with access to the expertise of the best test engineers.

Developed by SD-Scicon for the RAF.

ES for Road Improvement

TRANSPORT
PATTERN RECOGNITION

The Expert System for Road Improvement was developed to analyse the accident history of traffic "blackspots" and recommend road improvements. Whenever a road accident occurs a diagrammatic report is produced and information about the state of the road is noted as standard procedure. When a particular location is the site of many accidents, the relevant reports can be fed into the Expert System, which uses deductive analysis based on probability theory to examine the pattern of accidents and suggest which of a range of possible measures will most improve road safety at that site.

The Expert System for Road Improvement was developed by the Transport Studies Group of University College London, and tested using data provided by Essex County engineers.

EXCAP: Process Planning

ENGINEERING
PROCESS PLANNING

EXCAP is a knowledge-based system which automates process planning, and thus integrates the Computer Aided Design and Numerical Control machining of turned components. Process planning has proved difficult to computerize using conventional techniques, and can be a lengthy procedure for skilled operator. Working from the ISO-standard IGES description of a 2D CAD model, EXCAP produces a sequence of machining operations in reverse order by a process of "adding metal". Playback then allows the operator to watch a graphical simulation of machining operations in correct order, using Poplog's windowing facilities. The operator can interrogate EXCAP about the process plan. A further module converts the process plan to Numeric Control code to drive lathes and other machine tools.

Poplog's ability to combine rule-based and procedural programming, to support sophisticated windowing interfaces, and to integrate with existing technology, were vital to the development of EXCAP.

Based on research at UMIST, EXCAP is now being enhanced by a major UK engineering company which wishes to remain anonymous.

RESCU: Real-time Expert Systems Club of Users **RESEARCH**
PROCESS CONTROL

RESCU, a prototype expert system for process control, was one of Europe's first operational real-time expert systems. It was developed by SD-Scicon under the Alvey Real-Time Expert Systems Club of Users project, and installed at an ICI ethoxylate plant running on a VAX 11/730 interfaced to the Foxborough process control computer. RESCU consisted of a number of communicating processes including the operator interface, plant communication interface, inference engine, batch/blend, and it monitored the operation of the plant in order to advise operators on quality control. RESCU developers praised Poplog's efficient and compact code which allowed a multi-process, blackboard system to operate to plant timescales, on a VAX 11/730. During operational running over six months RESCU provided advice to the operators at a level approaching that of the most skilled plant supervisor. Following the success of RESCU, a consortium of 35 international companies went on to develop the RESCU technology into a general purpose tool for process monitoring and control: see COGSYS.

The Club consists of twenty-three industrial partners, who wanted to investigate the feasibility of deploying a process expert system, within an operational process environment.

COGSYS: Process Monitoring/Control **MANUFACTURING**
PROCESS CONTROL

COGSYS is a fully engineered real-time expert system for process control which was commissioned by a group of 35 international industrial companies following the success of the Alvey RESCU demonstrator. It provides a knowledge-based monitoring and control system capable of coping with the increasingly complex demands of modern industrial plants and processes. Designed as a generic system for developing a wide range of real-time process industry applications, COGSYS has already been used in live applications at British Gas and Metal Box. COGSYS utilises a process industry oriented Knowledge Representation Language built on Poplog, and benefits from Poplog's interactive development facilities. The COGSYS Generator produces data and rules which are executed by the run-time system, a compact C executive which interfaces directly to plant control and SCADA systems.

COGSYS Limited was formed in 1990 by British Gas, SD-Scicon and Salford University Business Enterprises Limited to market the COGSYS package.

ACKT: CASE Tool **SOFTWARE**
CASE

ACKT is a Computer Aided Software Engineering (CASE) tool which is designed to support the government standard software development methodology SSADM while providing intelligent support to the end-user.

The design and implementation of a computer data processing system can be error-prone and laborious. Structured Systems Analysis and Design Methodology (SSADM) was developed to formalize this process and is now a government standard. ACKT automates the SSADM methodology. In addition to providing intelligent diagram drawing facilities, ACKT offers further assistance to the systems developer by checking the logical consistency of the diagrams produced and reporting on the deficiencies of the design.

ACKT is based on Poplog and was developed byAppin Software. It will be commercially available in the near future.

IVE: Interactive Vision Environment **RESEARCH**

IMAGE PROCESSING

IVE is a Poplog-based image processing workbench, allowing users to experiment with a variety of image processing routines. IVE aims to assist in development of vision systems which "understand" an image with sufficient precision and reliability for industrial applications. By making it possible to experiment on sub-regions of images, IVE reduces the effort required to find the best algorithms and parameters for image processing. Having discovered, by experimentation, a suitable combination of low-level image processing routines, together with appropriate parameter settings, IVE allows the user to capture the strategy (a network of image processing routines) so that it can be applied to a new image as one operation. IVE is an object oriented implementation using POP-11 Flavours, and exploits Poplog's ability to link high-performance image processing routines written in C, Fortran etc. Using Poplog, IVE has also been linked to array processors for accelerating image processing.

IVE was developed at Reading University and follows on from research carried out on an Alvey project in collaboration with British Aerospace, Marconi, and RSRE.

