This is to certify that the

dissertation entitled

IPCA: An Intelligent Control Architecture
Based on the Generic Task Approach to
Knowledge-Based Systems

presented by

David Bruce Decker

has been accepted towards fulfillment
of the requirements for

Ph.D. degree in Computer Science

Major professor

Date November 13, 1995

IPCA: An Intelligent Control Architecture Based on the Generic Task Approach to Knowledge-Based Systems

By

David Bruce Decker

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

Department of Computer Science

1995

ABSTRACT

IPCA: An Intelligent Control Architecture Based on the Generic Task Approach to Knowledge-Based Systems

By

David Bruce Decker

A knowledge-based architecture for the control of complex dynamic processes is presented. The control architecture is rooted in the Generic Task (GT) approach to knowledge-based systems. The architecture extends the GT-based DSPL framework for routine design and planning by providing a complete architecture supporting knowledge-based control system synthesis and execution.

The **Intelligent Process Control Architecture** (**IPCA**) incorporates three primary components: a **Blueprint Generator**; a **Plan Assembler**; and a real-time reactive **Plan Execution Monitor**. The Blueprint Generator produces a specification for a complex process control plan which is designed for a specific execution environment. The blueprint includes information about overall control policies and processing goals but not their specific algorithmic implementations. The Plan Assembler implements each blueprint by combining pre-compiled plan and control method fragments. The selected fragments are assembled to form an executable process control plan. Process control plans are intended for use with the Plan Execution Monitor for execution within the environment to be controlled. The Plan Execution Monitor incrementally determines appropriate fragments from the process control plan based on sensory input gathered from real-time process monitoring, reactively selecting an appropriate path through the process control plan structure. As a result of fragment execution, processing goals are modified and process parameters (control outputs) are changed, thereby facilitating control of the process.

The usefulness of the approach is demonstrated by applying IPCA to the control of a microwave-based composite material fabricator. Experimental results are presented for selected polymeric composite structures which were cured using IPCA. These experiments demonstrate that IPCA effectively controls the fabricator, resulting in high-quality, high-speed composites processing.

IPCA was also tested in conjunction with a simulated control system interface in a second series of experiments. These experiments were used to demonstrate intelligent response to various hypothetical conditions which are difficult or impossible to create in the laboratory. The experiments also demonstrated the breadth of knowledge-based domain coverage included in the IPCA prototype.

ACKNOWLEDGMENTS

I would like to acknowledge the contributions of several faculty from the Michigan State University Department of Computer Science: Dr. William F. Punch; Dr. Jon Sticklen; and Dr. George Stockman. Their guidance, support and encouragement were essential in allowing the completion of this work. I would also like to thank Dr. Martin C. Hawley of the MSU Department of Chemical Engineering for providing the domain which has enabled a real-world test environment for this control architecture. I would also like to thank Dr. James K. McDowell for his many contributions and advice.

Financial support was provided by the Composite Materials and Structutres Center at Michigan State University and by the MSU Department of Computer Science. A significant portion of the computer hardware used during this research was supplied by Apple™ Computer Corporation.

Several graduate students were helpful through discussions and providing tools enabling the implementation:, including Kurt Patzer, Ahmed Kamel and Robert Hawkins. A special thanks goes to Valerie Adegbite and Yunchang Qiu for constructing the experimental microwave cavity, implementing the conventional control system and assisting with many of the fabrication experiments.

I would like to thank my family for their love, support and understanding in this long and sometimes trying endevor: Mr. Nestor J. Decker Jr.; Mrs. Marilyn A. Decker; Mr.

Nestor (Tory) J. Decker III; and Mrs. Gertrude Shier.

Finally, I would like to acknowedge the contributions of two of my closest friends: John Bauman and Keith Pyne. Their friendship provided much of the support I needed while completing this work. Thanks also go to Keith for proofreading the final copy.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

1.1 Motivation and Problem Introduction

The problem addressed by this research involves producing control algorithms which are suited for controlling specific physical systems. Controlling real-world phenomena often demands the successful integration of knowledge-based planning techniques and numerical control methods. There are many approaches from which to choose. Which techniques and approaches might be appropriate for the control of a particular real-world system or process? How can selected planning and control algorithms be generated and combined to form an appropriate control system for a given process? How can reasonable processing goals be determined for a particular process? Can knowledge about the process and domain help to guide the selection of appropriate goals and control strategies?

Although portions of the overall task have been extensively studied in the fields of control systems theory, planning and knowledge-based design, most approaches concentrate on only one or two of these areas. Within the domain of modern control systems theory, for example, physical systems are often assumed to be fixed and well-defined so that they may be modeled mathematically [Gopal 93]. The use of such simplifying assumptions can lead directly to elegant and efficient control algorithms based on well-established development methodologies. But these systems are almost exclusively hand-crafted and they cannot plan a complex set of actions to accomplish high-level goals. Moreover, classical control systems design methodologies may fail altogether when one attempts to model poorly-understood processes [Boyd 91].

Planning systems, as opposed to traditional control models, are much more effective at the task of reformulating and pursuing goals and in reacting to unforeseen complications [Winston 92]. Planners are capable of creating and revising complex plans designed to accomplish a large set of potentially interacting or conflicting goals. Many planning systems can apply their plans in real-time [Musliner 95]. For example, planning systems used in robotics and autonomous vehicle research cope well with extremely complex environments [Firby 88]. But planning systems also have drawbacks. **Classical** (static) planners do not work well in dynamic environments since the state of the world can change independently of applying selected actions from a plan [Fikes 71]. As a result, much of the planning research is not useful in "real-world" environments. **Reactive** planning approaches attempt to address this problem, but delays encountered in formulating and fixing plans during their execution may slow such systems to an unacceptable degree. **Universal planners** create plans which are size-exponential in the number of sensed parameters because they directly map all possible world-states into appropriate actions and goal re-formulations [Schoppers 89].

Another limitation of many planning systems is that well-understood subproblems of an overall planning and control task cannot be effectively addressed by many control algorithms since the planner must reason about operator interactions and potential goal conflicts. Most importantly, reactive planners must be provided with goals (which may need to be revised as replanning occurs) but goals are usually treated as either static or some-how provided from outside the planner.

The forte of knowledge-based **design systems** is their ability to synthesize artifacts which can meet specifications and satisfy design constraints [Coyne 90]. For example, most design frameworks could support a hypothetical problem-solver which produced control system designs given appropriate knowledge, specifications and design constraints. Such design systems, however, are almost exclusively intended solely for the production of artifacts. Thus, typical design systems could support the task of producing a control system design, but the computational activity of constructing and running the resulting artifact would be seen as a problem exterior to that of the design framework.

It seems reasonable to posit that a successful approach to creating control algorithms for the management of dynamic physical processes might combine the strengths of modern control theory and knowledge-based design and planning. The architecture described here, the Intelligent Process Control Architecture (IPCA), is based upon precisely that combination of approaches.

1.2 Approach

The overall problem-solving task addressed by IPCA involves constructing executable control systems from declarative domain specifications. Task-specific frameworks are used to specify appropriate knowledge structures and control strategies for the major elements of the overall task. The frameworks used are based on the Generic Task (GT) approach to knowledge-based systems, a task-specific approach to knowledge-based system analysis and construction [Chandrasekaran 86]. The IPCA approach considers the overall problem-solving task of creating and applying appropriate control algorithms to be composed of three major subtasks:

- **Subtask 1 - Determination of suitable control system features and processing goals based on design specifications, constraints and domain-specific knowledge of the system to be controlled.**

- **Subtask 2 - Automated implementation of control system features and goals by incorporating selected pre-compiled planning and control algorithms into an overall process control plan.**

- **Subtask 3 - Reactive execution of process control plans in a real-time environment capable of supporting control algorithm selection and goal maintenance.**

It will be shown that **Subtask 1** is a routine design and planning activity and that it is a suitable candidate for implementation under the task-specific knowledge-based DSPL framework for routine design and planning [Brown 89]. The task involves the creation of an overall structure for the control system and specification of the overall processing parameters and goals to be achieved by the system.

Subtask 2 involves the selection and assembly of appropriate goals, plan fragments (pre-compiled static plans, heuristics and/or dynamic planning algorithms) and control method fragments (pre-implemented, tested control algorithms) to implement process control plans meeting the criteria generated by Subtask 1. Knowledge regarding the applicability of each fragment (the conditions under which the fragment may be useful) is included in the resulting process control plan implementation. A knowledge-based selection mechanism is utilized to provide a framework for supporting real-time fragment selection based on situated applicability. The sponsor-selector system for real-time fragment selection is described in Chapter 4.

Subtask 3 involves facilitating the execution of the resulting process control plan. This task involves providing mechanisms for real-time fragment selection, data abstraction, goal maintenance and fragment library access. Communication with sensors, actuators and external control systems is also supported.

Together, the subtasks combine knowledge about various planning/control strategies and when each is appropriate, knowledge about the process to be controlled, and knowledge about the composition and physical characteristics of the material to be processed. The task-specific frameworks provide knowledge representations and control strategies which are appropriate for each of the subtasks. The IPCA architecture integrates the three subtasks into a comprehensive problem-solving framework which provides automatic construction and execution of control systems based on declarative domain specifications.

A prototype system has been constructed based on the IPCA architecture. The prototype is based on a particular real-world domain, that of microwave-based composite material fabrication. The prototype has been tested under a variety of fabrication conditions to prove the viability and flexibility of the IPCA approach. The details of the architecture, the underlying information-processing theory, the prototype implementation, the selected domain, the experiments conducted and the results obtained are summarized in the chapters to follow.

1.3 Outline

Chapter 2 provides background material including the research approach, an overview of IPCA system architecture, expected results and evaluation criteria. The literature review and related work section discusses task-specific architectures (including an introduction to the Generic Task approach to knowledge-based systems), classical and reactive planning systems, gives a brief overview of control theory, and introduces real-time knowledge-based control systems. The strengths and uniqueness of the IPCA research approach and contributions to the areas of AI and composite materials processing (the domain) are summarized. The overall structure of the IPCA architecture is briefly outlined in the architectural overview. Finally, the expected results and criteria for evaluation are presented.

Chapter 3 presents domain background regarding composite materials and fabrication technologies. A survey of processing technologies is provided, with emphasis on a novel single-mode microwave fabrication technique pioneered at Michigan State University. Related work in the area of high-speed low-cost composite materials processing is presented to provide a domain context for the IPCA research. A prototype single-mode laboratory-scale cavity and the associated external control system are described.

Chapter 4 includes theoretical background necessary for introducing the IPCA architecture. A complexity analysis of the overall IPCA information-processing task is presented. A knowledge-level analysis is then offered, followed by task descriptions of the major elements which comprise the IPCA architecture. The following sections describe the Generic Task (GT) approach to knowledge-based systems, including descriptions of sponsor-selector mechanisms and an introduction to the DSPL language for routine design.

Chapter 5 contains a detailed description of the IPCA system architecture. The Blueprint Generator, Plan Assembler, Fragment Librarian and Plan Execution Monitor modules are described. The organization of knowledge within each module and the format of information exchanged between modules are specified. The information processing responsibilities and implementation of each module are discussed. Finally, the mechanism for data exchange between the Plan Execution Monitor and the External Control system is outlined.

Chapter 6 contains introductory and supporting material describing the experimental laboratory-scale microwave-based fabrication setup. The initial section details the experimental scope and objectives for each experimental series. A preliminary considerations section describes sample composition, geometry, orientation and process instrumentation common to all the fabrication experiments. The electromagnetic modes selected for processing are given, together with heating profiles and mode maps for each experimental series. The mechanical limitations of the single-mode cavity are summarized. The evaluation criteria for the experiments is detailed in the final section.

Chapter 7 includes experimental data and analysis taken from material fabrication runs in the laboratory. Summaries for each experimental series are given. A summary and evaluation of experimental results and analysis is also presented for each series. Extensions to the approach are discussed in terms of potential applicability to alternate composite materials processing technologies (including variable-frequency tuning) and within a broader context of controlling dynamic physical processes outside the domain of composite materials processing.

Chapter 8 includes conclusions and a section addressing future work.

Appendix A contains a glossary of technical terms used in the dissertation.

Appendix B includes a user-level guide for the IPCA implementation, providing information regarding the operation of all central IPCA software features, including user-interface descriptions for all software components.

Appendix C contains a brief summary of the software utilities used in conjunction with IPCA.

Appendix D contains a description of the External Controller and the protocol specification used to support network-based communications with external control systems, sensors and actuators.

Appendix E contains background information regarding the domain of composite material fabrication.

Appendix F contains a case study involving the addition of a new control strategy (specifically, a new type of mode-switching control fragment) to the IPCA system. Considerations for adding capabilities to IPCA are presented. The considerations surrounding Plan Generator maintenance and fragment library maintenance are investigated.

CHAPTER II

2.1 Related Work and Literature Review

The research presented here builds on the strengths of selected approaches from four related areas and combines them into a unique comprehensive control architecture. The work which is foundational to IPCA arises from:

- **Control systems theory**
- **Knowledge-based planning systems**
- **Task-specific knowledge-based system architectures**
- **Knowledge-based routine design**

Following a brief introduction which explains the manner in which the four foundational areas contribute to the IPCA architecture, each of the areas will be introduced in the order listed. A brief literature overview of each area will be presented.

Although modern control system development methodologies are effective, there are many unrealized opportunities for exploiting knowledge to assist in the design, implementation, testing and tuning of such systems [Boyd 91]. Some complex control systems, for example, are designed to support interaction between low-level control and higher-level supervisory planning systems [Arzen 89]. At present, the design, implementation, testing and tuning of such systems is largely ad hoc. IPCA demonstrates that advanced computing techniques can be used to provide tools to assist the designer with much of the control system design and implementation process. IPCA reduces development time, promotes reusability of control and planning algorithms, heuristics and approaches, and provides a framework to capture knowledge about planning and control.

The knowledge-based IPCA architecture utilizes experience gained from control theory by providing a standard set of implemented and tested **control method fragments**, available for potential use in a variety of control system applications. Similarly, **plan fragments** provide IPCA with a set of useful planning algorithms. The set of fragments is easily extended to allow new control approaches, planning algorithms, techniques and heuristics to be incorporated into IPCA. Fragments are intended for use and reuse in many different combinations and configurations. They are used to produce a variety of complex, effective algorithms for various controlled systems. By providing mechanisms which intelligently select and assemble appropriate fragments into complete control system solutions, IPCA assists the control engineering task by exploiting knowledge about control and planning.

Because the IPCA architecture builds on solutions provided by modern approaches to state-based control, control theory provides an important framework for a portion the architecture and forms the basis for the algorithmic content of the control method fragments. Control theory is therefore introduced in section 2.1.1. Similarly, knowledge-based planning has an important role to play in the creation of effective plan fragments within the IPCA architecture. Knowledge-based planning is therefore reviewed in section 2.1.2.

The knowledge-based problem-solving frameworks upon which IPCA is based are task-specific. Task-specific approaches to knowledge-based systems are therefore described in section 2.1.3. Finally, the contributions from the area of knowledge-based design are treated in section 2.1.4.

2.1.1 The Contributions of Control Systems Theory

Classical linear control systems development methodologies call for a careful analysis and characterization of a particular physical system, leading to an initial estimation of the aspects of the system which must be monitored and controlled [Newton 57]. Many different approaches to control may be considered during the development process, the mechanics of which may include the construction of Bode plots, root-locus computations and consideration of control processor characteristics such as whether proportional (P), proportional-integrative (PI), proportional-derivative (PD) or proportional-integral-derivative (PID) control laws might be most effective [Boyd 91]. Any such approach can be selected to form the basis for a control system architecture. Appropriate sensors and actuators are selected and, in some cases, a model of the system is developed which specifies outputs (actuators) as a function of inputs (sensors). The resulting model is sometimes based on ordinary or partial differential equations. A control system implementation is then developed and tested. The completed implementation is often tuned for optimal performance. Performance may be evaluated against a number of criteria including regulation, response, robustness, sensitivity and margin.

Although modern **state-based** linear control system design follows a similar development methodology [Chen 87], the form of the model (i.e. the control laws) may be very different. State-based control systems exploit the availability of digital control processors which are capable of implementing more sophisticated control laws than could older vacuum tube-based or transistor-based control processors of the classical control systems era. With modern control system designs, discrete control laws may approximate the continuous forms covered under classical linear control systems theory or they may be much more complex [Brogan 91]. Some of the more interesting possibilities include linear quadratic regulators, estimated state feedback controllers and linear quadratic Gaussian controllers [Astrom 90].

Advances in control theory continue to add a variety of new control approaches. Control approaches based on neural networks and fuzzy logic have appeared. Adaptive control algorithms have become an entire area of study. All these control approaches are useful in particular circumstances, as well as many other types of control algorithms which were not explicitly mentioned here. The import is that the IPCA architecture utilizes control theory as a tool to provide a useful and appropriate repertory of control algorithms from which process control plans may be assembled. Control algorithms employed in the current IPCA implementation include single and multiple setpoint control, state-based control and proportional control. These algorithms are contained in the control method fragments and, taken together, fragments implement the control laws required to control the processes of interest. IPCA is designed to be extensible so that new control algorithms may be easily added at any time.

2.1.2 The Contributions of Knowledge-based Planning

Just as control theory contributes to the construction of control method fragments, knowledge-based planning research contributes to the construction of IPCA's plan fragments. Unlike control method fragments, which focus on controlling process parameters directly, plan fragments control process strategies through the manipulation of goals. In order to be effective at the task of planning and reformulating process goals, the plan fragments may draw upon the wealth of research which has been conducted in knowledge-based planning.

As briefly introduced in chapter 1, there are many types of knowledge-based planning systems. All of them, however, share certain characteristics. All planning systems attempt to generate sets of actions which will satisfy goals, given an appropriate formal description of the world.

Static planners, such as STRIPS, provided an effective approach to solving many types of planning problems in static worlds [Fikes 71]. Such planners rest on the assumption that the world will not change unless acted upon by the planner, but, unfortunately, many real-world environments do not enjoy this property. Many early planning systems such as STRIPS assumed a linear (completely ordered) plan structure. With the development of NOAH, Earl Sacerdoti reformulated the problem of planning from world states to one of a search through **plan space** and allowed the generation and manipulation of partially ordered plans [Sacerdoti 74]. From this development, it became clear that plans could be partially ordered and that it was often advantageous to defer placing planning actions into a total order. Hence, many successful non-linear (partial order) planning approaches appeared [Weld 95].