SPADE: Proof Checker/Validator **SOFTWARE**

THEOREM PROVING

The SPADE interactive Proof Checker assistant was developed to provide formal validation and verification of high-integrity software. It can be applied to software written in a variety of programming languages: Pascal, 68000 assembler, a provable Ada subset ("SPARK"). The Proof Checker polices proof attempts, preventing the user from making incorrect deductions or manipulations, and it assists the user in searching for a proof by applying a large library of replacement and inference rules using pattern matching techniques. Systematic validation by SPADE during a program implementation also significantly reduces the burden of subsequent testing and maintenance.

Poplog was selected for development of the Proof Checker because of its conformity to Prolog standards and its rigour in handling integers of arbitrary length.

The Proof Checker was developed by Program Validation Limited, based on research work at RSRE and Southampton University, and has been sold commercially since 1987. SPADE has been used to verify safety critical software in finance, defence, power, security, and automotive engineering.

TIC: Traffic Information Collator **TRANSPORT**

NATURAL LANGUAGE

TIC aims to improve management of traffic flow following a road-traffic incident by formulating up-to-the-minute advisory messages about road and traffic conditions. The input to TIC consists of messages about road traffic incidents which are transmitted back to headquarters by the police at the scene of an incident, as part of their standard incident logging procedures. Such messages are generally fragmentary but in a jargon which is fairly consistent across each police force. TIC uses Natural Language Processing in Poplog to analyse these messages and extract information such as location of incident, whether road is closed, estimated duration etc., which is pertinent to other road users. This information is then used to formulate messages suitable for broadcasting to motorists by local radio. Ultimately TIC will broadcast to autoguide beacons and motorists' on-board computers.

TIC was developed by Racal and Sussex University in association with the Roads Research Laboratory and Sussex Police Force. Current developments also involve the AA and IBA as information broadcasters.

SPRIDES: Spring Design **AEROSPACE**
POP-11

SPRIDES was an Expert System developed to assist designers of helical coil springs. Although standard formulae and processes exist to aid this process, in practice there are many interacting factors and constraints which make necessary several "trial and error" attempts before a satisfactory solution is found. By encapsulating the knowledge of spring designers SPRIDES made it possible to reduce average design time from 1.5 days to 30 minutes. The design terminal's driver module, written in Pascal, could be interfaced directly to Poplog, and the POP-11 language and VED text editor were used to program the application. Poplog's incremental compilation and interactive programming features contributed to speedy development.

SPRIDES was developed by Smiths Industries in a period of 4 man months.

DtoP: Design to Product **MANUFACTURING**
OBJECT ORIENTED PROGRAMMING

DtoP was developed to demonstrate the deployment of AI in all aspects of manufacturing, from initial conceptualisation through to production and maintenance. Traditionally procedures and operations in design and manufacture of engineered products are organised sequentially. Much of the information generated about a product may be lost as it proceeds through the manufacturing process, yet it may later be necessary to recreate this information. DtoP provides an integrated product description, which maintains consistency throughout the total product lifecycle, linking the "islands of automation". Poplog proved particularly apposite as it enabled each aspect of DtoP to be implemented in the most appropriate way: Poplog's Common Lisp, POP-11 and Prolog subsystems were all used, as well as C and Fortran.

DtoP was a large scale Alvey demonstrator developed over 5 years, with 200 man years of effort and £9m funding. The partners involved were GEC (Electrical Projects), Lucas CAV, and Edinburgh and Loughborough Universities. The Lucas factory in Gillingham was taken over for two weeks of demonstrations in February 1990.

Speech Synthesis **RESEARCH**
NATURAL LANGUAGE PROCESSING

This research project aimed to produce natural-sounding synthetic speech from an input of written text. This has hitherto proved an elusive goal. The task has typically been seen as one of engineering rather than knowledge representation, and the approach of most text-to-speech systems has been procedural and rule-based. Using the logic programming language Prolog on a MicroVAX 3400 in the Poplog programming environment it was possible to take a novel non-procedural declarative approach to knowledge representation in speech synthesis. The Unification Grammar formalism, which is now widely employed in natural language processing systems, can be adopted making use of node-labelled directed acyclic graphs as the data-structure for the representation of phonological words and phrases. The acoustic interpretation of such structures is determined by a simple constraint-satisfaction technique rather than a laboured set of destructive transformational rules. This approach avoids many of the practical problems encountered in conventional text-to-speech systems, and unlike most other systems yields high-quality synthetic speech.

Developed by York University for a major electronics company.

Recognising and Tracking Polyhedral Objects **RESEARCH**
IMAGE PROCESSING

This application addresses the problem of recognising a known 3D object (modelled as a polyhedron) in a monochrome image, and tracking it through a further sequence of images. The algorithms are efficient, and results obtained with natural sequences demonstrate that they are robust, even when the object is partly obscured.

In order to find the object, and continue to track it, the system integrates “bottom-up” and “top-down” processing in order to exploit the benefits of both.The object is located in the first frame using Model-Based Search. Image features are then tracked into the next frame using Optic Flow techniques, and their disparities are used to invert the Perspective Transform. The system can detect when object tracking is becoming inaccurate; it then recaptures its position through another model search that uses the reduced disparity information, and a “best-guess” at position, to constrain the size of the search space.Thus the system is highly predictive when it has good information concerning object position, but resorts to non-predictive strategies when it does not..

These ideas have been implemented as a suite of programs written in POP-11 which combine novel algorithms with established ones. A new method of matching 3D models against 2D data issued that would be suitable for implementation on a massively parallel architecture. The system also implements an algorithm for tracking objects between closely-spaced frames and it integrates the model-based search with tracking routines.

EDWIN: Financial Prediction **FINANCE**
PATTERN RECOGNITION

EDWIN is an expert system which provides advice about buy/sell decisions in the currency market. This combination defeated attempts by HP Researchers to develop EDWIN using the leading knowledge engineering environments. The developers turned to Poplog, which proved to have the capacity to cope with the large amounts of data involved and the need to integrate different paradigms within one environment. It uses chartist techniques to analyse recent currency-rate fluctuations and hence predict future currency movements. The problem of matching charts to the data involves combining detailed statistical processing with a knowledge-based approach. It is also necessary to handle multiple parallel hypotheses.In test cases using historical data EDWIN outperformed human chartist experts. EDWIN was developed by Hewlett Packard with chartist expertise from Barclays de Zoete Wedd.

Currency Movement Analyst **FINANCE**
PATTERN MATCHING

Foreign exchange dealers analyse trends in currency exchange rates. To maximise trading profits and manage risk effectively, forecasts of future trends must be more accurate and more timely than those of the competition. The Foreign Exchange operation of the First Interstate Bank of California (FICAL) stored information about the exchange rates between five major currencies, measured every five minutes over a period of eight years. They wished to use this information to model movements in exchange rates. Using POP-11 FICAL’s research group built a temporal spreadsheet with up to a million lines of data: this is possible because the size of Poplog applications is limited only by the capacity of the host machine. The information can be accessed quickly, and the flexibility of Poplog makes it possible to compare movements between a number of currencies at the same time. Currently FICAL’s research group presents results from the system in a daily newsletter for the bank’s traders, but dealers will soon have instant access to the system on the dealing room floor. The same system has also been used in options trading.

Developed with three man months’ effort, FICAL’s system has been in use in daily trading since October 1990.

FAUST: On-line Fault Finding

ELECTRICAL UTILITY
OBJECT-ORIENTED REPRESENTATION
DIAGNOSIS

FAUST provides on-line diagnosis of faults in electricity supply on the National Grid. FAUST uses an object-oriented representation of the grid which is written in POP-11 flavours (POP++). As input data FAUST reads in real-time the telemetry information concerning equipment status changes (including faults and automatic switching actions) which is received at the regional grid control rooms as standard. FAUST then hypothesises where possible faults lie, and uses a simulator to predict confirmatory telemetry. FAUST is able to cope with partial data by formulating a number of hypotheses and adjusting their probability according to the amount of confirmatory data. FAUST can also cope with the large volumes of data which arise during, for example, adverse weather conditions, when very high fault rates (up to 400 faults/minute) due to lightning strikes, falling trees etc. make it difficult to identify individual faults. In live monitoring FAUST has operated successfully in real time.

FAUST was developed in Poplog by ERDC Capenhurst, in association with Thames Polytechnic.

TASKER

OIL EXPLORATION
SCHEDULING

TASKER plans and schedules underwater maintenance operations for North Sea oil platforms

TASKER is provided with details of each platform to be covered in an inspection programme, and of the support vessels available including their capabilities for different types of diving. Taking into account other activities going on at the same time on the platforms, TASKER prepares a plan for carrying out a season's workscope of tasks, each task representing a single operation on one platform. The planning is in two stages: first assembling tasks into 'task groups' which can be carried out together, because they are, for example, on the same face of the same platform and have the same diving requirements; and then scheduling the task groups to produce a complete plan. As a highly interactive support tool, TASKER can be used as much or as little as the operator wishes. Thus an operator may choose to amend a plan generated by TASKER, or TASKER may complete a plan started by an operator, or re-plan, following progress reporting. TASKER can also analyse the schedule and report on vessel costs and utilisation.

Developed using POP-11, TASKER is able to manipulate the complex data involved easily and efficiently. It uses the Poplog window manager (PWM) to provide a sophisticated graphical interface, with multiple views of plans and form-based data browsing. Evaluation of the system has shown that workspace inspection schedules can be developed more rapidly and more reliably than with conventional manual approaches. TASKER was developed by the CAD Centre, University of Strathclyde, in conjunction with BP Limited and System Designers Limited.

CLORIS

VIDEO-BASED INTELLIGENT TUTORING SYSTEM
TRAINING

CLORIS was developed to enable trainees to make more active use of video sequences designed to demonstrate a specific skill. It is often instructive for trainees to watch an expert performing a task which they need to master, but in real life it is not always possible for them to interrupt. For example it would be unwise to delay an emergency operation while medical students asked questions. Trainees may be reluctant to raise queries "on-line", but if questions are not raised when relevant, then misunderstandings may multiply and it may not be possible for a tutor to explain the procedure satisfactorily in subsequent discussion. Interactive video enables trainees to observe an expert, and interrupt whenever they need clarification, either about events in progress and objects on the screen, or when they want to refer back to other related objects or processes. CLORIS, written in Poplog/POP-11, controls a videodisc containing still frames. Any frame can be accessed and displayed, and groups of frames can be displayed as moving films. Generated graphics and text can be overlayed onto the video image. Trainees' queries are answered by generated natural language explanations.

CLORIS was developed at Lancaster University.