These encouraging developments led planning researchers in many directions. Some developed **reactive planning** systems, which fixed plans only as they interacted with their respective real-world environments [Allen 90]. These systems became popular in robotics and autonomous vehicle research, since they could rapidly react to many unexpected conditions without the need for detailed or extensive preplanning [Bihari 89]. Some approaches were based on incremental planning, using least commitment strategies to plan only as far ahead as required [Dean 88]. Other approaches called for pre-computation of so-called **universal plans**, which function as pre-compiled reactive plan caches [Schoppers 87]. Many ideas developed out of the reactive planning community, including the somewhat controversial theory of **situated action** [Brooks 86].

The RAPS (Reactive Action PackageS) system allows control over complex dynamic processes by bundling reactive actions into "packages," each of which competes for resources until explicit goals are achieved. [Firby 88]. The RAPS system is similar to a parallel blackboard architecture. The system provides a facility for defining and maintaining goals (a blackboard), and each RAP (agent) competes for execution time in an attempt to satisfy specific goals. The execution of IPCA process control plans is similar, but IPCA does permit several actions to execute in parallel.

The Procedural Reasoning System (PRS) [Georgeff 87] is also similar to the reactive execution environment of IPCA. The PRS system is composed of a database of beliefs about the state of the world (using first-order predicate calculus), a set of goals to be achieved, a set of knowledge areas (plans, KAs) which are presently available, an intention structure containing those plans which have been chosen for potential execution, and an interpreter which manipulates all these components and selects KAs for execution. KAs consist of an invocation condition and a procedural body. Unlike IPCA, PRS allows plans to execute in parallel. Overall, the execution mechanisms of RAPS, PRS and IPCA are similar, although RAPS and PRS are somewhat more sophisticated. However, IPCA differs from both RAPS and PRS by providing an additional mechanism which supports the automatic generation of plans.

Advances in knowledge-based planning continue to produce a variety of new planning approaches. Each of the approaches described here is useful in particular circumstances, but there are also many other approaches to knowledge-based planning which were not explicitly mentioned. The import is that the IPCA architecture utilizes knowledge-based planning as a basis for providing a useful and appropriate repertory of planning algorithms from which process control plans may be assembled.

2.1.3 The Contributions of Task-Specific Architectures

Task-specific architectures (TSAs) adopt the position that knowledge representation and modeling should not be arbitrary "ad hoc" processes [Chandrasekaran 86]. Rather, TSA approaches to knowledge-based systems hypothesize the existence of commonalities between certain classes of tasks that can be exploited when performing analysis, modeling and knowledge acquisition. The commonalities or "shared properties" of the tasks must be such that they are meaningful across a large set of systems (i.e. they must be somewhat generic), yet specific enough to offer a reasonable exploitation of constraints during analysis and modeling (i.e. they must offer leverage during problem-solving). This "grain size" problem is one of the major ways in which common TSA approaches differ.

Grain size has been the focus of much debate, especially between groups working on different TSA approaches. But even within specific TSA approaches, there are problems with achieving agreement on what constitutes a reasonable (and disjoint) set of problem-solving mechanisms [Chandrasekaran 86].

Some of the common task-specific approaches to knowledge-level modelling include **Role-limiting Methods** (i.e. Mole) [Marcus 88], **Components of Expertise** [Steels 90], **KADS** [Wielinga 92], and **Generic Tasks** [Chandrasekaran 86].

The **Mole** system (based on role-limiting methods) is used for knowledge acquisition and explanation for cover-and-differentiate tasks often associated with diagnosis. The grain size of role-limiting methods is quite small, and the resulting models are very well specified. In fact, the models can be machine-translated into sets of rules for directly producing the desired behavior.

Components of Expertise, a componential framework directed at the knowledge and knowledge-use levels, attempts to combine four viewpoints into a complete methodology for expert system design. The components of expertise include tasks/subtasks, domain models, case models and problem-solving methods. The objective of the work is to create a "handbook" which can assist by mapping task features into expert system solutions [Steels 90].

KADS uses a "componential methodology" to construct layered models of tasks, including distinct layers for domain, inference and task knowledge. KADS serves as a knowledge engineering methodology by guiding the modification of an interpretation model chosen by the designer. However, following the KADS methodology results in the production of a high-level design that cannot be used to automatically produce executable code. KADS is therefore primarily a tool for knowledge-acquisition and task analysis, not for supporting implementation directly.

The **Generic Task** (GT) approach to knowledge-based systems differs in several ways from other TSA approaches. First, the grain size is comparatively large. Common tasks are directly supported by specific tools designed to help with problem-solving. There is considerable debate over the number and nature of the Generic Tasks, yet several tools have emerged to assist with common tasks such as hierarchical classification (CSRL)

[Bylander 86], functional reasoning (FM) [Sticklen 87] and routine design (DSPL) [Brown 89]. These tools have been used to solve a variety of problems from many domains [Chandrasekaran 86]. Each tool contributes knowledge structures and control strategies appropriate to the task, thus providing an underlying theory of commonality between tasks which are similar and easing system development by providing tools to assist with problem-solving. The GT methodology will play a central role in the implementation of several key problem-solving components of IPCA.

2.1.4 The Contributions of Knowledge-based Design

One GT tool, the Design Structures and Plans Language (DSPL), makes a central contribution to IPCA by providing a problem-solving mechanism in which to implement one of the components in the architecture. The problem-solving task of selecting appropriate fragments and goals to form the overall control plan implementation has been assigned to a DSPL-based component. That is, the synthesis of lists of fragments and goals which are suitable for inclusion in the final, generated control plan is treated by IPCA as a "routine design activity."

In general, the knowledge-based design process is concerned with the timely production of designs which meet specific criteria [Coyne 90]. The resulting design is an abstraction from which an artifact can be constructed. In particular, routine or "class 3" design is concerned with a particular subclass of the overall design problem [Brown 89]. Routine design is concerned only with the production of designs which do not require "innovation." Some requirements for GT-based class 3 design problem-solving include the existence of effective problem-decompositions for the task, the availability of compiled design plans for the lowest-level "component" parts of the design subtasks and explicit knowledge about how to handle design component failures. It is hypothesized that much of design falls under the type III designation [Brown 89].

A DSPL problem-solver contains a complex problem-solving mechanism which is described in detail in chapter 4. DSPL supports a hierarchical decomposition of an overall routine design task into a series of smaller design tasks, including suggested plans to accomplish each of those tasks and knowledge to recognize when design failure of a sub-component might occur. The problem-solver selects and attempts to follow design plans to accomplish each of the subtasks in turn until the overall design is complete. This process normally involves backtracking whenever failures occur, i.e. when one or more of the design constraints are not satisfied. If it is determined that the constraints cannot be met with any of the available design plans, the design effort ends in failure and either the constraints must be relaxed or one or more new design plans must be created which might prove more successful.

The underlying problem-solving framework of GT-based class 3 design is one of state-space search in a large but finite problem space [Newell 80]. Although the overall process of routine design is theoretically a constraint-satisfaction problem within such a space, conventional approaches to constraint satisfaction are not efficient at solving large, practical, complex problems in design. Propose and revise strategies can potentially cause the entire state-space to be searched while other methods, such as hill-climbing, can result in problems with local optima [Brown 89]. DSPL seeks to avoid these pitfalls by providing appropriate, structured knowledge to guide the design process.

2.2 Research Approach and Architectural Introduction

The IPCA approach to intelligent control combines some of the strengths of task-specific architectures (in particular, the GT approach to knowledge-based systems), the DSPL language for routine design, a knowledge-based sponsor-selector mechanism [Punch 89], classical and reactive planning approaches and control theory [Decker 94]. The resulting architecture is made possible by contributions from all of these areas.

The IPCA architecture provides a unique combination of design, control and planning activities and has been tested in a practical, real-world domain to determine its effectiveness. The domain chosen involves the control of a microwave-based fabrication process for composite materials, though the IPCA architecture is not restricted to this domain.

The overall information processing task addressed by IPCA involves the construction of planning and control algorithms for complex (process control) tasks.

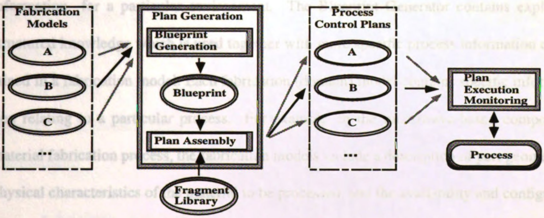

Figure 2.1 - Overall IPCA System Architecture

As shown in figure 2.1, the overall problem is separated into three components:

- **The blueprint generation task**
- **The plan assembly task**
- **The plan execution monitoring task**

Although the context of composite material fabrication will be assumed in order to describe the specific functions of the components in greater detail, the IPCA architecture can be applied to many other process control domains.

The **blueprint generation task** is one of routine (class 3) design and is accomplished by a hierarchical collection of specialists. Each specialist contains a set of design plans addressing a particular sub-task of the overall design problem. Knowledge about the process to be controlled is used in conjunction with knowledge about the applicability of particular control and planning approaches to synthesize an applicable plan blueprint for a specific environment. The plan blueprint specifies a collection of plan and control method fragments (strategies, policies, approaches, algorithms) and appropriate goals and control information for a particular environment. The Blueprint Generator contains explicit, structured knowledge which is used together with more specific process information contained in a fabrication model. Each fabrication (domain) model contains specific information relating to a particular process. For example, in the microwave-based composite material fabrication process, the fabrication models include a description of part geometry, physical characteristics of the material to be processed, and the availability and configuration of processing hardware, actuators, and sensors.

The **plan assembly task** involves retrieving the appropriate plan and control method fragments referenced in the plan blueprint and assembling them into a hierarchical structure based on knowledge about their applicability. The resulting structure (which is specific to a single process, material and geometry) is called the **process control plan**.

The **plan execution monitoring task** involves the application of the process control plan, based on information gathered through real-time process monitoring. One fragment from the process control plan is chosen for execution during each cycle, based on estimated utility in the current process state. The selected fragment is executed, thereby producing a series of control outputs and goal state changes which effect the process. The execution cycle then repeats after all sensed information (and thus, the estimated process state) is updated. The same process control plan is used throughout the execution process, so process control plans must typically contain a variety of control methods, plans and strategies for a variety of problem-state contingencies.

Chapter 4 contains a task-level analysis of each of the three highest-level problem-solving tasks. The architectural details of IPCA are reviewed in chapter 5.

2.3 Expected Results and Evaluation Criteria

Expected results and evaluation criteria are divided into two areas: those expected within the area of knowledge-based systems and those expected in the domain.

With regard to knowledge-based systems research, particularly within the GT approach to knowledge-based systems, this research is expected to offer the following contributions:

- **Analysis, architecture and implementation of a complete knowledge-based system including process control plan generation, assembly and execution.**
- **An extension of the GT approach to knowledge-based design and planning through the incorporation of a knowledge representation and control strategy for the synthesis and execution of process control plans.**
- **A structured framework for representing knowledge regarding composite materials processing.**
- **Broad applicability and ease of knowledge transfer to other process control domains.**

An intelligent control system built on the IPCA architecture has been implemented in VisualWorks SmallTalk™ and integrated with an existing implementation of DSPL [Kamel 94]. The resulting implementation has met certain subjective criteria, including responsiveness, ease of use and interface consistency. It will be shown that the implementation demonstrates the viability of the IPCA architecture and validates the appropriateness of the architecture in the chosen domain.

Within the domain of tuned single-mode microwave-based polymeric composite material fabrication, the research is expected to produce the following contributions:

- **A facility to support the automated synthesis and execution of discrete process control plans yielding composite structures of acceptable quality within a rapid timeframe.**
- **Knowledge-based frameworks for effectively representing domain-specific knowledge of single-mode microwave-based composite materials processing.**
- **Breadth and depth of domain coverage sufficient to demonstrate the applicability and effectiveness of IPCA in this domain.**
- **A reduction in the average thermal gradient within the material during curing beyond that achieved through the use of previous control approaches.**
- **Theoretical proof-of-concept for rapid, controlled material heating.**
- **Exploration of the limitations of mechanical mode tuning.**
- **Ability to support additional fabrication technologies including variable-frequency microwave tuning.**

The initial implementation supports one polymeric matrix material and three geometric structures as a proof-of-principle. Quality judgements of the resulting composite structures are based largely upon physical examination and analysis of temperature gradients measured during fabrication (see chapter 7). IPCA showed a significant reduction in the measured thermal gradient when compared to fixed-mode microwave fabrication techniques. It is further demonstrated that a reduction in thermal gradient has been achieved when compared to static “complementary” mode switching approaches, and that the gradient is ultimately limited by the mechanical response of the mode switching apparatus. A degree of breadth in domain coverage is also demonstrated. Finally, a foundation is offered for the ability to conduct rapid materials processing using higher power and variable-frequency tuning.

CHAPTER III

3.1 Domain Introduction

The effectiveness of the IPCA prototype has been demonstrated by applying the implementation to a series of related real-world process control problems from a particular domain. All experimental evaluations conducted with IPCA address the domain of polymeric composite material fabrication. In this section, the domain characteristics which are responsible for the suitability of IPCA to the chosen domain are examined.

The majority of composite materials are composed of two constituents. The **matrix** material is used to support and protect the **reinforcement**. There are several classes of matrix, including polymeric, metal and ceramic types. Similarly, there are several types of reinforcements, based on shape (including fibers, particulates and whiskers) and on material (such as glass or carbon). Polymeric carbon-fiber reinforced composite fabrication has been used as a focus for demonstrating the effectiveness of the IPCA architecture. The process to be controlled is based on a particular class of polymer composites known as **thermoset** polymeric composites. For additional detail regarding composite materials, refer to appendix E.

Fabricating polymeric thermoset composite structures is accomplished by initiating and maintaining a chemical crosslinking reaction through the application of heat and/or pressure. Most commonly, composite structures are fabricated or "cured" in autoclaves under heat and pressure, and these approaches are briefly introduced in section 3.3. However, the processing technology used for demonstrating IPCA uses microwave energy to initiate and maintain the exothermic crosslinking reaction. Microwave processing has a number of desirable features which are described in detail in section 3.4.

3.2 Related Domain Research

Research into low-cost, high-speed composites processing addresses a wide range of disciplines, technologies and approaches, including high-speed processing, basic materials research and the use of composites as structural components. The related research forms a domain context which extends well beyond intelligent control of microwave-based fabrication for polymeric thermoset composites. For completeness, an overview of the surrounding domain context is provided. This section discusses some of the related research into composite materials and structures and the issues involved with low-cost, high-speed materials processing. Closely related research conducted at the Composite Materials and Structures Center (CMSC) at Michigan State University is presented.

3.2.1 Composite Materials Research at Michigan State University

The CMSC is a cooperative multi-disciplinary research center for composite materials and processing research, which also serves as a technology transfer agent to the industrial sector. The work conducted at the CMSC includes research sponsored by the NSF Center on Low-Cost, High-Speed Polymer Composites Processing and the Michigan Materials and Processing Institute (MMPI). The relationship between CMSC, MMPI and the NSF center is demonstrated in figure 3.1.

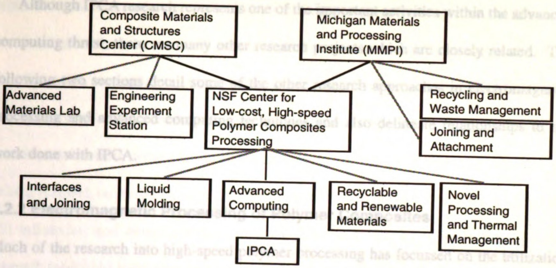

Figure 3.1 - Relationships Between CMSC, MMPI and the NSF Center

MMPI is a non-profit institute focused on polymer composites for the durable goods industries. The board of directors and executive committee oversee four technical committees. Technical committee 1 (TCOM1) addresses Design and Performance Issues in Composites. The NSF Center is considered to be TCOM2. TCOM3 addresses Recycling and Waste Management. TCOM4 is Technology Transfer. Each committee has a board which oversees the projects.

The NSF center is focussed on increasing the competitiveness of US-based industry and encouraging local economic development through a federal/state/industry partnership. The NSF center is focused on basic and applied research which will help to provide high-quality polymer composites at high-speed and low-cost for consumer goods applications. The technical thrust areas for the NSF center include novel processing and thermal management, liquid molding, recyclable and renewable materials, advanced computing techniques for composite design and processing, and interfaces and joining. The IPCA research is included under the thrust of advanced computing techniques.

Although IPCA research represents one of the important activities within the advanced computing thrust, there are many other research projects which are closely related. The following two sections detail some of the other research approaches in electromagnetic processing and advanced computing techniques and also delineate relationships to the work done with IPCA.

3.2.2 Electromagnetic Processing of Polymer Composites

Much of the research into high-speed polymer processing has focussed on the utilization of electromagnetic fields for rapid material heating. Electromagnetic processing forms an important basis for the research conducted with IPCA.

Research in electromagnetic processing includes work on cavity characterization and automation detailed in [Adegbite 95]. This work formed one of the starting points for the work with IPCA and also influenced the development of the IPCA architecture. The cavity described in chapter 3 was built while the IPCA research was in progress, and the cavity was characterized as IPCA was in the later stages of development. The goals of this related work were to automate the single-mode microwave process studied previously in [Jow 88] and to study the high-speed processing capabilities of the automated microwave fabricator.

The microwave cavity described in detail in chapter 3 was built and outfitted with an external control system as one part of the work described in [Adegbite 95]. In addition, the electromagnetic modes of the resulting cavity were empirically characterized through a series of experiments and the results were compared with theoretical EM field models [Adegbite 95]. A closely-related variant of the conventional control system developed and described in [Adegbite 95] was developed and used in conjunction with the IPCA prototype. This conventional control system is described in chapter 5 and appendix D. The conventional control system was designed to perform control and sensing functions which are specific to the controlled (domain) process. As described in section 5.4, IPCA cannot function without an External Controller. It is therefore important to emphasize that the conventional control system forms an integral portion of the IPCA prototype, and that the related research on cavity characterization and closed-loop control described in [Adegbite 95] influenced and complemented the work done with IPCA. This interdisciplinary research model has been used extensively and with great success within MMPI.