Identifying Weld Defects **ENGINEERING/MANUFACTURING**
NON-DESTRUCTIVE TESTING

Non-destructive testing of materials usually requires significant use of human expertise to interpret results. While it may be technically possible, say, to scan every weld in an oil pipeline ultrasonically, the vast amount of output data produced would require extremely large investment of time by expensive and possibly rare experts. A computer based interpretation of the output data, even if it only detected the areas worthy of expert inspection, could make the process far more efficient and cost effective.

A Poplog-based system, built by the Open University's KBSEng (Knowledge Based Systems in Engineering) group goes far beyond this. Using a blackboard architecture implemented in POP-11 ("ARBS" - Algorithmic and Rule-based Blackboard System), their system inputs raw data from ultrasonic scans, finds low-level features and identifies potential defect areas. Then, using a "hypothesise and test" strategy, it goes on to work out the likely nature of the defect. Poplog's flexibility made it a ideal vehicle for supporting multiple paradigms within a blackboard framework.

The demonstrator of the system runs on data supplied by the UK Atomic Energy Authority.

Aircraft Maintenance Scheduling **AEROSPACE**
SCHEDULING

Airworthiness requirements set down fixed intervals for routine checks and maintenance of aircraft systems. The intervals are based on reliability data. Hence for a given aircraft requiring maintenance of many different systems there can be many different maintenance intervals.

So that airlines can plan their work load in advance maintenance checks are grouped together into tasks that should be carried out after a fixed number of flying hours. These tasks should be arranged to require about the same time to complete so that maintenance staff are not faced with erratic variations in work load. Generally checks must be carried out at or before the required time. However, there is a limit on the degree to which checks may be moved forward apart from the first occurrence of a check.

BAe supplies customers with a maintenance schedule which is customized to the airlines' operations and equipment fit. Formerly such schedules were produced by hand. Now, a computerised scheduler written in Poplog Prolog is providing improved schedules. This system provides airlines with maintenance tasks that are approximately equal in work content, minimises the number of aircraft zones accessed during a task and ensures that checks are scheduled for the appropriate intervals. Introduction of the system has reduced the time taken to produce a schedule by more than an order of magnitude.

The aircraft maintenance scheduling system was produced by British Aerospace Avionics Computer Systems and has been in use since early 1991.

Classification of Plankton in Neural Networks **MARINE BIOLOGY**
PATTERN RECOGNITION

Marine biologists wished to classify groups of plankton which may be causing increasing incidence of paralytic shellfish poisoning in coastal waters around Europe. Traditionally plankton analysis is carried out by highly trained analysts using a microscope with a magnification of around 300 times. Skilled analysts are scarce and it is becoming necessary to automate techniques for classifying plankton. This type of pattern recognition task has proved difficult to automate using traditional techniques relying on knowledge elicitation, as experts may be unable to give adequate formal descriptions of the patterns they use for analysis. Neural networks are well suited to the task, as they are not programmed but "taught": they are exposed to "training" data and learn by being "told" the appropriate responses to the different inputs. The neural network can then apply this learning to respond to further data.

Using Poplog, Plymouth Marine Laboratory built a neural network based system which can distinguish between images of two species of plankton with a similar degree of accuracy to that of human experts.

GRIPE: A Graphical Interface for Protein Exploration **MOLECULAR BIOLOGY**
SEARCH

Analysis of the topological features of proteins has traditionally been performed by human experts inspecting diagrams of structures. However with the increase in the number of known protein structures it has become difficult to maintain a comprehensive list of topological motifs. Moreover an increasing number of workers is now examining protein structure. Using Poplog Prolog a very large database of facts about protein structures was established which can be queried to establish the presence of a particular topological motif - such as "Greek keys" - in a protein. The interactive graphical interface GRIPE, implemented in POP-11 and Prolog, allows the user to construct graphical queries about the linear and topological structures of selected proteins. GRIPE also provides facilities for viewing the three dimensional and topological structure of the proteins. The interface provides an easy and effective way for molecular biologists to examine protein structure, without first needing lengthy training in the use of the system.

GRIPE was developed by the Imperial Cancer Research Fund and the Department of Crystallography at Birkbeck College.

HACM: Helicopter Air Combat Modelling **DEFENCE**
SIMULATION

With the increasing use of helicopters on the battlefield, both military planners and helicopter designers needed to study the unique form of air combat that has resulted. The planners wished to develop strategies which would utilize helicopters' special capabilities to best effect, whilst the designers needed to develop performance characteristics best suited to the roles helicopters would be likely to play. Westland Systems Assessment Limited initially developed a Helicopter Air Combat Model written in Fortran, which enabled them to investigate the technological factors required for air combat. However this model was not suitable for investigating tactical variation, and did not permit any human intervention once the model was running. As it was clear that the model had potential for studying tactics development, the high level tactical decision code was re-written as a set of combat rules in Poplog Prolog. Using Poplog's external language interface, this Air Combat Shell was then interfaced to the main physical simulation. Users are now able to interrupt the simulation at decision points, and interrogate the current decision logic, or simulate the consequences of an alternative decision. They can also modify or add to the rule set, and again explore the consequences. HACM exploits one of Poplog's key design objectives, interactive programming.

HACM was developed by Westland System Assessment Limited.