The mode switching research of [Fellows 91] is also closely related to the work conducted with IPCA. Fellow's research explored a technique known as "mode switching" which involves alternating between two or more resonant modes (see section 3.4) to help ensure uniform heating. The work was extended to allow for the processing of complex, non-planar shapes, and showed an encouraging improvement over the previous microwave-based processing. However, the approach required the manual generation of heating profiles for each mode to be considered, which proved difficult and time-consuming. The approach also relied upon pre-defined static thresholds for controlling the selection of modes. A variety of materials could not be processed without changes to the controller software, nor could a variety of microwave fabricators be supported without creating multiple control programs. Although the work demonstrated the viability of the mode switching approach, it also pointed toward a need for a more sophisticated support system for composite material fabrication. One of Fellow's observations was that more sophisticated control algorithms were needed to effectively control material processing using mode switching techniques [Fellows 91]. The quantitative results obtained in [Fellows 91] are valuable, and are later used as a basis for evaluating the laboratory fabrication results obtained in this research.

3.2.3 Research in Advanced Computing Techniques

A variety of advanced computing techniques for polymer composites processing have been explored and are related to the research which eventually produced the IPCA prototype. The research projects described in this section share the goal of using advanced computing techniques to assist with the design and manufacture of composite structures in some manner.

The research described in [Kamel 94] resulted in a knowledge-based design system which assists a human designer with the task of generating multiple composite material selections (designs) based on specific design requirements. The underlying system architecture is based on a tool developed to support routine design activities [Brown 89]. The research showed that knowledge-based material selection was possible, and that knowledge-based material selection was a feasible approach to a difficult problem. In addition, Kamel made an important contribution in extending the underlying routine design framework to generate multiple designs. As with many MMPI research efforts, Kamel's work made contributions in the domain of composite material design and also within the knowledge-based systems community.

The system described in [Kamel 94] was subsequently extended to cover a wider range of materials as a result of work done by [Lenz 94]. The resulting third-generation decision support system (COMADE) was not an architectural extension, but information regarding many alternate materials was included to provide broader knowledge-based support for material selection. Support for thermoplastic matrices, expanded thermoset coverage and chopped-fiber matrix materials were included in the COMADE effort, proving out the extensibility of the knowledge-based design approach.

The approach of developing advanced computing architectures and subsequently extending their capabilities within their respective domains has proven successful within the NSF Center. This approach forms the working model for many of the projects within the Center, and allows the interdisciplinary development of advanced computing architectures and small-scale prototypes (such as IPCA). The approach also provides extensible computing frameworks which are typically used to represent additional knowledge from domain experts to increase their scope and their usefulness.

Other intelligent systems for decision support have been developed, including a second-generation fabrication technology selection system described in [Moy 94]. This fabrication technology selector is used to recommend appropriate processing technologies when given a set of design requirements including material types, economic factors, process throughput and part geometry. The system uses a knowledge-based hierarchical selection mechanism to choose an appropriate fabrication method or multiple methods.

The ultimate goal of the Advanced Computing research thrust involves combining the intelligent support systems into an automated end-to-end support system suite for composite material selection, part design, fabrication technology selection and fabrication control [McDowell 93]. The IPCA research contributes to this goal by providing a framework for effective knowledge-based fabrication control for a variety of composite materials and geometries. Since it is a first generation system, the current IPCA prototype does not yet support a wide variety of materials or processing technologies. However, the IPCA system ca

ner ve

has al

the ca

3.3 C

Therm

[Loos

coolin

qualita

Th

functio

used to

feedba

dielect

imate s

typical

tem can be extended to cover additional materials and processing technologies in a manner very similar to the effort which ultimately produced COMADE. The IPCA prototype has already been extended on a small scale, and those extensions are described in detail in the case study supplied as appendix F.

3.3 Conventional Composite Processing

Thermoset polymeric composites are typically cured in autoclaves under heat and pressure [Loos 83]. Key process variables include time, temperature, heating rate, pressure and cooling rate. The relationships between these variables are complex, but it is possible to qualitatively characterize major interactions to a useful extent.

The curing profile specifies temperature and pressure to be applied to the material as a function of time. Thermocouples and other temperature-sensing devices are commonly used to monitor the internal temperatures in the material being cured in order to provide feedback to either a human operator or a control system. Also directly measurable is the dielectric constant of the material, which indicates ion mobility, and, indirectly, an approximate state of the cure. Viscosity may also be measured through the use of rheometers. A typical curing profile for an epoxy resin is shown in figure 3.2.

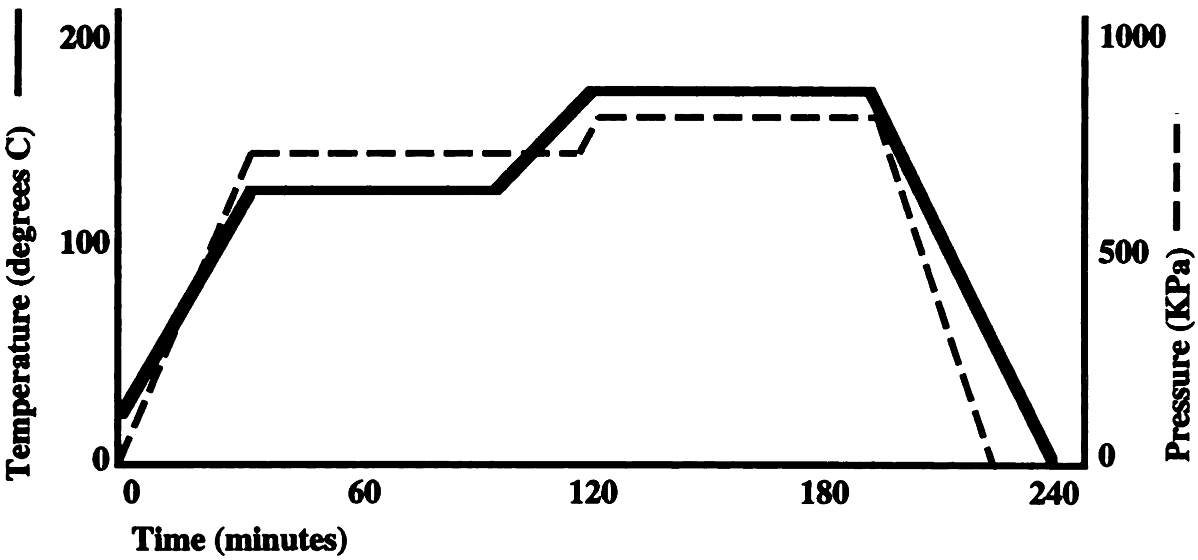

Figure 3.2 - Typical Autoclave Curing Profile

The typical autoclave curing profile depicted in figure 3.2 includes a ramp-up time for allowing the material to come up to a desired processing temperature while minimizing the formation of large voids, caused by trapped water vapor in the material. Pressure is applied to attempt to consolidate the material and minimize void formation. The application of heat and pressure continue until the curing cycle is complete. Often, the material is **postcured** in an attempt to ensure that the crosslinking reaction is complete and the finished structure has the best material properties possible.

3.4 Microwave-based Processing

A number of alternate fabrication methods for polymer composite structures have been investigated, including the microwave curing of thermoset polymers [Asmussen 87]. Some of the advantages of microwave curing are rapid, controlled heating, high energy efficiency, low cost and high throughput [Wei 93]. It is for these reasons that microwave processing technologies are being investigated as an energy-efficient means to produce low-cost, high-quality composite structures at high speeds.

Microwave processing involves curing composite structures through interaction with microwave-frequency electromagnetic (EM) fields. The applied EM field causes rapid heating within the composite structure. The interaction between composite structures and the applied EM field is highly complex, however, and dependant upon a number of factors, including the shape, volume and physical characteristics of the cavity used and the characteristics of the composite structure. The geometry, mass distribution, placement and dielectric constant of the composite structure all affect how the EM field will be perturbed. There are so many factors to account for that predicting a pattern of interaction between a given EM field and a given composite structure remains an extremely difficult and unsolved problem. However, some generalities may be stated. For example, "lossy" reinforcement materials such as glass-fiber are distinguished by their higher dielectric constants. Lossy materials therefore heat more slowly than lossless (e.g. carbon-fiber) reinforced materials. This is true because lossy materials absorb less energy from the applied EM field. The overall dielectric constant of the material thus plays a key role in determining how the structure will interact with an applied EM field.

The cavity characteristics are also central in determining EM field interactions, since the cavity behaves as a waveguide which focuses EM energy. Depending on the physical characteristics of the cavity and the frequency of the applied EM field, microwave heating may be single-mode or multi-mode in nature. The two types of heating differ in how the EM energy is focused.

The applied EM field used for multi-mode heating may be visualized as a linear combination of several distinct EM "basis" vectors. The cavity is detuned in the sense that a stable standing-wave pattern does not exist and the energy is "randomly" distributed. The pattern is not actually random, of course, but depends upon so many factors that it is impossible to predict without making many simplifying assumptions. This means that the energy delivered through multi-mode heating is largely unfocused. Power absorption is generally quite low and therefore multi-mode systems do not heat efficiently, although the overall heating distribution may be more uniform than could be achieved through the use of any single basis field vector. The main advantage to multi-mode heating is that it can be used to heat many different types and shapes of materials with no need for any feedback, tuning or focusing. For this reason, commercial microwave ovens use multi-mode heating exclusively. Although multi-mode heating is not efficient, it can be used to heat a wide range of materials to moderate temperatures fairly evenly.

Single-mode heating offers the ability to heat materials more rapidly and much more selectively. When using a single-mode heating technique, a standing wave is set up in the cavity so that maximum energy may be focused into the material. This approach has both a major advantage and a disadvantage. The advantage is that the material heats rapidly and, since efficiency is very high, that the material may be heated to high temperatures if desired. The disadvantage is that all the EM energy is doing work on the material in a specific, static, and generally uneven fashion. This means that the pattern of heating within the material is often uneven, since the applied EM field is more focused on some parts of the material than on others. This is troublesome for composite materials manufacturing because thermal gradients in a composite structure are undesirable since non-uniform heating negatively impacts the physical properties of the finished structure. Single-mode heating techniques produce efficient, rapid energy delivery, but power application and mode selection must be carefully monitored and adjusted to minimize thermal gradients and maintain specific material processing temperatures. IPCA is designed to specifically address this problem by providing automated process control plan construction and execution.

Research on mode-switching for processing 3-dimensional low-loss polyester/glass composite structures based on empirically-gathered heating data can be found in [Fellows 95]. Mode-switching techniques proved useful in this work for maintaining sample temperatures between the onset of reaction and the onset of thermal degradation for complex 3-dimensional structures. Two or more "complimentary" modes were selected based on heating experiment results with a cured sample. The complimentary modes were then chosen during subsequent processing runs in an attempt to achieve uniform curing. Mechanical testing confirmed that the properties of the resulting structures were not as desirable as those achieved through autoclave curing. Best-case results from 12-ply 2-dimensional polyester/glass samples and 60-minute cure cycles involved temperature gradients of approximately 30 degrees C. Fellows concluded that variable-frequency tuning would not only reduce these gradients further, but would allow the application of pressure during processing.

A diagram of a sample composite layup is shown in figure 3.3. The layup includes a hand-laid composite structure (the sample) enveloped by several layers of various films and enclosed in a teflon mold.

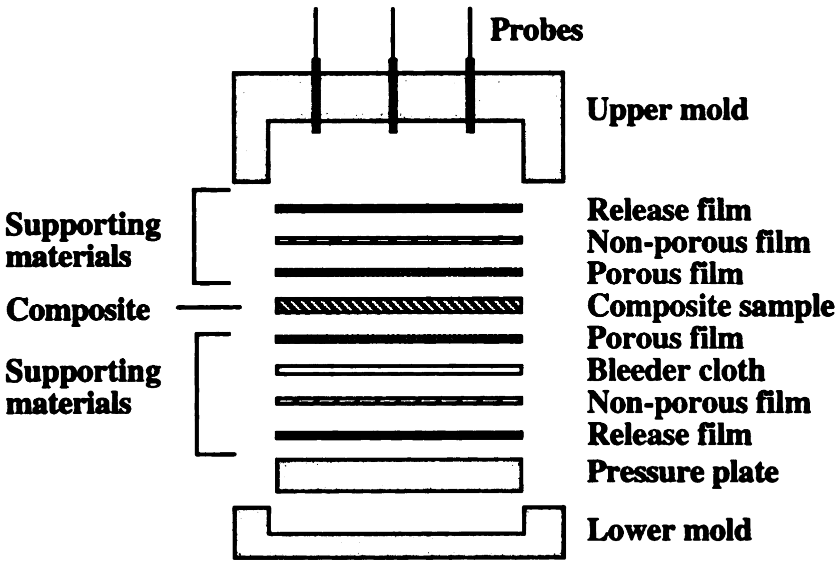

Figure 3.3 - Composite Layup for Microwave Processing

The porous film on either side of the composite allows excess resin to flow away from the sample while preventing adhesion. Excess resin flows through the porous film into layers of bleeder cloth. One layer of bleeder cloth is included for every 3 plies in the sample. Under the bleeder cloth is a layer of non-porous film, designed to contain any excess resin which is not absorbed by the bleeder cloth layers. A release film is used between the non-porous film and the mold as a final guard against leakage of epoxy. The entire layup is enclosed in a teflon mold, which ensures that the material will remain in the desired shape during processing, while simultaneously providing access for temperature probes to the composite structure within. The structure of a typical mold is shown in figure 3.4. This mold was used in the fabrication of 1" square composite structures, detailed in chapter 7. Probe placement and identification are given in chapter 6.

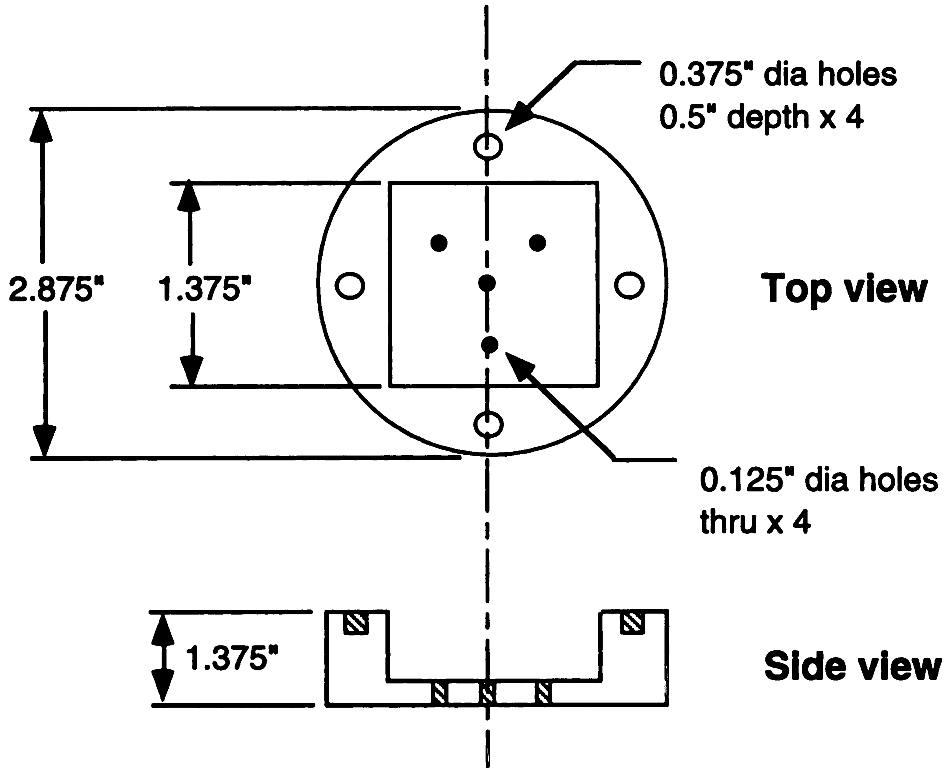

Figure 3.4 - Teflon Mold for Microwave Processing

There are several methods for creating and maintaining a particular EM mode within a given cavity. If the cavity dimensions are assumed fixed, various modes can be selected by changing the frequency (and therefore the wavelength) of the applied EM field until it is an integral multiple of the cavity length. For example, if the cavity were 40cm long, theoretical (empty-cavity) modes would be present at wavelengths of 40cm (n=1), 20cm (n=2), 13.33cm (n=3), 10 cm (n=4), and so on. The exact wavelengths required to produce standing waves in a loaded cavity would depend upon many other factors, as mentioned earlier. Nonetheless, it would be a simple matter to switch from one mode to another by changing frequencies. Although this approach would offer the ability to rapidly and efficiently change modes without altering the physical characteristics of the cavity, it would require an extremely costly variable-frequency microwave power supply. Many research budgets do not allow for such an expenditure, so a lower-cost cavity tuning approach was developed.

3.5 The Experimental Single-mode Resonant Cavity

An experimental single-mode variable-length cavity has been developed for laboratory-scale microwave-based composite materials processing [Adegbite 95]. The cavity features a pair of motor-controlled actuators which alter the cavity length and the coupling probe depth to facilitate the selection of specific EM field patterns (modes) in conjunction with a fixed-frequency microwave power supply. Although this technique does not allow instantaneous switching between modes, it has the advantage of low cost and is nearly as flexible as a variable-frequency power source. However, there are limitations imposed by the time required to switch between modes. These limitations are explored in chapter 6. A diagram of the experimental cavity is shown in figure 3.5.

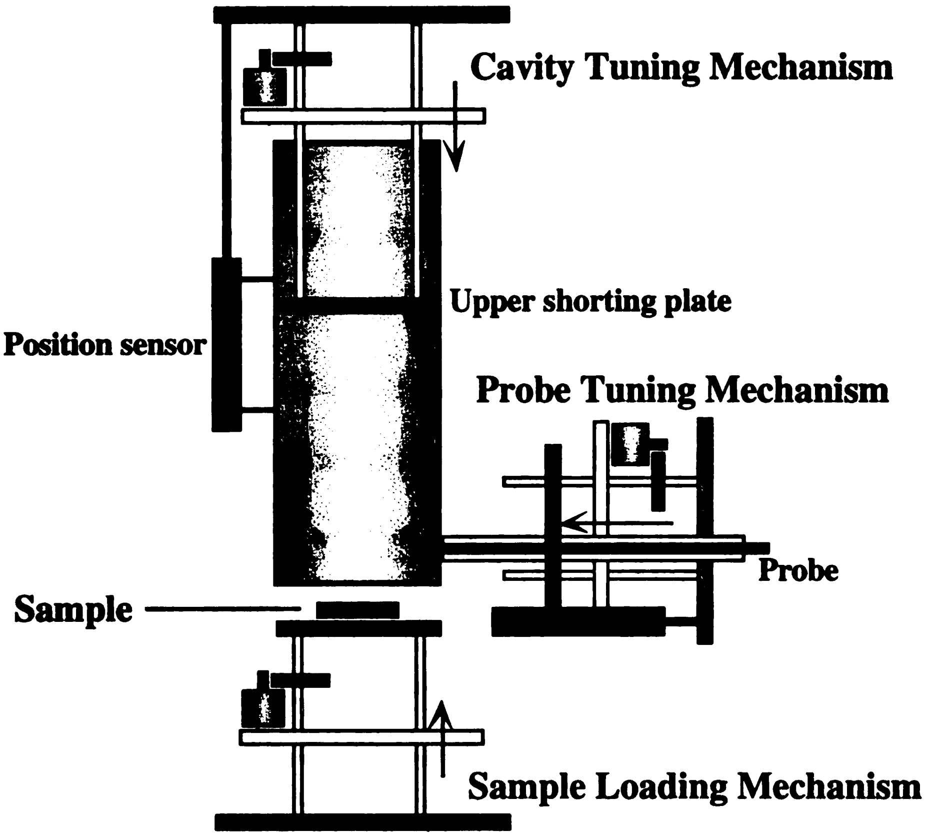

Figure 3.5 - Experimental Single-Mode Microwave Cavity

The experimental cavity is a hollow, cylindrical brass structure 17.78 cm in diameter and 40.00 cm in length. Two brass shorting plates are used to seal the ends of the cavity. The lower shorting plate is raised into position by a motor-controlled sample loading mechanism and remains fixed in position until the sample is removed after curing. The lower shorting plate serves only to support the sample and seal one end of the waveguide. It does not move at any time during material processing. The upper shorting plate is also equipped with a motor-controlled cavity tuning mechanism and is used to select particular EM field patterns during processing. Each desired mode corresponds to a particular cavity length (L_C). In addition, a probe-tuning mechanism is used to modify the depth of the microwave coupling probe (P_D). This adjustment is useful when fine-tuning the cavity modes, since slight movements of the coupling probe can compensate for small variations in sample placement, geometry and composition without the need to modify the cavity length.

A number of sensors are included in the design of the cavity. A pair of precision 5K-Ohm linear potentiometers are used as linear sensing elements for the cavity length and probe depth positions. These sensors are employed to give feedback on the position of the upper shorting plate and the microwave coupling probe. The incident and reflected power are monitored through the use of a pair of power sensors and analog meters.

In addition to these dedicated sensors, a variety of invasive and non-invasive sensors may be used to monitor the temperature of samples, including fiber-optic probes and non-invasive IR sensors. Most experiments conducted with this equipment have employed a Luxtron 755 four-channel fluoroptic temperature measuring system. Two fluoroptic thermometers (each with four temperature probes) can be used to monitor up to eight temperatures inside the cavity. Electrode-based dielectric sensors are not used since metal cannot reside inside the cavity during processing, but alternative methods for estimating dielectric constant of the material are under investigation.

Refer to appendix D for more detail about the External Controller and the equipment setup. Chapter 6 contains a description of probe placement and configuration for the various fabrication experiments.

3.6 External Controller Overview

In the current implementation, the External Controller preforms sensor monitoring, preconditioning, low-level mode tuning, and other activities thus relieving the remainder of the architecture of these responsibilities. The prototype External Controller is written in the National Instruments LabVIEW™ development environment [National 91]. LabVIEW™ provides graphical programming tools and an extensive "virtual instrument" library for communication with a wide variety of laboratory instruments connected to an IEEE-488 bus or attached to a National Instruments™ I/O interface. The External Controller monitors sensors attached to the I/O interface including temperature, power and position and is designed to maintain a tuning lock on whichever mode is selected by the intelligent controller. Since the External Controller runs on a dedicated machine, it can rapidly monitor reflected power and adjust the cavity length and probe depth positions to keep the reflected power to a minimum as the dielectric constant of the sample changes during fabrication. A PID control algorithm is employed to accomplish this task [Astrom 90].

Since the External Controller and IPCA run concurrently, the External Controller has time to perform sample averaging and sensor preconditioning to reduce sensor noise. For example, the External Controller can sample the A/D converters on the I/O interface hundreds of times per second. Spurious or noisy readings can be discarded while others are averaged since the intelligent controller only needs the sensor readings once every 5-10 seconds.

The External Controller serves to provide a vital link between IPCA and the physical sensors and actuators required to control most real-world processes. Although neither the form nor the implementation of the External Controller is specified by the IPCA architecture, the services which it typically provides are described. The interface protocol between the External Controller and IPCA is described in appendix D. Additional detail regarding the External Controller can also be found in appendix D.

CHAPTER IV

4.1 Theory of Operation

The overall problem-solving task addressed by IPCA involves creating a mapping from a set of sensor values to an set of actuator values for use throughout process control plan execution. The underlying nature of the process control task will therefore be examined. An example diagram showing the control problem in continuous time is given in figure 4.1.

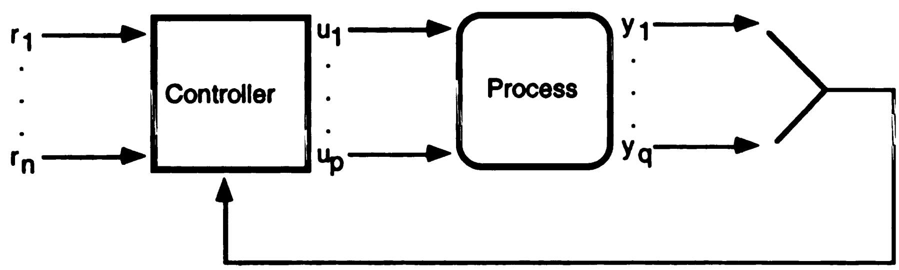

Figure 4.1 - Structure of a Continuous-time Control System

Referring to this continuous-time model, the controller maps from control inputs $y_1..y_q$ (sensed values representing the state of the process) and reference commands $r_1..r_n$ (describing desired states for the process) to produce appropriate control forces $u_1..u_p$ to be exerted on the process. This relationship is demonstrated by the equation

$$\wp(t) = \left(\left(\begin{bmatrix} r_1(t) \\ r_2(t) \\ \dots \\ r_n(t) \end{bmatrix} \times \begin{bmatrix} y_1(t) \\ y_2(t) \\ \dots \\ y_q(t) \end{bmatrix} \Rightarrow \begin{bmatrix} u_1(t) \\ u_2(t) \\ \dots \\ u_p(t) \end{bmatrix} \right), \forall (t \in \Re, t \geq 0) \right)$$

which assumes a continuous (linear or non-linear) time-variant system. As stated in chapter 2, classical control theory assumes linear, time-invariant systems. Therefore, it can already be seen that the production of such a relation for non-linear, time-variant systems will require more than classical control theory can offer. There are $n \cdot q$ continuous-valued time-dependant equations required to generate the p control forces for any given value of t. The complexity of producing this mapping for the continuous-time case will not be investigated, since modern state-space based digital control has been selected to model processes investigated by this research.

Modern digital controllers are designed as discrete-time systems using quantized (sampled) inputs and outputs. Discrete-time process control systems have a general structure as depicted in figure 4.2.

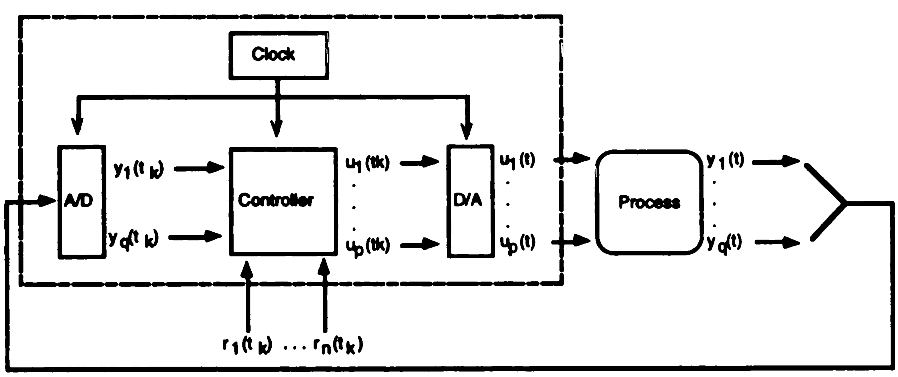

Figure 4.2 - Structure of a Digital Control System

The discrete-time digital controller samples the continuous process outputs $y_i(t)$ at regular time intervals to form discretized control inputs $y_i(t_k)$. These inputs, together with reference commands $r_i(t_k)$ are used by the controller to generate a set of discrete control forces $u_p(t_k)$. The discrete control forces are then transformed into continuous values $u_i(t)$ through the use of a D/A converter which holds the control force outputs constant between conversions to maintain a continuous output signal. The controlled process reacts to the control force outputs and is sampled again at the next discrete time interval.

4.2 Controller State-Space Complexity Analysis

Modern control theory makes a number of assumptions to simplify the task of creating an effective digital control system. The relationship between input and output may be expressed for the discrete-time case as

$$\wp(t_k) = \left(\begin{bmatrix} r_1(t_k) \\ r_2(t_k) \\ \dots \\ r_n(t_k) \end{bmatrix} \times \begin{bmatrix} y_1(t_k) \\ y_2(t_k) \\ \dots \\ y_q(t_k) \end{bmatrix} \Rightarrow \begin{bmatrix} u_1(t_k) \\ u_2(t_k) \\ \dots \\ u_p(t_k) \end{bmatrix} \right), \forall (k \in Z, k \geq 0)$$

where the t_k represent discrete time instances dependant on the sampling period, h. Thus, the relationship

$$t_k = k \cdot h, k \in Z$$

holds, and the sampling frequency (in samples per second, or Hz) is given by

$$F_s = \frac{1}{h}$$

Assuming that we are interested in controlling the time-variant process for T_f seconds, we have

$$k = \left\lfloor \frac{T_f}{h} \right\rfloor = \lfloor T_f \cdot F_s \rfloor$$

Further, assuming each control input y_i is quantized into Q_y levels (using a $\log_2(Q_y)$ bit A/D converter) and each of the reference commands r_i can be represented using Q_r discrete states (using $\log_2(Q_r)$ bits), then the size of the domain in the resulting control state-space is given by

$$k \cdot 2^{(Q_y \cdot q \cdot Q_r \cdot n)}$$

and, assuming Q_p quantization levels ($\log_2 (Q_p)$ bit D/A converters) for the control forces u_i, the co-domain of the control space is of size

$$2^{(Q_u \cdot p)}$$

Therefore, we find a set of

$$(k \cdot 2^{(Q_y \cdot q \cdot Q_r \cdot n)}) \cdot (2^{(Q_u \cdot p)}) = \lfloor T_f \cdot F_s \rfloor \cdot 2^{(Q_y \cdot q \cdot Q_r \cdot n \cdot Q_u \cdot p)}$$

distinct mappings from the domain to the co-domain.

The overall information-processing task of IPCA involves the design, synthesis and application of a selected mapping from among this set, using knowledge about the process to be controlled. But even for the most modest control task, the size of this space is enormous. For example, using only 4 process inputs and 4 control force outputs with 4 bit A/D and D/A converters, a sampling period of 1 second, and two 2-bit reference commands to control a time-variant process for 2 seconds would yield

$$\lfloor T_f \cdot F_s \rfloor \cdot 2^{(Q_y \cdot q \cdot Q_r \cdot n \cdot Q_u \cdot p)} = \lfloor 2 \cdot 1 \rfloor \cdot 2^{(2^4 \cdot 4 \cdot 2^2 \cdot 2 \cdot 2^4 \cdot 4)} = 2^{(2^{17} + 1)} \sim 4.5 \times 10^{39}$$

possible control mappings from which to select.

Of course, IPCA does not represent process control plans as simple mappings. The complexity analysis is intended only to show the size of the problem space. The method by which IPCA uses knowledge to generate process control plans is described in the following sections.

4.3 Knowledge Level Analysis

The proposal of the knowledge level (KL) hypothesis has been extremely influential in modern artificial intelligence research [Newell 82]. The KL hypothesis states that an implementation-independent computer systems level (the knowledge level) exists immediately above the symbol level. The symbol level, in turn, exists immediately above the implementation level. Each computer systems level is associated with particular primitive constructs and properties. For example, the gate level contains primitive elements such as AND, OR and XOR logic gates, flip-flops, clocks, and so on. The elements of the register transfer level (comprising registers, adders, latches, etc.) are constructed from the elements from the gate level. It is possible to discuss structures at the register transfer level without including details about the next lower (gate) level.

Similarly, the elements of the knowledge-level are composed of primitives from the symbol and implementation levels, making it possible to describe activities at the knowledge level without committing to any particular implementation. As with the other levels, the knowledge level is a distinct computer systems level complete with its own properties and constructs. The KL hypothesis argues that a problem-solving analysis of a task can be given without any reference to the particular implementation details of how knowledge is represented (frames, rules, semantic networks, etc.). In fact, a knowledge level analysis of the behavior of a system is based on an "external observer" standpoint, treating the system as a black box. The KL analysis, then, attempts to explain the behavior of the system in terms of what is does (observation), how it works (mechanization) and why it works the way it does (rationalization) [Van deVelde 93].

Therefore, the following analysis of the overall problem-solving task addressed by IPCA is given. The task analysis is from the viewpoint of an external observer, and avoids discussion of the implementation details of task decomposition and knowledge representation.

At the knowledge level, the overall information-processing task addressed by IPCA involves the generation of control forces which effectively manage a variety of complex controlled systems based on high-level descriptions of those systems and real-time feedback from the systems. The specific requirements for the IPCA task are thus subject to the nature of the input and output descriptions. These descriptions will therefore be made more explicit, without a commitment to any representation, before the analysis proceeds.

As input, we observe that the IPCA task accepts a description of a particular system to be controlled. This description is relatively high-level, i.e. it makes specific references to the real-world equipment involved and information about how it is configured into an overall processing system, but does not include any formal mathematical description of the process. For the domain of microwave composite material fabrication, the input description also contains information about the material to be processed, including the material composition and geometry. Again, these descriptions are high-level, including the names of the materials, for example, but not their chemical or structural descriptions. The output requirements for the task are somewhat modest. The system needs to generate a series of control force outputs which will effectively manage the controlled system in real-time.

Using the given information about the nature of the input and output of the task, we can describe some of the knowledge which must be present to support problem-solving. Obviously, the task must contain knowledge which supports a mapping from the current system configuration into a set of control force outputs. In order to generate meaningful control force outputs, the system must contain knowledge about the physical sensors and actuators detailed in the input description. Since the system is described in terms of physical equipment, materials and geometry, the task must utilize knowledge about the nature of the specific configuration of these items, including the relationships between sensed inputs and appropriate outputs. Knowledge about the characteristics of various materials must be included.

A more detailed KL analysis cannot be performed without making assumptions about how the system is structured and how knowledge is used inside the "black-box." Therefore, the overall task will be decomposed into three subtasks for further analysis.

4.4 IPCA Task Descriptions

The IPCA architecture views the overall problem-solving task as consisting of three distinct information-processing subtasks. As shown in figure 4.3, the tasks include **synthesis** of appropriate process control plan specifications, **assembly** of process control plans from compiled fragments and plan **execution** based on information gathered from real-time process monitoring.

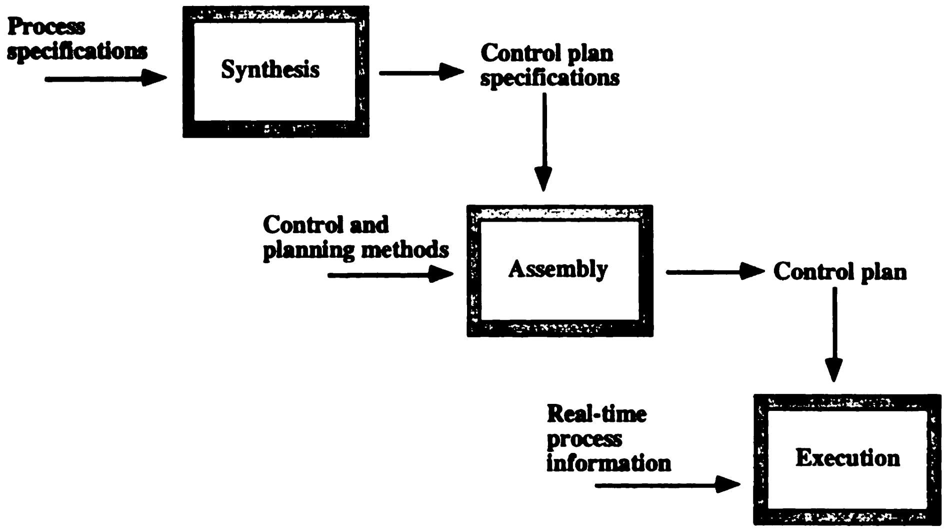

Figure 4.3 - IPCA Task Structure

The **synthesis task** involves using process specifications and compiled knowledge to construct a control plan specification. Process specifications contain high-level descriptions of equipment and materials, sensor types, placement and actuator configurations. The control plan specification, called the **plan blueprint**, contains information about which control and planning methods are applicable and how they relate to the associated policies and goals required for process management. Knowledge required to synthesize the plan blueprint includes:

- **Knowledge of appropriate fabrication goals for the process.**
- **Knowledge of the applicability and utility of various planning and control methods in terms of ability to achieve the goals provided.**
- **Knowledge about how the methods could be used together to form policies and processing strategies.**
- **Knowledge regarding the physical properties of the material to be processed.**

The **assembly task** involves the construction of a complete, executable process control plan which incorporate the methods, goals and policies set out in the plan blueprint. Knowledge required in order to accomplish this task includes:

- **Knowledge of the applicability and utility of each planning and control method in terms of ability to meet the goals provided.**
- **Declarative procedural (executable) representations for each planning method and control method.**
- **Knowledge of dependencies among planning and control methods.**

The **execution task** involves applying the process control plan while gathering and reacting to process information in real-time. This task requires:

- **Knowledge about how to process raw sensor data into meaningful abstract characterizations of the process state.**
- **Knowledge about how to select appropriate planning and control methods at appropriate times during execution.**
- **A facility for supporting goal revision, thereby allowing planning methods to alter the initial goals provided in the plan blueprint.**

IPCA goals are described in detail in chapter 5.