RAP: Resource Allocation Prototype **DEFENCE**
COMMAND AND CONTROL

Success or failure of a naval mission often depends on a few key decisions taken rapidly after a significant event such as the detection of noise from an enemy submarine. The commander must quickly decide which of his available assets (sensors and weapon types) to deploy in response. Selection is based on factors such as current position, maximum speed, remaining fuel, and readiness for deployment. He therefore needs a timely and accurate picture of events, pre-planning of the resources available for any possible contingency, and real-time decision-making capability. The Admiralty Research Establishment recognised the need for effective automated support for command and control, and commissioned SD-Scicon to develop a prototype reactive resource allocation system. The aim of the prototype was to evaluate the techniques and establish user requirements. RAP was developed on a Sun Workstation, and interfaced to external processes which provided scenario data. Object oriented programming based on Poplog was used to provide the necessary factual and problem solving knowledge. The Poplog Window Manager was used to implement the user interface, which had to be fully interactive and provide fast and clear access to the system's knowledge, to give optimum decision support. By means of rapid prototyping the developers were able to understand users' requirements better, and modify the system accordingly.

RAP was developed by SD-Scicon and the Admiralty Research Establishment.

PANTECHNICON **INTELLIGENT DOCUMENTATION**
POP-11

Pantechnicon is an environment which supports the generation and production of documents. It enables the user to compose a document freely, using a combination of keystrokes and mouse actions. But it can cope with the constraints imposed on documents by professional requirements: engineering designs, mathematical formulae, computer programs, legal documents etc.

Pantechnicon is integrated with Poplog's text editor VED. Material in the edit buffer is parsed into an abstract syntax which is then further transformed into a typographic specification before being presented on the screen to the user. Inverse mapping from the typographic presentation on the screen to the cursor position in the VED buffer is provided. Currently a presenter for the LISP language is under development. LISP appears with scope delineated by boxes, and with standard mathematical typography for algebraic and transcendental functions. Other applications include presentations of technical material for engineering design.

Pantechnicon was produced at the University of Massachusetts at Amherst.

Automating Electrical FMEA **CAR MANUFACTURING**
MIXING LANGUAGES

Failure mode effects analysis (FMEA) involves the investigation and assessment of the effects of all possible failure modes on a system. FMEA is of growing importance in the automotive, aerospace and other advanced manufacturing industries, where increasingly complex electrical, electronic and mechanical systems are combined in safety-critical applications.

FMEA work should be carried out during the design stage to ensure that designs are analysed for hazardous and safety-critical situations. It is a tedious process demanding detailed and systematic examination of all aspects and parts of the design. Yet FMEA must be carried out by professional engineers because it requires extensive experience of the domain. These two factors, painstaking work and expert judgement, indicate the great benefits to design engineers of automated support for FMEA. The only commercial software available is a spreadsheet-style program to help keep track of the clerical side of the task.

Aberystwyth researchers used Poplog to model the cruise control system of a modern car. The system reasons from the structure of the electrical circuit, and from knowledge of the circuit's function. With it the design engineer can investigate the effects of all possible faults on an electrical subsystem (the most difficult task in the FMEA process).

The system was developed in Pop-11, but integration of languages within Poplog made it possible to include two existing pieces of software written in Common Lisp: a qualitative simulator capable of reasoning about circuit structure originally developed on Macintosh, and a functional reasoning system from Michigan State University together with its persistent object store written in Portable Common Loops.

Evaluation of the system has shown that it provides a basis for the construction of an automated assistant for the whole of the FMEA process. This will enable automotive engineers to specify the functionality of subsystems, and will incorporate rule-based and case-based reasoning.

The system was developed by Chris Price and John Hunt at the University of Wales, Aberystwyth working with Jaguar Cars, EDS-Scicon, and ICAD Ltd.

Polytechnic Timetabling System **EDUCATION**
GENETIC ALGORITHMS

The problem of producing timetables for large educational establishments has been approached in a variety of ways. There are many heuristic algorithms for assigning lessons and applying interchange operations. Mathematical methods like linear programming and integer programming have been attempted. However few have considered simultaneously all possible dimensions of the problem.

Genetic Algorithms have recently been reported as a new technique for solving the timetable problem. However where the search space is large and the problem highly constrained, Genetic Algorithms require very complex cost functions which are difficult to implement and evaluate.

Using mixed language programming within Poplog, Ngee Ann Polytechnic was able to use Genetic Algorithms more effectively. A Prolog assignment programme using a logic/knowledge-based assignment algorithm first sets up a major part of the timetable. Then Genetic Algorithms written in POP-11 continue to search for a complete solution within the established framework, which has already cut down much of the unpromising search space. Three genetic operations, crossover, mutation and filtering are used, while “elitism” ensures that good partial solutions are preserved in the next generation.

Developed by Si-Eng Ling of Ngee Ann Polytechnic.

Ada Compiler Register Allocation **SOFTWARE**
PROLOG

One of the critical problems facing the designer of highly optimising code generators for compilers is the optimum allocation of the computers’ registers. The problem is compounded when a compiler is required to be retargettable to a number of different computer architectures.
One technique is to define a virtual target machine which has a set of virtual registers. To complete code generation the virtual register usage must be mapped onto the actual registers of the target machine.

Systems Designers (now EDS-SCICON) developed one of the first validated retargettable cross-compilers for Ada. Actual target machines pose many constraints, for example register pairs are used to handle long arithmetic, and some registers are reserved for special purposes, as on Mil.Std. 1750 where R2 is used implicitly for some instructions. “Back-end” compiler designer, Chris Nettleton, used Poplog to develop a graph colouring algorithm to model the mapping of virtual to machine registers while observing all constraints. With Poplog’s incremental development it was simple to try out different register allocation strategies.

The efficiency of code generation of the Ada compiler was such that SD’s retargettable code generator was selected by Digital as the basis for their XD-Ada family of cross-compilers. “I don’t belive that we could have developed the compiler without the rapid-prototyping of the register allocation strategy in Poplog” reports Nettleton.

The Ada Compiler was developed by EDS-Scicon.

GAUSS: A Generalised Approach to Updating Scanner Strategies **MEDICINE**
EXPERT SYSTEMS

GAUSS is an Expert System designed to aid hospital radiographers in selecting parameters for magnetic resonance scans on hospital patients. As MR imaging is a relatively new field, still being actively researched, the criteria used in selecting scan sequences are constantly changing. The knowledge base may need to be changed, to correct errors which have been discovered, or to include new research findings, such as new disease types or new imaging techniques; or because individual medical experts using it have differing views of the best criteria for scan selection.

Thus the expert system must be structured in a way which allows the expert to make changes to the knowledge base whilst always leaving it correct and consistent. GAUSS is designed in two parts: the expert system which advises on scan selection, and the update program which permits changes to the knowledge base. The scope of possible modifications has been kept as wide as possible, whilst ensuring that potentially damaging changes are detected.

Changes can be made to the User Interface by modifying the Contents or grouping of menus, or to add criteria corresponding to new research. The knowledge base can be updated by modifying or adding new rules or facts. Checks are made, for example, that the rules at leaf nodes are attempting to produce an answer. GAUSS also allows new concepts to be added to the knowledge base. There are several ways of validating the modified knowledge base. The expert can query the system and check its justification for choice of a particular scan setting, while the output list enables him to check how his changes are affecting the reasoning process. And the updated system can be tested on a set of stored test cases where the correct results are known.

GAUSS was implemented on a Sun SPARCstation using Poplog Prolog. The system was developed by Greenwich University School of Computing and Information Technology, Plymouth University School of Information Science, and Guy's Hospital Division of Radiological Sciences.

An Expertext System for Building Standards **REGULATORY AUTHORITY**
EXPERTEXT

Acts of Parliament, codes and standards regulate numerous fields of human activity and continue to increase in scope and complexity. The Building Directorate of the Scottish Office is responsible for the drafting and maintenance of the Building Standards Regulations for Scotland and their associated Technical Standards which set out the statutory, technical requirements for governing building design and construction. The Technical Standards document references a range of other published codes and standards, and its authors must also take into account other legislation which may affect the particular issues addressed by the Regulations. The information in the Standards documents is subject to frequent review and revision.

A conventional expert systems approach to the development and management of Standards information has proved to be of limited value. The Building Directorate is currently engaged in a collaborative development project which seeks to combine the representational paradigms of both expert systems and hypertext, a combination which has been termed expertext. A system is being implemented based on the concept of headed record expertext in which formalised information can be attached to the nodes in a hypertext version of the Standards. The system handles navigation rules and provides intelligent guidance for the reader through the hyperdocuments. The system is intended to provide support for three information domains: the documentation itself; the knowledge accumulated by the development and use of the Standards, termed "argumentation"; and general knowledge about the Standards domain, termed "meta-knowledge".

The system is written entirely in POP-11 using the Poplog development environment and runs on SUN SPARCStations.

The Building Directorate's academic collaborators are the AI Department, Edinburgh University, the Department of Architecture, the Queen's University Belfast and the Department of Civil Engineering, Leeds University.

Interactive Materials Requirement Planning **MANUFACTURING**
PLANNING

Manufacturers plan ahead so that capacity can be established and materials ordered in time to meet demand. But as customers change orders, competitors launch initiatives, suppliers fail to deliver, and production lines break down, plans can easily become invalid. Planners and buyers try to cope with changing circumstances by changing the plans, but in an environment where their only tool is the fortnightly run of MRP, consequences such as material shortages can easily be overlooked.

The Knowledge Based Programming Department at HPLabs Bristol developed an interactive materials re-planning system, I-MRP, allowing production plans to be changed and evaluated in minutes. Potential material problems can be investigated and resolved interactively, and alternative plans explored. I-MRP users can maintain production plans that keep in step with the changing world.

I-MRP was developed in collaboration with HP's Computer Peripherals Division in Bristol, but the team set out to develop a system that could be quickly adapted to the needs of other users. The crucial issue was dealing with change - in user requirements. The approach was to exploit declarative programming techniques. Crucial parts of I-MRP, such as the equations that determine how many parts of each type will be needed in the months ahead, are represented explicitly in a declarative, spreadsheet-like language that can be easily modified, and then transformed to produce efficient code.

This approach proved its worth when the team came to re-apply I-MRP at HP's Vancouver Division (VCD). VCD were using HPJIT to determine materials requirements, involving a completely different material demand model. The declarative nature of I-MRP allowed the required modifications to be made in a few hours. Subsequently, the VCD installation of I-MRP underwent further changes as VCD moved to a structured bill-of-materials. All changes have gone smoothly.

Since February 1992 I-MRP has been responsible for all materials requirements computation at VCD, for both interactive re-planning and the regular MRP run. To provide continuing support for I-MRP, HPLabs and VCD established a novel technology transfer project in July 1992. The responsibility for I-MRP development and maintenance is being transferred from HPLabs to India Software Operation. A supported, product-quality, multi-site version of I-MRP will be available in February 1993.