4.5 GT Approach

As introduced in chapter 2, the GT approach to knowledge-based systems forms a basis for IPCA. The architecture, however, also adds to the set of GT tools by providing an extension to the routine design/planning framework of DSPL. The architectural implementation of IPCA adds knowledge-based support tools for plan assembly and execution to the existing DSPL system. The full functionality of the routine design mechanism provided by DSPL is still available, but is enhanced with mechanisms designed to support plan assembly and execution.

A second GT-style problem-solving mechanism is used in the IPCA architecture to support the knowledge-based selection of appropriate plan alternatives during plan execution. A sponsor-selector mechanism will provide this support, using matchgroup tables to yield appropriateness measures under various conditions. Both mechanisms are discussed in detail in the following sections.

4.5.1 DSPL and the Blueprint Generator

As briefly introduced in section 2.2, the design structures and plans language (DSPL) provides a problem-solving mechanism for routine design and planning tasks [Brown 89]. Specific knowledge representations and control strategies are included to facilitate this task-specific problem-solving approach for routine design problems. A diagram of the major DSPL constructs is given in figure 4.4 below.

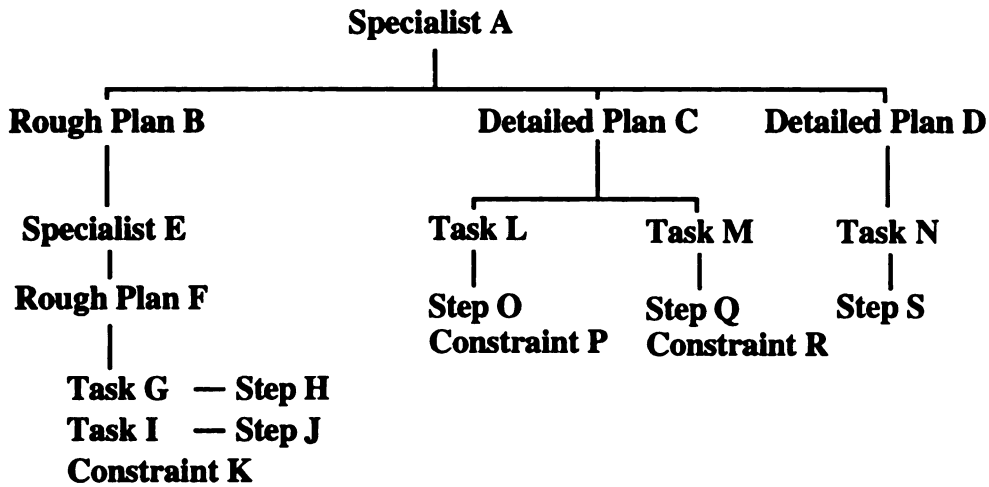

Figure 4.4 - DSPL Design Agents

Using DSPL terminology, an overall design task is hierarchically decomposed into a number of sub-tasks, each of which is supported by a specific **design specialist**. Each specialist includes a list of one or more plans from which to select in order to design a specific sub-component for which it is responsible. Specialists may also contain a list of constraints which are used to determine whether a particular plan was successful. Plans can contain references to specialists, so that any portion of the design task can be decomposed into many levels of sub-tasks if required. Plans can also contain direct references to tasks, possibly interspersed with constraints. Each DSPL task, in turn, is a design agent which contains a list of steps and constraints. Steps are used to set the value of a single design variable, while constraints are used to determine whether steps were successful.

There are two types of plans: rough-design plans and detailed design plans. The DSPL system makes a first pass through the entire specialist hierarchy in an attempt to complete a rough design. During this **rough design phase**, each specialist tries to apply it's rough-design plans, tasks and steps in an attempt to meet all design constraints. Referring to figure 4.4, Rough Plan B and Rough Plan F would be utilized during the rough design phase.

The rough design phase is used to determine whether a successful design is possible. This task is accomplished by determinng whether the critical design variables can be set in such a manner so as not to violate important design constraints. If the rough design phase indicates that a successful design is not possible, the detailed design phase is never attempted. Referring to figure 4.4, if Constraint K is not met, the rough design will fail since there are no alternate rough design plans from which to choose. This two-phase design approach generally saves time since clearly impossible designs are not be attempted in detail. However, if the top-level specialist ultimately reports success (meaning all rough design constraints have been satisfied) the detailed design phase is initiated.

Having established critical design values during the rough design phase, a **detailed design phase** is subsequently attempted during which each specialist selects from a set of detailed design plans. The detailed design phase proceeds just as the rough design phase did, with each specialist attempting to use the detailed design plans (from most to least suitable) to design each sub-component in turn. Referring to figure 4.4, either Detailed Plan C or Detailed Plan D could make an initial attempt to design the component for which Specialist A is responsible. If the more suitable of the two plans should fail, the other plan is then invoked (unless the plan sponsor indicates that it should be ruled out). If the detailed design pass ends in success, meaning that all design constraints are satisfied, the design is complete. Note (also from figure 4.4) that Detailed Plan D will always succeed since there are no constraints within any subordinate agents. Also, Detailed Plan C will always attempt both Task L and Task M, requiring that both Constraint P and Constraint R be satisfied for success. For additional detail on the problem-solving mechanism of DSPL, see [Brown 89].

A simplified structural representation of the DSPL specialist, plan, task hierarchy used for a portion of the prototype IPCA implementation is shown in figure 4.5. Many of the design agents have been omitted for clarity and only a single fabrication technology is shown. The top-level specialist, called the "fabrication specialist," is responsible for choosing an appropriate processing technology and directing the overall process control plan design through the selection of appropriate rough and detailed design plans. In this example, only the microwave discrete processing design plans are shown. Additional plans which might be appropriate for other processing technologies (such as autoclave processing or pultrusion) would appear under the top-level specialist.

Each specialist is responsible for some element of the design. For example, the "MD process structure" specialist is responsible for decisions regarding the policies which will be used to implement mode switching and power control. The tasks under this specialist choose the strategies which are most appropriate. For example, static, scheduled mode switching is appropriate if temperature sensors are unavailable but the mode interactions with the part (heating patterns) are well-understood. Similarly, reactive, adaptive mode switching is appropriate if little is known about the heating patterns, but temperature information is available.

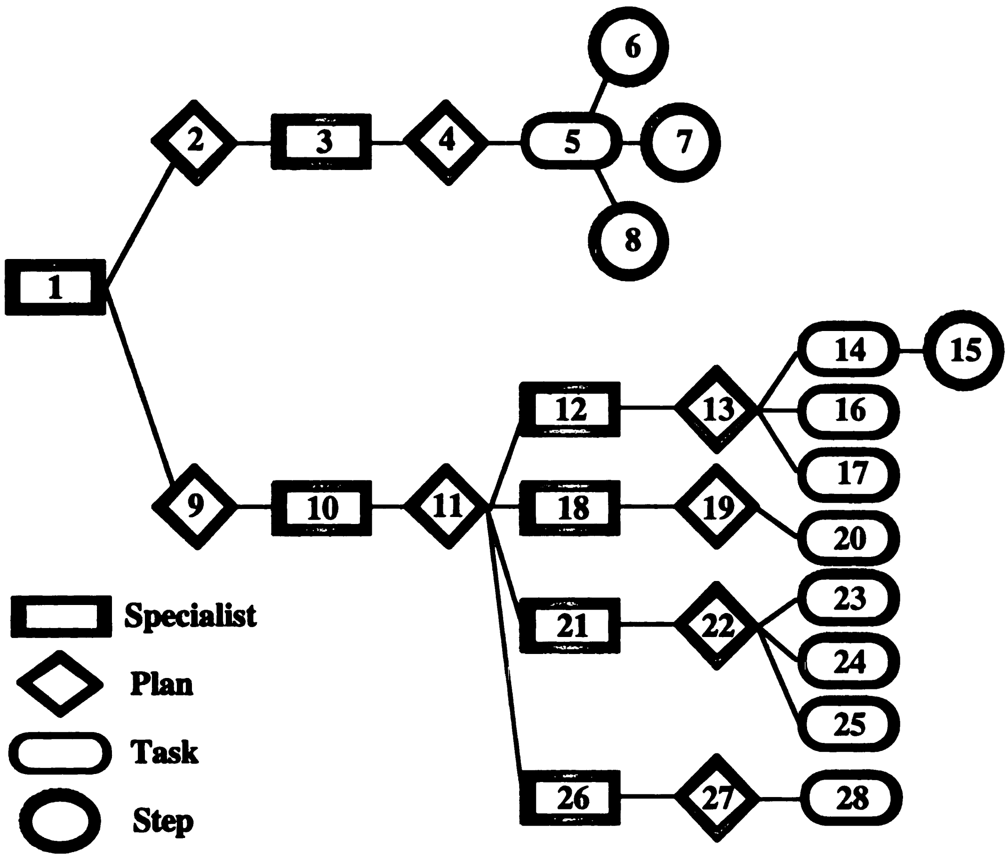

1. Fabrication technology specialist
2. MD rough plan
3. MD process structure
4. MD rough plan structure
5. Set process policies
6. Fabrication technology
7. Mode control policy
8. Power control policy
9. MD detail plan
10. MD detail specialist
11.Thermo-power-timing plan
12. Thermo specialist
13. Thermo plan
14. T-deg temp
15. Set T-deg temp
16. Max temp
17. Min temp
18. Power specialist
19. Power plan
20. Set recommended power
21. Static processing
22. Schedules
23. Processing time
24. Mode switch schedule
25. Power schedule
26. Reactive processing
27.Completion events
28. Time limits

Figure 4.5 - Structure of IPCA Design Specialists, Plans and Tasks

The specialist hierarchy used within the Blueprint Generator coordinates the problem-solving activities needed to fill out the detail in the blueprint, including knowledge about the utility of various control methods for maintaining tuned modes, appropriate power levels and schedules for movement between processing phases. The overall processing strategy is selected by the fabrication specialist and subsequently influenced by information about sensor availability, processing experience, initial material state, material type and geometry. Other agents determine the EM field types which are best suited to the initial material setting requirements and the criteria for making transitions between the processing phases. Each design decision adds a key/value pair to the blueprint. For example, if reactive, adaptive mode switching is chosen, the pair *(ModeSwitchPolicy MWDiscreteModeSwitchPolicyC)* is added to the blueprint.

This naturally leads to the question of how knowledge about the appropriateness of DSPL plans, tasks and steps is represented. This representational mechanism, called a sponsor-selector system, is described in the next section.

4.5.2 Sponsor-Selector Mechanisms

IPCA uses a sponsor-selector mechanism to support knowledge about the applicability of agents under various conditions [Punch 95]. The mechanism is used for two distinct purposes within the architecture:

- **Within the DSPL problem-solver to represent the appropriateness of specialists, plans, tasks and steps.**
- **Within the generated process control plan to represent the appropriateness of fragments.**

For both of these purposes, the mechanism functions identically. A sponsor-selector mechanism consists of a hierarchy of selectors, sponsors and selection items. Each selector may be associated with any number of sponsors, which are grouped together under the selector. Each sponsor, however, can represent only a single selection item. In multi-level sponsor-selector hierarchies, selection items are sometimes other selectors.

Figure 4.6 shows an example sponsor-selector structure containing two levels. The top-level selector chooses from a set of three sponsors, one of which has another selector as its selection item. This second-level selector, in turn, chooses from a set of three sponsors. Each of these sponsors has one associated selection item.

Selectors effectively choose from among rating values returned by the sponsors to determine which selection items are most appropriate for use under the current problem-solving conditions. The selectors also contain a mechanism for deciding which sponsor is to receive preference in the event of a tie. In IPCA, this mechanism is implemented by ordering the sponsors by priority, but many other implementations are possible, including breaking ties randomly.

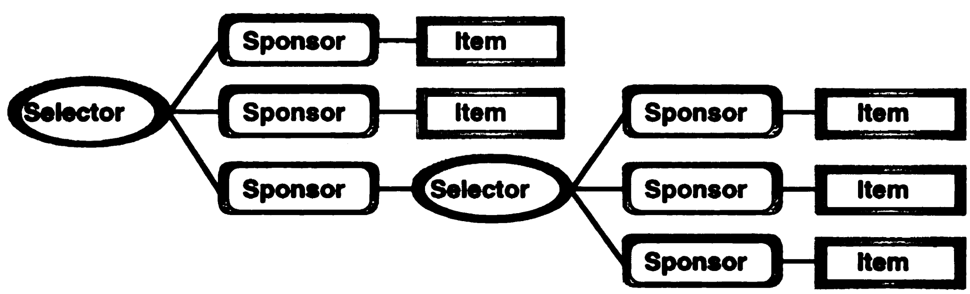

Figure 4.6 - Sponsor-Selector Structure

Each sponsor contains knowledge about the appropriateness of a single selection item. Though it can be done in any appropriate manner, this implementation encodes knowledge in a **match-group (MG) table**, as depicted in figure 4.7.

X < 0	Y >= 0	Z < 1	Result
true	true	true	A
true	false	*	C
false	true	false	B
false	*	*	C
false	false	false	D

Figure 4.7 - Match Group Table Structure

An MG table consists of one row of column header expressions, followed by one or more rows of Booleans and don't-cares forming the body of the table. Every MG table contains at least two columns. The right-most column always contains the result to be returned if conditions cause the corresponding row to be matched. The set of possible values that these result values can assume is called a **rating system**, which will be discussed in chapter 5.

When evaluating an MG table, the column header expression truth values are determined first (ignoring the last column, which is always labelled "Result"). Rows in the table are examined to determine which row(s) match the values determined for the header expressions. Since don't-cares may be used (indicated by a "*" in figure 4.7), multiple initial row matches are possible. The most specific (fewest don't-cares) matching row is selected from this group and it's result value is returned. For example, using the MG table from figure 4.7, assume X=3, Y=1 and Z=2. The column-header expressions evaluate to false, true, false. This pattern matches both rows four and five of the six-row table. Row four, however, is more specific because it contains fewer don't-cares. Therefore, the result for row four (the symbol "B") is returned. Also, note that some values of X, Y and Z will cause no row to be matched. In such cases, a default result is returned.

Selection item characteristics are not restricted by the sponsor-selector architecture, and in fact many strategies have been encoded through the use of this mechanism (e.g., DSPL plans [Brown 84], hypotheses [Punch 90], IPCA control methods [Decker 94], etc.). In each case, the sponsors act as "local" measures of how appropriate their associated items are for achieving the current goals, while the selector takes the more "global" view of choosing which of the fragments is the most appropriate under the given problem-solving situation.

CHAPTER V

5.1 Architectural Overview

This chapter contains an architectural overview of IPCA. Implementation details are provided in appendix B.

The IPCA architecture exploits the benefits of the Generic Task methodology by integrating established approaches to knowledge-based design and planning with real-time plan execution monitoring. The plan generation component is based on the DSPL language for routine design and planning [Brown 89]. The plan execution monitoring component is based on another problem-solving mechanism termed a sponsor-selector system [Punch 89].

Figure 5.1 shows the detailed system architecture in the context of controlling a microwave composite curing process, the present laboratory system on which the IPCA implementation is being evaluated. **Fabrication models** are used by IPCA during a **plan generation** step, wherein a **blueprint** is generated and a **process control plan** is subsequently assembled from **fragments** contained in a **fragment library**. The resulting process control plan is then applied in real-time by a **plan execution monitoring** facility. During this plan execution monitoring activity, the IPCA system interacts with the controlled process, obtaining sensor information and producing control actions.

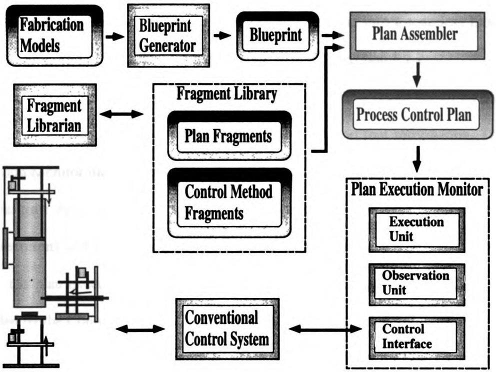

Figure 5.1 - System Architecture

Process control plan generation and execution, though integrated in the overall system operation, are activities supported by separate problem-solving modules, each with its particular knowledge representations and control strategies.

The **Blueprint Generator** creates a plan blueprint which specifies processing goals and potentially applicable fragments to be used during fabrication. Every fabrication model contains information about a specific process which is used to produce a unique blueprint. Blueprints are specific to the process described in their corresponding fabrication model. Further detail is provided in section 5.2.1.

The **Plan Assembler** integrates the plan and control method fragments specified in the blueprint and creates the final process control plan by retrieving from the fragment library those methods called for in the blueprint. The Plan Assembler also creates sponsoring mechanisms for the specified fragments, and integrates them into a seamless fabrication plan structure. The Plan Assembler is described in detail in section 5.2.2.