I-MRP is implemented in a mixture of Prolog and POP-11. An application-specific language was designed and implemented in Prolog. POP-11 is used to store large data structure and to interface to the MOTIF user interface.

A System for Knowledge based segmentation of remotely-sensed images. **AEROSPACE**
IMAGE PROCESSING

Remote sensing platforms are providing an increasing wealth of data which can be invaluable for monitoring land use. The images produced need to be interpreted before information can be extracted. The image data must first be segmented into connected homogeneous regions, then it can be classified. Segmentation has proved difficult to automate, as there is often insufficient information in the images, and the task has been best performed by human experts. Under the ALVEY Man Machine Interface Project a consortium developed a knowledge based system for automated segmentation, written in POP-11 in the Poplog environment. A set of modules makes it possible to integrate the information contained in the image with existing maps, other sensor data, and the expertise of human analysts, to produce a region-based representation. The initial segmentation can be refined by a process which evaluates each region of the image separately and updates the overall segmentation.

The system was developed in a collaboration between Systems Designers, Sussex University, NERC Unit for Thematic Information Systems and CSATA.

Information Fusion on the Battlefield **DEFENCE**

OBJECT ORIENTED REPRESENTATION

With the advances in sensors and communications technologies, information sources become more detailed and their complexity more overwhelming. At the same time the increasing emphasis on rapid response in a more mobile battlefield will lead to information overload on the crewmen. Therefore a system is required to process and interpret the incoming information and present it in a more easily assimilated form to the crewmen.

EDS-Scicon developed a prototype Information Fusion System using Poplog and EDS-Scicon's proprietary software TOOLKIT on a Sun SPARCstation. Object-Oriented techniques were used to encapsulate knowledge about real world entities such as vehicle and sensor models, whilst the use of Knowledge-Based systems techniques enable expert knowledge to be used for the processing and interpretation of information. A fused picture of the current scenario can then be presented to the crewman, who has the capability to integrate all the hypotheses and the justifications for them.

The Information Fusion System utilises a scenario generator module, also implemented in Poplog and TOOLKIT, to simulate the real world and provide sensor reports to the system. These reports will simulate real life problems such as the effect of weather conditions, detection range, terrain features and communication delays on the fused picture. The system enables both the real world picture and the fused picture to be observed simultaneously during each reporting cycle for comparison.

The Information Fusion System was developed by ED-Scicon in conjunction with the Mission Management Techniques Section at DRA Chertsey.

Optimising the Placement Sequence of Surface Mounted Devices on Printed Circuit Boards

MANUFACTURING

RULE INDUCTION

Printed circuit boards (PCBs) are assembled by a placement machine which picks up the surface mounted devices (SMDs) and places them on the board, which is mounted on a table. Planning is complicated as it is essential to prevent potential collisions, and to minimise table travel time. The ability to shorten the time taken for placement is significant in enabling a company to meet production targets, improve the process yield, and so achieve a better return on investment in expensive pieces of machinery.

The assembly machines have a built-in controller with an embedded component placement algorithm. However it is not easy for the user to check its effectiveness, or to compare performance between different machines.

Researchers at Nanyang Technological University studied a range of machines and built a prototype optimiser for component placement on PCBs. This also enables them to supplement or validate the algorithms built into the controller.

Poplog was chosen because it is an integrated toolset which gives the researchers the ability to explore a number of approaches.

Rule induction was used to study production tactics employed in the industry and derive a satisfactory rule set. The placement sequencing rules can be embedded within a rule-based system and linked with programmed, company-specific, heuristics to generate an optimal cycle time. All coordinate points for different boards and machines can be stored in a database.

The work has been carried out at Nanyang Technological University Singapore, in collaboration with industrial users.

The Writer's Assistant

GENERAL
INTELLIGENT AUTHORING

The Writer's Assistant is an integrated writing environment for people who create complex documents - technical reports, newspaper articles, research papers - as part of their professional lives.

Existing tools for this purpose tend to be based on text editors to which various refinements and extra functions have been added, in an attempt to adapt the tool to the user. This may constrain the writer's process of working. The Writer's Assistant was designed to support the wide variety of strategies that writers actually use, and perhaps to offer them new ways of working.

Thus the first stage of the project was to study writers at work, and carry out a task analysis of the writing process - note taking, planning, revising, editing etc. The Writer's Assistant was then developed to an initial specification, which was refined as users were able to test and comment on the tool.

The Writer's Assistant provides three views of the emerging document; a linear view, which shows the flow of text from beginning to end, allowing the writer to perform standard editing operations; a structure view that allows the writer to create and manipulate a structural outline of the text; and a notes network view which provides a medium for the capture and manipulation of idea labels.

Each of these views is enhanced by a set of presentations to tailor the way the document is displayed. The structure view can be shown as a linear outline or as a tree diagram. Linear text may be given a fisheye presentation, where the text at the current focus of attention is shown in full and text further away is compressed to keywords or section headings, giving both detail and an overview of the document on the screen. The Writer can specify text types such as 'research paper/in-house style' or ‘company report’, which will provide a guide to the structure of the document.

The Writer's Assistant was developed using Poplog on a Sun workstation. Poplog, with its support for incremental development, was well suited to the application because it was not possible to specify the system fully in advance. Little was known about writers' strategies, and in any case the aim was to offer writers new ways of working. As the project developed writers' techniques became better understood, and users were able to provide feedback on features of the Writer's Assistant they found useful, or new features which might be desirable. Using Poplog it was possible to modify the functional specification of the system in the course of development.

The Writer’s Assistant was developed by British Telecom and Sussex University.

PLINTH: An Expertext Shell

REGULATORY AUTHORITY
EXPERTEXT

PLINTH (the platform for Intelligent Hypertext) is an expertext shell. It integrates the technologies of hypertext, rule-based expert systems and semantic nets to provide tools and intelligent support for the authors and readers of regulations and other technical documents.

PLINTH was originally designed to support the work of the Building Directorate of the Scottish Office on the Building Standards Regulations for Scotland and the associated Technical Standards.

Both hypertext and rule-based systems are well established individually as tools for work on regulations and standards. Hypertext systems provide structured access to technical documents, and can maintain and cross-reference large bodies of text, diagrams and tables. Rule-based expert systems exploit the inherently rule-oriented nature of regulations. However, both approaches have limitations.

- Hypertext has the known problems of disorientation and digression, where readers fail to maintain a mental picture of their location and path in large networks, and get side-trackedor lost.
- Expert system rules are limited in expressiveness compared with texts and diagrams.
- Conventional expert systems must reduce all knowledge to rules, whether or not this is appropriate.
- Creating large rule-bases for regulations, and maintaining the rules as the regulations are revised, is a time-consuming process requiring skilled programmers.

- Neither conventional hypertext nor rule-based systems are well suited to providing support for handling the "design rationale" underlying the technical documents. The ability to structure and manipulate this information easily is very important for authors.

PLINTH solves these problems by using 2 strategies.

Firstly, PLINTH augments hypertext networks with semantic properties. The author can assign types and slots to nodes and links to mark the function of the nodes in the network and the structural, logical and rhetorical relations between them.

PLINTH then supports this augmented hypertext with rule-based intelligent navigation, to interactively compute a customised path through the document for the reader, based on the semantic properties and directed by commands and queries. The author can write sets of navigation rules to perform different text retrieval or consultation functions for each view of the network.

PLINTH supports authors and readers of technical documents. It allows authors to: - construct semantically augmented hypertext networks for the documents and their associated design rationale; - write rule-based intelligent navigation and consultation functions in order to assist the reader in using the documents effectively and correctly.

PLINTH allows readers to: - display different structural, logical and rhetorical views of a document and manually select nodes for reading; - explore the design rationale behind a document to better understand its purpose; and - browse and consult documents aided by the intelligent navigation functions supplied by the authors.

PLINTH is implemented in POPLOG POP-11 on Sun SPARCstations, using the X Toolkit and the Sun Open Look widget set.

The PLINTH project at AIAI is currently funded jointly by the Science and Engineering Research Council (SERC Grant GR/J59197) and the Scottish Office.

Model-based Understanding of Traffic Scenes **TRANSPORT**
IMAGE PROCESSING

Traffic managers responsible for traffic surveillance and monitoring require answers to questions such as:

- What is the state of congestion at a junction?
- Why is the traffic behaviour abnormal?
- Are traffic regulations being broken?

Vision systems able to address such questions need to consider relationships between vehicles and the traffic scene as 3-D objects interacting in a 3-D world.

The University of Reading's research within the Perception Component of the VIEWS Esprit project concentrates specifically on methods for recovering a 3-D description of traffic scenes. The recovered description is suitable for use by the Situation Assessment Component of VIEWS which reasons about events and behaviours of vehicles of interest to the end-user of the system.

The system makes very extensive use of scene-specific "a priori" knowledge of:

- The traffic scene, and the position of the camera(s) within the scene.
- The types of vehicles expected within the scene.
- The dynamic behaviour of vehicles in the scene.
- The expected routes of vehicles in the scene.

Such knowledge is essential for two reasons:
(i) to define the specific requirements of a traffic understanding system
(ii) to make the tasks to be solved by the vision system tractable.

In particular, the ability to use the "ground-plane constraint" - i.e. the knowledge that vehicles will (normally) stand on the road surface - makes the task feasible with current vision technology.

A range of knowledge acquisition tools have been developed to allow application specific knowledge to be assembled by an engineer during the development phase of a system. These enable the generic capabilities developed in VIEWS to be tailored rapidly to specific end-user needs.

The methods developed at Reading give a dramatically improved performance of vehicle detection and tracking, which creates a full 3-D understanding of the traffic scene. The Figures (overleaf) show typical results, obtained at Newcastle Airport. The position and type of each vehicle has been instantiated from the camera view. Position of accuracy is typically within 1 metre, and directional accuracy within 5 degrees.

In particular it is easy to show the visual interpretation in plan view, and provide a representation in real world terms. This allows "high level" question about the scene to be answered by the Situation Assessment Component of VIEWS reasoning in 3-D.

The system was prototyped using Poplog on a Sun SPARCstation. Part of the system was ported into C code, but the control module remains in Poplog, with external linking to C and Fortran.

The work was carried out by Reading University as part of the Esprit Project P2152 VIEWS.

Integral Solutions Limited,
Berk House,
Basing View,
Basingstoke,
Hampshire.
RG21 4RG
U.K.
Telephone: +44 (0) 1256 55899
Fax: +44 (0) 1256 63467
E-Mail: isl@isl.co.uk

Contact ISL for details of Distributors outside the UK.

Form Ref: ISL0139a