The responsibilities of the **Plan Execution Monitor** are to gather sensor data in real-time and to select appropriate fragment sequences based on that information. The Plan Execution Monitor therefore consists of three interacting components: an observation unit (for real-time data capture and analysis), an execution unit (for plan fragment selection and execution) and a Control Interface Unit for data acquisition. The Observation Unit within the plan execution monitor gathers sensor data using network interface services provided by the Control Interface Unit. The Control Interface Unit communicates with an External Controller through a network interface. A description of the Plan Execution Monitor is given in section 5.3.

The **External Controller** is responsible for monitoring a collection of sensors which measure various aspects of the material state during fabrication and for reporting these values to the Control Interface Unit when requested. Currently, sensed parameters include temperature, incident and reflected power and physical positioning information for the motor-controlled cavity tuning mechanisms. The External Controller also communicates control outputs directly to physical actuators. The roles of the External Controller are described in section 5.4 and further detail is available in appendix D.

5.2 The Plan Generator

The Plan Generator is comprised of three interacting components. The **Blueprint Generator** is described first, followed by the **Plan Assembler** and the **Fragment Librarian**.

The information processing task of the plan generator is to produce process control plans based on information about the process to be controlled (provided by fabrication models), knowledge of various fabrication processes (contained within the blueprint generator) and a library of pre-compiled fragments (contained within the fragment library).

5.2.1 The Blueprint Generator

The Blueprint Generator produces blueprints by utilizing knowledge incorporated in fabrication models and compiled domain-specific knowledge of composite material fabrication, thermodynamics and heat transfer. Blueprints (see table 5.1) consist of a set of references to potentially applicable fragments and the conditions under which they can be executed. As suggested by table 5.1, the blueprint is implemented as a data structure containing a number of unordered key/value pairs.

The final process control plan may be completely reactive in situations where a full complement of sensors is available but little or no experience has been gained in processing similar structures or materials. Conversely, process control plans may also be completely static in cases where there are few or no sensors but previous processing experience is available. Typically, however, process control plans contain plan fragments of both types, and this information is reflected in the key/value pairs contained in each blueprint.

Key	Value	Key	Value
PartMass	30	PlyType	Crossply
Matrix	DGEBPA	Plies	12
Technology	MicrowaveDiscrete	ModeControlPolicy	ReactiveAdaptive
PowerControlPol-icy	Reactive Non-Adaptive	PhaseChangePolicy	Event-TimeBased
TargetTemp	155	TDegTemp	180
Modes	1-5	PreviousRuns	5
TargetGrad	5	HoldTime	45

Table 5.1 - Typical Blueprint Format

5.2.2 The Plan Assembler and Fragment Librarian

The Plan Assembler generates a sponsor for every fragment specified in the blueprint, encoding in each sponsor the knowledge of when that fragment is appropriate. Each key/value pair contained in the blueprint is examined and matched by the Plan Assembler so that appropriate fragments are included in the final process control plan. However, note that the resulting process control plan is not "flat" but consists of a hierarchy of fragments. Thus the Plan Assembler arranges fragments into an appropriate hierarchical configuration. The Plan Assembler generates a sponsor for each of the fragments retrieved by the Fragment Librarian for inclusion in the process control plan, placing it into the process control plan hierarchy under the appropriate selector. The resulting plan structure is in a format utilizable by the knowledge-based selection system contained within the real-time plan execution monitor.

All fragments are managed by a Fragment Librarian. The Fragment Librarian provides a mechanism for organizing plan and control method fragments into specific libraries. The Fragment Librarian also provides library retrieval services to the Plan Assembler. Other elements within the same fragment contain lists of various parameters used by the fragment, dependencies on other fragments and methods for computing required parameters. Additional information regarding fragments is given in section 5.2.3.

The Plan Assembler is designed to examine all the key/value pairs contained in a blueprint and to request appropriate fragments from the Fragment Librarian. The Plan Assembler combines the required fragments into a process control plan by sponsoring each fragment and placing it within the process control plan hierarchy at an appropriate position. The Plan Assembler utilizes information about dependencies between fragments to ensure that all required fragments are included in the process control plan and that appropriate parameters are created and initialized.

It is important that all fragments referenced in the blueprint exist in the fragment library at assembly time. In the event that the blueprint contains a key/value pair which is not found in the fragment library, an error is generated and the assembly operation will end in failure.

The typical (simplified) structure of a process control plan is shown in figure 5.2. The process control plan consists of several structures including a root-level selector, some number of fragment sponsors, and a fragment associated with each sponsor. Associated with each group of sponsors is a single selector which can be used to resolve conflicts that can arise when multiple fragments appear equally appropriate. The structure of the process control plan reflects the nature of the control task in the composite material processing domain. The control functions are decentralized, with, for example, power control and mode switching responsibilities residing in different sponsors. This structure reflects both the modularity of the IPCA fragment library and the nature of the controlled process.

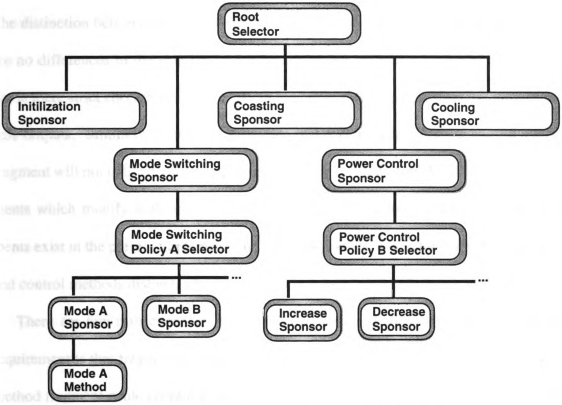

Figure 5.2 - Typical Process Control Plan Structure

5.2.3 Fragments

Two types of fragments are contained in the fragment library: plan fragments and control method fragments. The fragment types are identical except for their effects when executed by the plan execution monitor. Plan fragments are distinguished by their direct influence over processing goals, as detailed in section 5.2.5. When a plan fragment is executed, processing goals may be directly modified as a result. Plan fragment execution does not directly modify any other process parameters, such as outputs, although the modification of processing goals may subsequently influence particular control method fragments to be executed, which might subsequently change control outputs. Control method fragments directly control process outputs when executed but they do not modify goals. The distinction between fragment types is made only for ontological convenience. There are no differences in the way the fragments are stored, selected, sponsored or executed, nor is there strict checking to ensure that an executing plan fragment does not modify process outputs. Similarly, there is no guarantee that executing a particular control method fragment will not indirectly modify processing goals. It is therefore possible to create fragments which modify both processing goals and control outputs, although no such fragments exist in the present implementation. This dichotomy was enforced to keep planning and control methods distinct, which helps ensure greater generality and reusability.

There are few restrictions on the algorithms contained within fragments. The only requirement is that fragments contain compilable SmallTalk source code. Thus, a control method fragment could contain a fuzzy-logic-based control algorithm, another could contain a complete neural network, a third could contain a PID control algorithm and so on. Neither are there restrictions placed on combining fragments of many different types within the same process control plan. Thus, some control method fragments could be based on PID algorithms while others could be based on simple setpoint control. Similarly, there are few restrictions on plan fragments. Each plan fragment can be constructed to act on processing goals using any available algorithm. Thus, a plan fragment could contain a complete STRIPS planner implementation, a reactive planning algorithm, a simple heuristic and so on. This approach adds flexibility and ensures that process control plans are not restricted to a single planning or control approach.

5.2.3.1 Fragment Library Content

Currently, the Fragment Librarian contains a variety of static, reactive and adaptive control and planning fragments which are used for microwave composites processing, including single and multiple setpoint control methods for managing power, proportional control methods, and a variety of qualitative state-based methods for mode selection, cure maintenance and fabrication planning. The IPCA control method fragment content is summarized in table 5.2.

The IPCA power control policies evolved as experience was gained from fabrication experiments. The earlier algorithms such as single setpoint control were supplemented by increasingly sophisticated versions for all but material heating applications, which still use the single setpoint algorithm to ensure a preset power level. The various mode control policies are used under different fabrication conditions. For example, plans for gathering heating profiles for a new part utilize MDModeControlPolicyE, which provides a static, timed schedule which cycles through each of the available modes. Standard reactive mode switching, adaptive-reactive switching, timed static switching are all available and their selection depends upon sensor availability, previous processing experience and other factors.

Library Name	Fragment Name	Notes
PowerControlPolicy	MDPowerControlPolicyA	Single setpoint (on/off) power control used mainly for heating trials
	MDPowerControlPolicyB	Multi-level setpoint control. Power sets back 10% at each excursion.
	MDPowerControlPolicyC	Proportional power control
ModeControlPolicy	MDModeControlPolicyA	Static, timed mode switching - no probes required.
	MDModeControlPolicyB	Reactive mode switching with hold time
	MDModeControlPolicyC	Reactive-adaptive mode switching with hold time
	MDModeControlPolicyD	Reactive mode switching using uniform energy strategy, with hold time
	MDModeControlPolicyE	Used for heating trials - cycles through all modes
EmergencyPolicy	MDEmergencyPolicyA	Shut down if any problems

Table 5.2 - Control Method Fragments in the IPCA Fragment Library

The IPCA fragment library also contains a number of useful planning methods for microwave-based polymeric composite fabrication. The IPCA plan fragment content is summarized in table 5.3.

r

a

c

th

5

IP

vi

Library Name	Fragment Name	Notes
InitilizationPolicy	MDInitilizationPolicyA	MW discrete processing
Setup0001	Master (&Mode1-5)	3" square DGEBPA modes for experimental series F1
Setup0002	Master (&Mode1-4)	1" square DGEBPA modes for experimental series F2
Setup0003	Master (&Mode1-4)	2" square DGEBPA modes for experimental series F3
Setup0004	Master (&Mode1-3)	2" square DGEBPA com-plimentary modes for experimental series F4
FabricationTechnology	MicrowaveDiscrete	Fabrication planner for fixed-time cure
PhaseChangePolicy	MDPhaseChangePolicyA	Maintains processing goals (e.g. warmup com-plete when MinProcess-ingTemp reached)

Table 5.3 - Plan Fragments in the IPCA Fragment Library

Setup0001 through Setup0004 contain fragments for a Master and associated mode selection fragments, one per mode to be used during processing. The applicability of each mode is included within the appropriate mode fragment. Since one fabrication technology and material is currently supported by IPCA, there is one phase change policy (material dependant) and one initialization policy (fabrication technology dependant) contained in the plan fragment library.

5.2.3.2 Adaptive Fragments

IPCA contains a facility for supporting adaptive fragments. An access mechanism is provided to allow fragments to modify the contents of the fragment library during process

co
th
tal
m
ex

m
du
m
m
tiv
ma
m
pa
sta
ity
pr

the
me

5.2

Th
Ge

control plan execution. The access mechanism is provided by a pair of messages to which the fragment library will respond. One message provides query access to the matchgroup tables within any fragment in the library. The second message allows matchgroup elements to be replaced with new values. Together, the messages provide read and write execution-time access to the library contents.

An adaptive mode switching fragment has been developed and tested. This fragment modifies the estimated applicability entries of mode fragments based on their performance during execution. The MDModeSwitchPolicyB and MDModeSwitchPolicyC control method fragments estimate the qualitative state of the sample and select a processing mode which is most appropriate. Whenever a mode is selected (under the adaptive-reactive mode switch policy, MDModeControlPolicyC) which causes an increase in the thermal gradient, the rating for the mode in that qualitative state is weakened by the adaptive mode switching algorithm. Similarly, if the gradient is reduced as a result of selecting a particular mode, the rating for that mode is strengthened for the appropriate qualitative state and mode. Through the use of this adaptive mode switching method, the applicability of selected modes in various qualitative states are automatically refined during material processing.

For detail regarding mode fragment generation, see appendix C. For details regarding the adaptive mode switching fragment, see "MDModeControlPolicyC" in the IPCA fragment library.

5.2.4 Process Control Plan Format

The process control plan takes the form of a pair of ASCII files created by the Blueprint Generator. The **plan file** contains information about the process control plan structure, match groups used for determining sponsor applicability and the SmallTalk code for all the fragments included in the process control plan. The **parameter file** contains declarations for the parameters used by the process control plan, including names, types, initial values, mappings to physical hardware and methods for computing abstract parameter values. A sample process control plan and parameter file are given in appendix A.

One portion of a fragment from a generated process control plan is shown in textual form in figure 5.3. A typical process control plan contains many fragments, most of which have more elaborate control method algorithms than the one included this example. However, the example clearly shows the sponsor (emergencySponsor), the matchgroup table (with column headers, wildcards and ratings), a default row match return value, and the procedural body of the method.

```
"Sponsor for emergency conditions: applicable if part is too hot"
Sponsor: emergencySponsor.
Headers: ('T1Sensor > ThermalDegradationTemp' 'T2Sensor >
ThermalDegradationTemp' 'T3Sensor > ThermalDegradationTemp' 'T4Sensor >
ThermalDegradationTemp' 'PartTemp > ThermalDegradationTemp' 'Rating').
Contents: (
        (true * * * * 'APlus')
        (* true * * * 'APlus')
        (* * true * * 'APlus')
        (* * * true * 'APlus')
        (* * * * true 'APlus')
).
DefaultReturnValue: 'E'.
Item: (method partTooHotMethod).
MethodName: partTooHotMethod.
Method: '
partTooHotMethod
        Transcript show: "Thermal degradation limit exceeded".
        PowerActuator1 := 0.
'.
```

Figure 5.3 - Example Fragment from MDEmergencyPolicyA

the

is s

mat

is t

then

The

upd

tual

to a

able

but

para

calc

obta

meth

inclu

"Pov

para

curre

5.2.5

IPCA

The process control plan fragment shown switches off the microwave power source if the sample temperature exceeds a preset limit. The value of "ThermalDegradationTemp" is set by the Blueprint Generator and included in the plan blueprint as a parameter. The match group table indicates that the method is applicable (with a rating of "APlus" which is the highest rating available) when the temperature of any part of the sample exceeds a thermal degradation limit.

All four types of IPCA parameters are present in this simple control method fragment. The parameter "T1Sensor" is an **input** parameter. The values of all input parameters are updated at the beginning of each execution cycle. Each input parameter is assigned a virtual channel by the Blueprint Generator and Plan Assembler, which is eventually mapped to a physical sensor. The parameter "ThermalDegradationTemp" is an example of a **variable** parameter. These parameters have an initial value (set by the Blueprint Generator) but the value does not change unless modified as the result of fragment execution. The parameter "PartTemp" is an **abstract** parameter. Abstract parameters have their values recalculated before every execution cycle. Immediately after new input parameters are obtained, the Observation Unit updates the values of all abstract parameters by invoking methods included in the parameter file. SmallTalk code implementing these methods is included in the parameter file by the Plan Assembler. Finally, the parameter "PowerActuator1" is an example of an **output** parameter. Output parameters, like input parameters, are assigned to virtual channels. At the completion of each control cycle, the current values of all output parameters are routed to the appropriate physical actuators.

5.2.5 IPCA Goals

IPCA goals are defined and initialized by the Blueprint Generator and maintained by the var

are

ma

ete

par

tin

by

ma

pai

dis

ure

ma

que

a g

var

He

tial

pla

crea

izes

various planning fragments found within the generated process control plans. IPCA goals are not supported by any specialized framework within the architecture, but are simply maintained by the Plan Execution Monitor in the same manner as all other process parameters. The Plan Execution Monitor provides a generic facility for maintaining all process parameters, including goals. Therefore, there is no architectural or implementational distinction between goals and variable parameters. The distinction between goals (modified by plan fragments) and variable parameters (modified by control method fragments) is made for notational convenience.

Initial goals for processing are setup by the Blueprint Generator. For example, the pair *(ProcessingGoal InitialMaterialHeating)* is added to the blueprint for all microwave discrete processing by a particular step under the Thermo-power-timing plan (refer to figure 4.5). The pair *(MinProcessingTemp 125)* might also be added to a blueprint (this is material dependant) by a step under the Thermo plan. Although both tuples are subsequently treated as variable parameters by IPCA, the first would generally be referred to as a goal and the second as a setpoint.

When the resulting blueprint is subsequently assembled, the Plan Assembler creates a variable parameter named *ProcessingGoal* and fills in an initial value of *InitialMaterialHeating*. Similarly, a variable parameter named *MinProcessingTemp* is created and is initialized to a value of *125*. The plan fragments included in the assembled process control plan will presumably modify the variable parameter *ProcessingGoal* during execution.

Before the assembled process control plan is executed, the Plan Execution Monitor creates a variable parameters named *ProcessingGoal* and *MinProcessingTemp* and initializes their values to *InitialMaterialHeating* and *125*, respectively. During the execution of the process control plan, a plan fragment is eventually selected when the average sample temperature exceeds *MinProcessingTemp*. The purpose of the plan fragment is to watch for the event which signals that the sample has reached a minimum processing temperature and to change *ProcessingGoal* to reflect this fact. When the selected plan fragment is executed, it changes the value of *ProcessingGoal* to *MaintainMinimalGradient*, subsequently affecting the selection and operation of other plan and control method fragments.

5.3 The Plan Execution Monitor

The plan execution monitor is composed of three interacting components. The **Execution Unit** described in section 5.3.1 selects the plan fragment sequence in real-time. The **Observation Unit** described in section 5.3.2 processes sensed process information and provides abstractions from sensor data gathered during real-time process monitoring. The **Control Interface Unit**, which provides network access to the External Controller, is described in section 5.3.3.

5.3.1 The Execution Unit

The Execution Unit applies the process control plan by selecting fragments for execution based on data gathered by the Observation Unit. The Execution Unit must coordinate execution of the process control plan with sensing and control functions. The External Controller relieves the execution unit of the need to supervise low-level control loops involved with cavity tuning operations.

Upon initialization, the Execution Unit parses the selected process control plan and instantiates the sponsor-selector mechanisms to be used during plan fragment selection. The Execution Unit pre-compiles all the fragment code contained within the process control plan, so that fragments are ready to execute immediately upon selection.

Execution of the processes control plan is cyclic. At the beginning of each execution cycle, the Execution Unit requests the latest sensed process parameters from the Observation Unit. The Execution Unit then evaluates all the sponsors of the root-level selector and gathers their appropriateness measures for their associated fragments. The root-level selector then chooses a winning fragment based on the sponsor values and, if necessary, the selector priority list. The process control plan structure may be of considerable depth, since the selected fragment may be another selector. The search through the process control plan proceeds best-first until a single plan or control method fragment is reached. The winning fragment is then executed in a protected mode as a background process so that any run-time errors which might result do not interfere with the execution of the overall process control plan. The Execution Unit kills the background process if an error or timeout occurs. The executing fragment may make adjustments to the controlled process parameters or change the processing goals. In either case, the execution cycle repeats immediately after the fragment finishes execution or a timeout or error occurs. Plan execution continues indefinitely until suspended by an executing fragment or by the operator.

5.3.2 The Observation Unit

The Observation Unit communicates directly with the Control Interface Unit to obtain input parameter values. The Observation Unit also manages several other types of process parameters and performs analysis during process control plan execution. Process parameters are used in conjunction with the process control plan to effectively control the fabrication process through appropriate selection of sponsored plan and control method fragments.

As introduced in section 5.2.4, there are four types of parameters maintained by the Observation Unit: input, output, abstract and variable. Their relationships to the Observation Unit will now be explored in more detail.

Input parameters are those values sensed by the External Controller and subsequently passed to the Observation Unit through the Control Interface Unit. **Output** parameters are modified by control method fragments and are passed to the external control system through the same network interface and Control Interface Unit services. **Abstract** parameters have their values calculated once per execution cycle according to computational algorithms included in the process control plan parameter file. These values are not communicated to the Control Interface Unit or External Controller. Finally, **variable** parameters are used to maintain all other values required for fragment execution. Typically, goals and global variables are encoded as variable parameters and modified during fragment execution. Variable parameters are used for many purposes, including plan goal maintenance and providing global data structures for fragment use. A variable parameter can contain any SmallTalk data type and is therefore a highly flexible and completely adaptable facility for persistent storage.

Before each execution cycle begins, the Observation Unit requests the most recent sampled sensor data from the Control Interface Unit. The requested information is normally transmitted by the dedicated External Controller to the Control Interface Unit through the services of a network interface. From these sensed parameters, the Observation Unit computes a set of abstract parameters by invoking specific computational algorithms included in the process control plan. The specifications for the abstract parameters required for a given process and the algorithms used to compute them are provided within the process control plan parameter file. Some examples of typical abstract parameter computations for the microwave fabrication process are elapsed time, average temperature and thermal gradient. Abstract parameters are computed once per control cycle prior to initiation of the selection process by the Plan Execution Monitor.

All sensed parameters are timestamped by the Observation Unit so that recency of data can be determined during fragment execution.

5.3.3 The Control Interface Unit

The Control Interface Unit is used to provide a standard virtual interface between the Execution Unit, Observation Unit and the External Controller. All plan and control method fragments are designed to address input and output parameters on virtual channels so that they are not bound to the details of how physical equipment is configured or how I/O channels are assigned. The Control Interface Unit maps virtual channel designations into physical sensor and actuator channels needed to access the hardware.

As a secondary function, the Control Interface Unit provides data logging services that allow process inputs and outputs to be captured for later playback or analysis. Also, the Control Interface Unit provides switching between multiple physical hardware configurations and the simulated control unit interface used for testing plan and control method fragments. The simulated control unit is described in section 5.5 below.

5.4 External Controllers

An **External Controller** is used to provide facilities for connecting physical sensors, actuators and processing equipment to the IPCA prototype. The architecture and implementation of external controllers are not dictated by IPCA, since the structure and duties of each external controller are domain and process-specific. Many external controllers may be used, potentially accommodating several different processes. Each external controller provides a conduit between the Control Interface and the processing hardware.

Although the internal design and implementation of external controllers are not specified by the IPCA architecture, certain duties are expected of every external controller. Each external controller is expected to monitor available physical sensors and to report sensor values to the IPCA Control Interface on a timely, periodic, synchronous basis. Similarly, each external controller is expected to accept timely, periodic, synchronous actuator outputs from the IPCA Control Interface and to pass these values to the appropriate physical actuators. These responsibilities are described in detail in appendix D, which contains a protocol document detailing the mechanism of information exchange.

External controllers are generally designed to reduce the computational load on the IPCA Plan Execution Monitor by providing local, low-level control operations and by providing sensor preconditioning facilities. The external controller implemented for use with the IPCA prototype contains these features, using sensor preconditioning to reduce the effective sensor sampling rate and PID control algorithms to facilitate mode tuning operations.

External controllers typically execute on dedicated hardware and communicate with IPCA through the use of a shared memory facility. In the present implementation, the external controller executes on a Macintosh IIci™ and the IPCA prototype executes on a Macintosh PowerBook™ 180. A network interface provides mutual access to a shared memory area, in this case, a RAM disk. Messages are posted and retrieved from this shared medium to facilitate communications, as shown in figure 5.4. Additional details re

5.

IPo

wi

pro

act

sin

pla

sett

inpu

regarding the format of exchanged information may be found in appendix D.

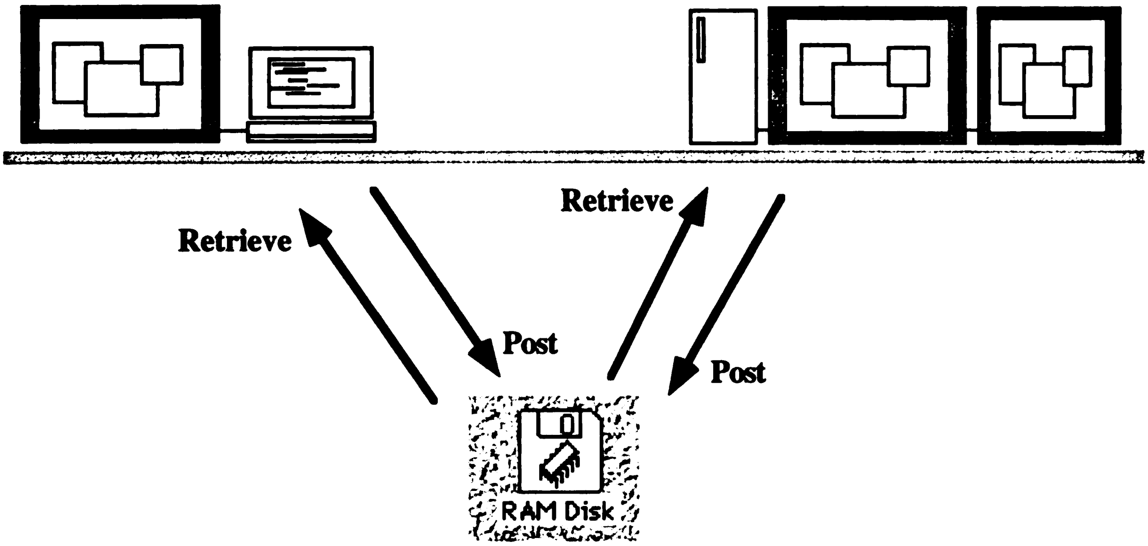

Figure 5.4 - Communication with an External Controller

5.5 The Simulated Control Unit

IPCA includes a simulated control unit which allows the execution monitor to function without need for an external controller. The simulated control unit is useful for testing process control plan reactions to a variety of exceptional conditions (such as sensor and actuator failures) which would be difficult or impossible to otherwise demonstrate. The simulated control unit is also useful for verifying the proper operation of process control plans before they are connected to physical equipment and control systems in laboratory settings.

The simulated control unit allows process control plans to be executed while process inputs and outputs are routed normally through the Control Interface Unit. From the viewpoint of the Execution and Observation Units, the Control Interface Unit is connected to a generic source of process data. The simulated control unit features 16 virtual sensors and 8 virtual actuators which can be used to simulate up to 16 input channels and 8 output channels.

For implementation details regarding the simulated control unit, see appendix B.

CHAPTER VI

6.1 Experimental Scope and Objectives

The functionality, utility, domain applicability and performance of the IPCA prototype has been evaluated through a collection of experiments. The experiments are of two types: **C-series** and **F-series**. Both types of experiments are introduced and an overview of the scope and objectives of each experimental series is presented. The criteria used to evaluate experiments from each series is presented in the final section.

6.1.1 C-Series Experiments

The C-series experiments are designed to demonstrate the breadth and capability of the IPCA prototype Plan Generator and Plan Assembler by using the Simulated Control Unit described in section 5.5. C-series experiments are conducted by presenting specific fabrication models for blueprint generation and plan assembly. The resulting blueprint and/or process control plan are examined for fragment content. Process control plans may be exercised on the simulated control unit to demonstrate capabilities under simulated conditions. C-series experiments do not involve real-world processing hardware, but are designed to demonstrate the following capabilities:

- **The Blueprint Generator and Plan Assembler generate executable process control plan structures for various fabrication scenarios (series C1).**

- **The range of process control plans produced is sufficiently broad, including reasonable material coverage (series C2).**

- **Process control plans respond effectively to potentially dangerous or hazardous conditions (such as emergency handling) which would be difficult to observe in the laboratory (series C3).**

6.1.2 F-Series Experiments

F-series experiments are designed to verify the viability of IPCA prototype and the applicability of the planning and control method fragments to real-world laboratory-scale composite structure fabrication. The F-series experiments are used to demonstrate the following capabilities:

- **Microwave-based fabrication processes can be effectively managed by IPCA-generated process control plans.**

- **Quality composite structures can be rapidly fabricated under a variety of conditions.**

- **Rapid, controlled heating is possible using small, low-mass samples.**

- **Thermal gradients during processing are lower than those associated with other control techniques. Gradients are dependant on a number of factors which are investigated in detail.**

- **IPCA effectively communicates with sensors, actuators and external control systems.**

- **The adaptive control methods are useful and effective. They improve their performance with experience.**

These capabilities are collectively demonstrated through the use of four experimental series featuring the same composite material, but different structural geometries.

6.2 Experimental Preliminaries: F-Series

This section addresses experimental considerations for the F-series experiments. The sample composition, geometry, layup, sensor placement, electromagnetic mode selection and characterization and mechanical limitations of experimental fabricator are presented in detail.

6.2.1 Sample Geometry and Composition

The composite material used in all F-series experiments was a carbon-fiber reinforced Di-Glycidyl Ester of Bis-Phenol A (DGEBPA) prepreg. All samples were hand-laid 12-ply unidirectional squares. Seven 3-inch square 12-ply samples were processed in experiments F1.1 through F1.7 from series F1. One 3-inch square 24-ply sample was processed in experiment F1.8. Eight 1-inch square 12-ply samples of a second geometry were processed in experiments F2.1 through F2.10. Three 2-inch square 12-ply samples were processed in experiments F3.1 through F3.3. Four 2-inch square 12-ply samples were processed in experiments F4.1 through F4.4. In each of the series F1 through F4, heating data was gathered using additional samples.

6.2.2 Temperature Probe Placement

All F-series experiments were conducted using a Luxtron™ four-channel fiber-optic thermometry system to measure sample temperatures at four sites during processing. The probe channels 1 through 4 are designated as T1, T2, T3 and T4 for all experimental series.

Each thin-section composite sample was processed using a 90 degree fiber orientation, indicating that the fibers in the unidirectional layups were perpendicular to the coupling probe axis. The fiber-optic temperature probes were inserted into glass capillary tubes designed to protect the tips from physical damage during processing. The protective tubes were inserted through 1/8" holes in the Teflon™ mold and rested directly on the upper release film (see layup diagram, chapter 3). Probe placement is the same for all experimental series, with T1 in the center of the layup, surrounded by T2, T3 and T4. A diagram of the fiber-optic temperature probe placement is shown in figure 6.1.

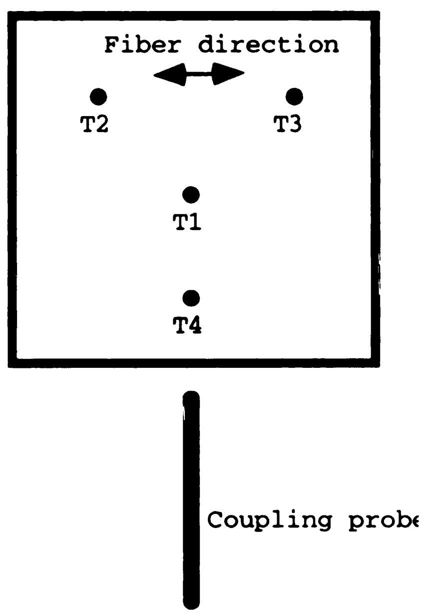

Figure 6.1 - Fiber-Optic Temperature Probe Placement

6.2.3 Electromagnetic Modes

Each fabrication experiment was conducted using a particular set of electromagnetic modes from which the Execution Monitor can ultimately select. Electromagnetic processing modes were initially identified using a cavity-sweeping procedure described in [Fellows 91]. The cavity-sweeping procedure involved loading the cavity with an appropriate

s
c
u
r
n
l
in

re
E
p
E
en
in
co

ri
pr
se
pr
ne
an
mo

sample (1", 2" or 3" square) and applying a low-power microwave-frequency signal to the coupling probe. The cavity length was changed while the reflected power was monitored using a dual-channel oscilloscope. Potential modes were identified based on minimized reflected power, indicating effective "tuned" EM coupling. In an effort to further minimize the reflected power, the probe depth was swept after each suitable cavity length was located. The cavity length and probe depth of each candidate mode were noted for further investigation.

Merely knowing the location of a mode does not give information about how the resulting EM field will interact with the sample. Since some modes are associated with EM field patterns which "surround" or "miss" the sample (particularly for very small samples), these modes were deemed as poor candidates for composite materials processing. Even modes having EM patterns which effectively heat the sample generally distribute energy into the material in an uneven fashion. Since this information is critical to making intelligent decisions about the utility of each mode, a series of heating experiments was conducted on each candidate mode using higher power.

The modes selected for the part geometries used in series F1 through F4 are summarized in table 6.1. The table includes a heating pattern for each mode, which lists the probe sites in order of decreasing EM coupling effectiveness. For example, mode 1 from series F1 tends to heat T2 the most, followed by T1, T3 and T4. The type of mode is also provided, and designations of **TE** (for transverse electric) and **TM** (for transverse magnetic) modes are used together with the number of standing half-wavelengths in the x,y and z planes to form the complete mode designation. For more information about EM modes, see [Herrington 61].

us

S
F
F
F
F-

seri

mat

at ex

Five modes were used during the F1-series fabrication runs, while four modes were used in series F2 and F3. Series F4 used three modes.

Series	Mode #	Mode Type	L_c(cm)	L_p(mm)	Heating pattern	Notes
F1	1	TE_{111}	13.5	21.5	T2, T1, T3, T4	Fast; Edge/Center
	2	TE_{311}	15.6	21.0	T3, T4, T1, T2	High Pr; Edge
	3	TE_{113}	18.5	21.0	T2, T1, T4, T3	Slow
	4	TM_{112}	21.1	21.0	T4, T1, T3, T2	Fast; Edge/Center
	5	TE_{213}	23.1	33.0	T3, T1, T4, T2	Even; slow
F2	1	TE_{211}	8.7	15.7	T1, T2, T3, T4	Very fast; low Pr
	2	TE_{011}	12.2	24.4	T1, T3, T4, T2	Highly favors T1
	3	TE_{212}	18.2	30.8	T1, T3, T4, T2	Even; med. rate
	4	TM_{013}	20.6	12.9	T3, T1, T4, T2	Even; slow rate
F3	1	TE_{012}	22.6	30.0	T2, T1, T3, T4	Ring; fast heating
	2	TE_{212}	17.2	27.2	T1, T2, T3, T4	Even
	3	TM_{012}	14.3	30.6	T1, T2, T3, T4	Fast; T1,T2 close
	4	TE_{211}	9.0	30.5	T2, T1, T4, T3	Even
F4	1	TM_{014}	27.4	28.0	T4, T1, T2, T3	Highly favors T4
	2	TE_{012}	22.9	28.0	T2, T4, T1, T3	Ring
	3	TM_{112}	21.7	28.0	T4, T2, T3, T1	T2, T4 close

Table 6.1 - Modes Used for Fabrication Series F1 - F4

A total of 15 candidate modes were identified for series F1, 13 for series F2, 8 for series F3 and 9 for series F4. The majority of these candidate modes were unsuitable for material processing either because they did not heat the sample effectively or were located at extreme cavity and/or probe positions which were difficult or impossible to tune due to physical clearance problems with the mold.

6.2.4 Heating Profiles: F-Series

Heating profiles were created for each candidate mode by applying power while simultaneously recording information from the fiber-optic temperature probes. Each profile reveals the manner in which EM energy is transferred into the sample, and thus reveals information about the utility of the mode during processing. Information from the resulting heating profiles was encoded into IPCA either manually (the method used for experimental series F1) or through the use of the IPCA mode map generator (used for experimental series F2 through F4).

The heating profiles for the experimental series F1 through F4 are given in figures 6.2 through 6.17. One set of profiles is given for each experimental series, since the profiles are dependant on sample geometry. Each profile reveals the behavior of each fiber-optic probe during the heating experiments for the sample in the respective mode. The profiles are generally divergent, as most modes preferentially heat one or more areas of the sample.

6.2.4.1 Series F1

Five modes were selected for use with the 3" square samples in the F1-series experiments. The heating profiles for modes 1 through 5 are given below in figures 6.2 through 6.6. The time scales for these heating profiles vary because the average heating rates in each mode are not identical.

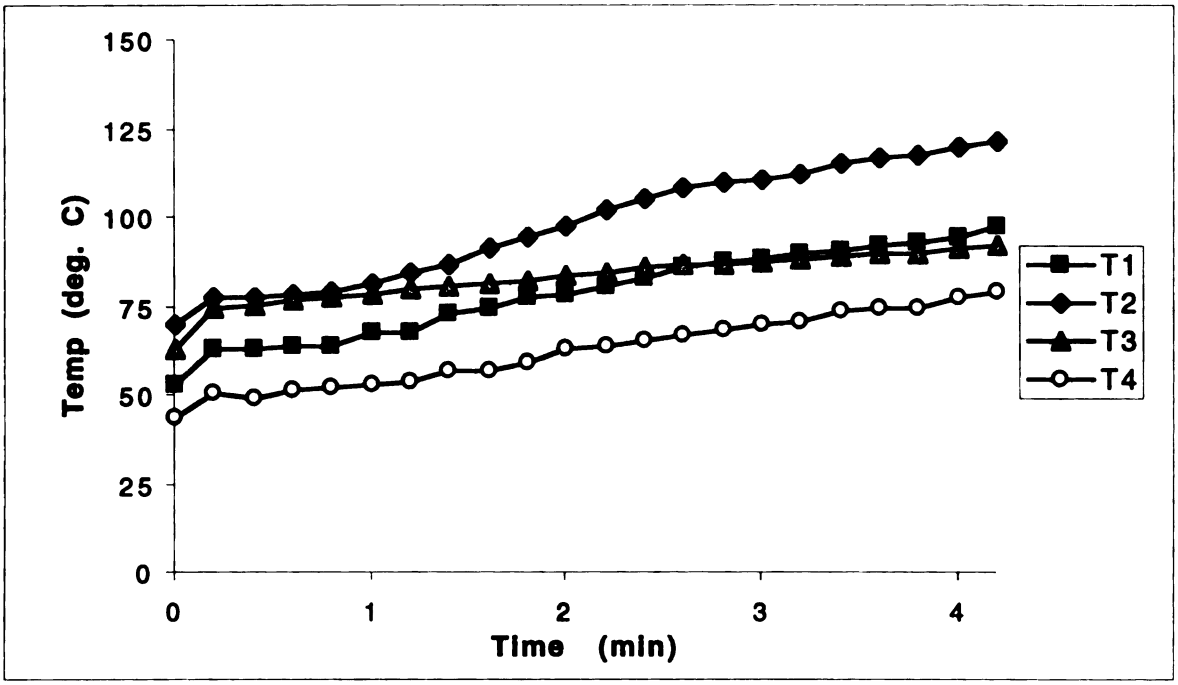

Figure 6.2 - Heating Profile: F1 (Mode 1)

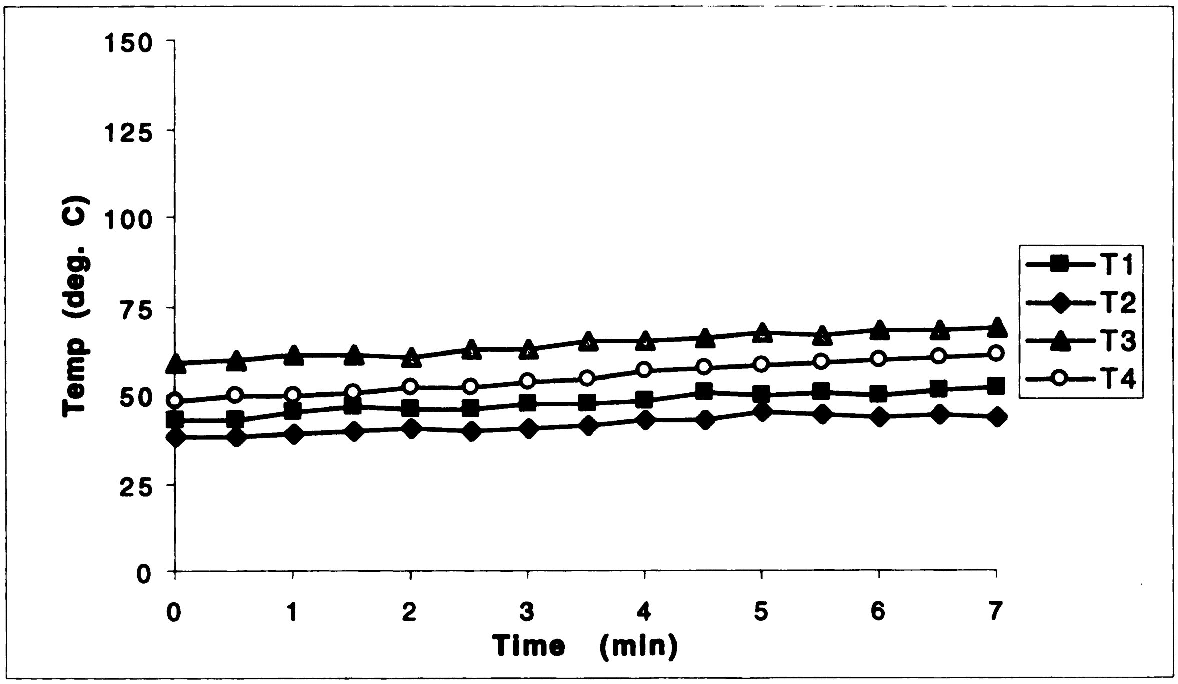

Figure 6.3 - Heating Profile: F1 (Mode 2)

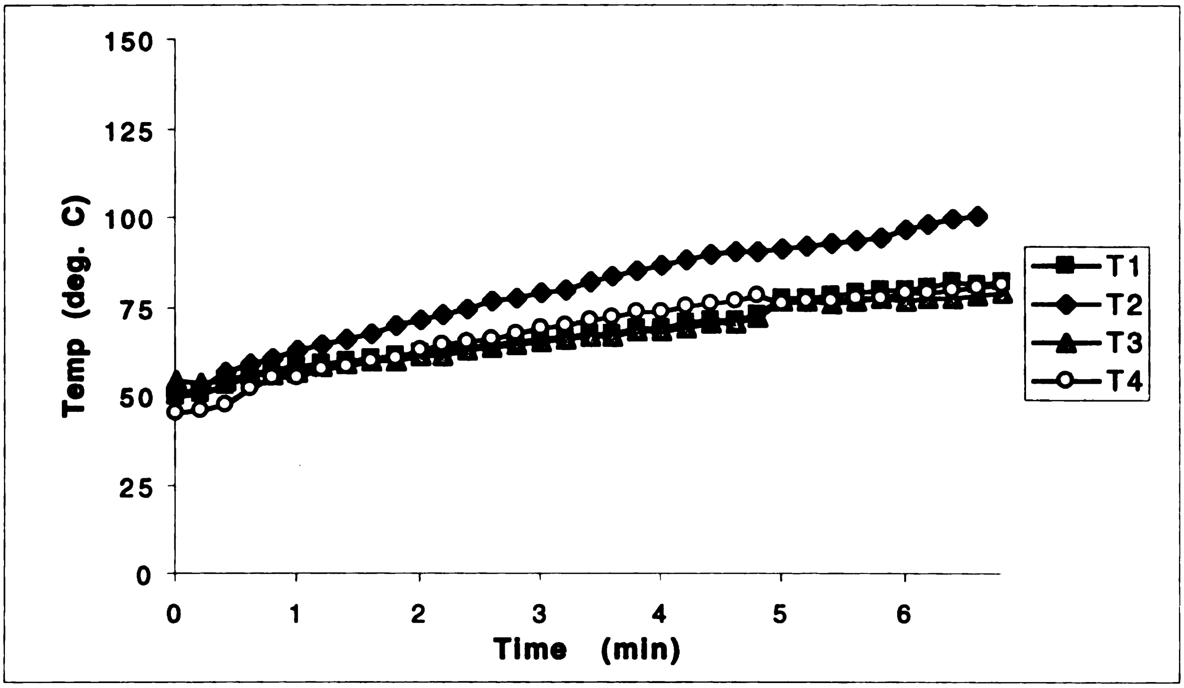

Figure 6.4 - Heating Profile: F1 (Mode 3)

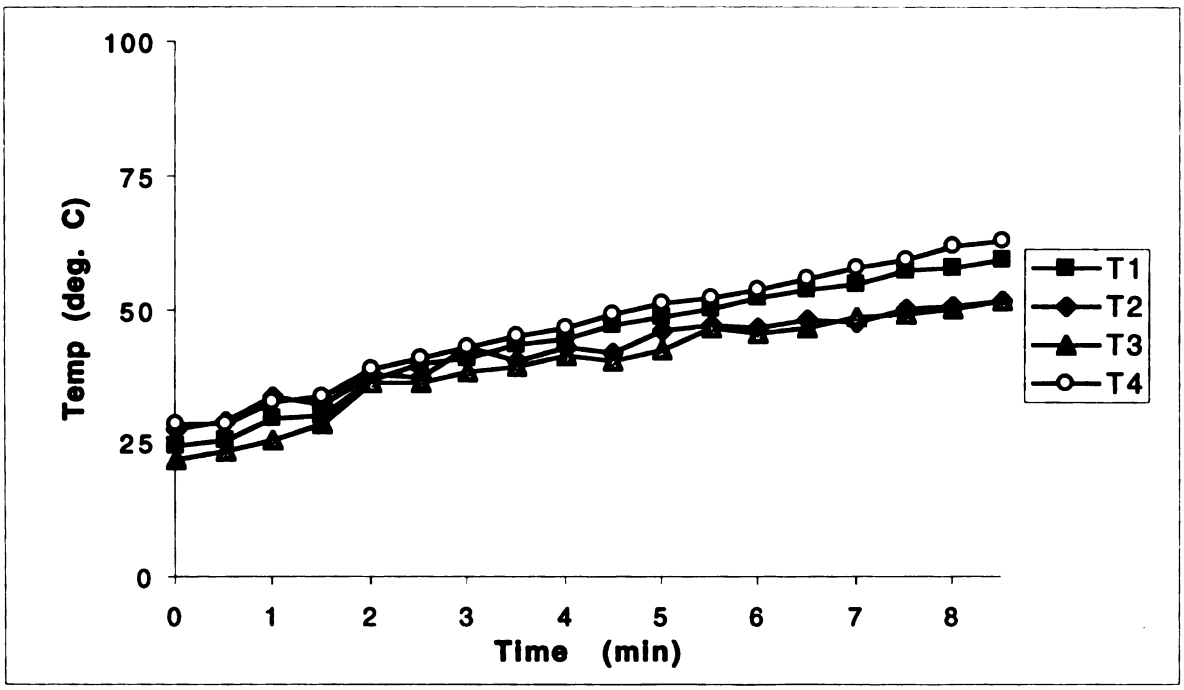

Figure 6.5 - Heating Profile: F1 (Mode 4)

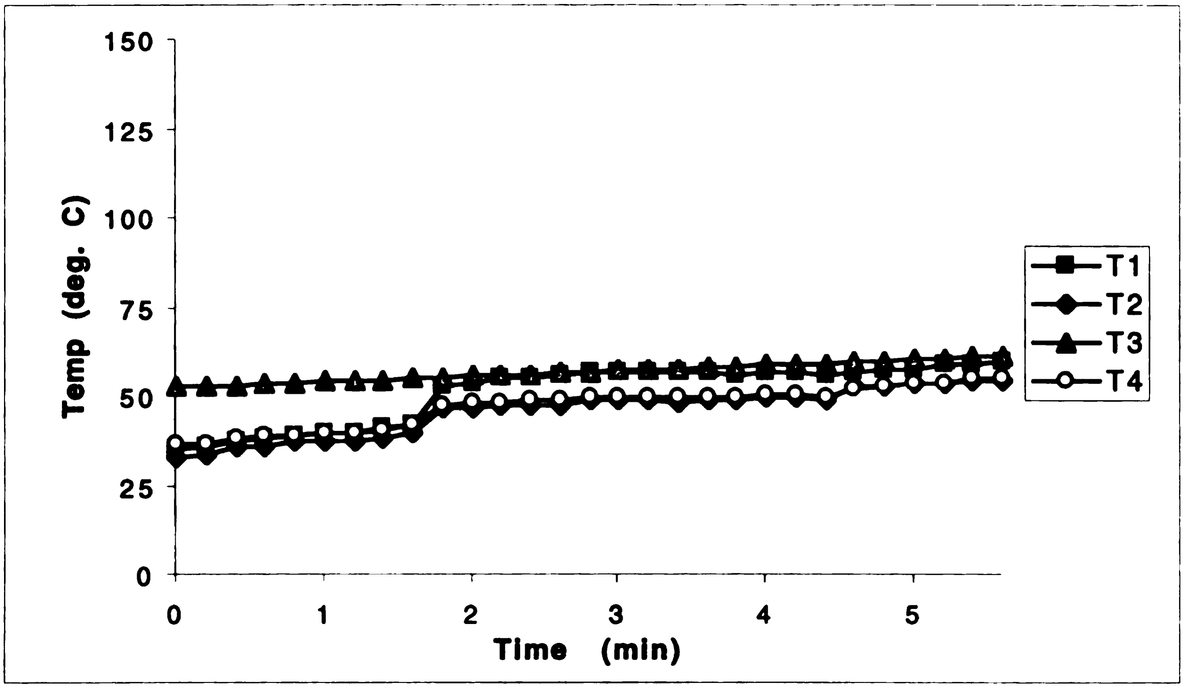

Figure 6.6 - Heating Profile: F1 (Mode 5)

6.2.4.2 Series F2

Four modes were selected for use with the 1" square samples in the F2-series experiments. The heating profiles for modes 1 through 4 are given below in figures 6.7 through 6.10.

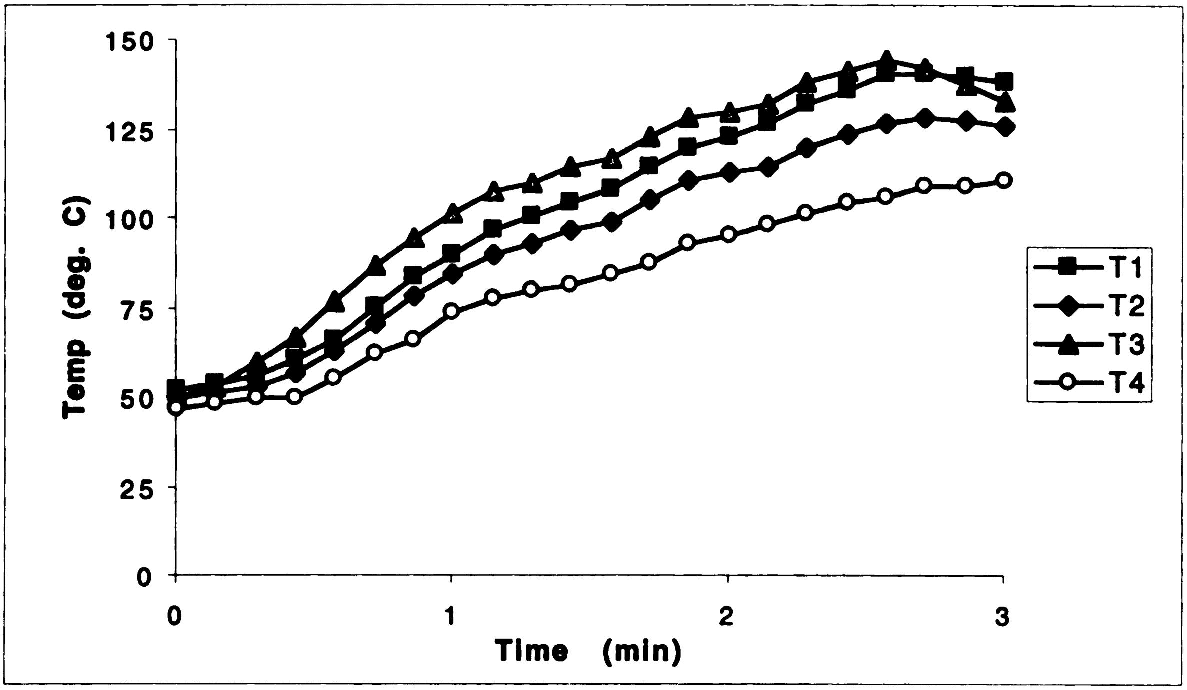

Figure 6.7 - Heating Profile: F2 (Mode 1)

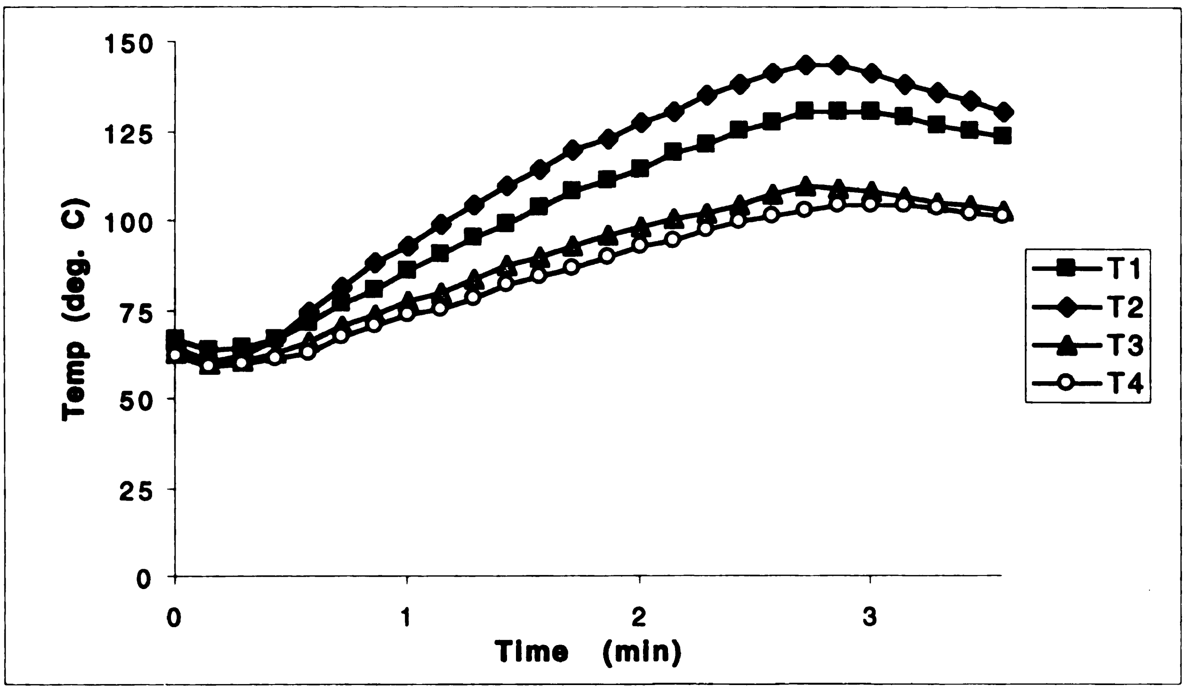

Figure 6.8 - Heating Profile: F2 (Mode 2)

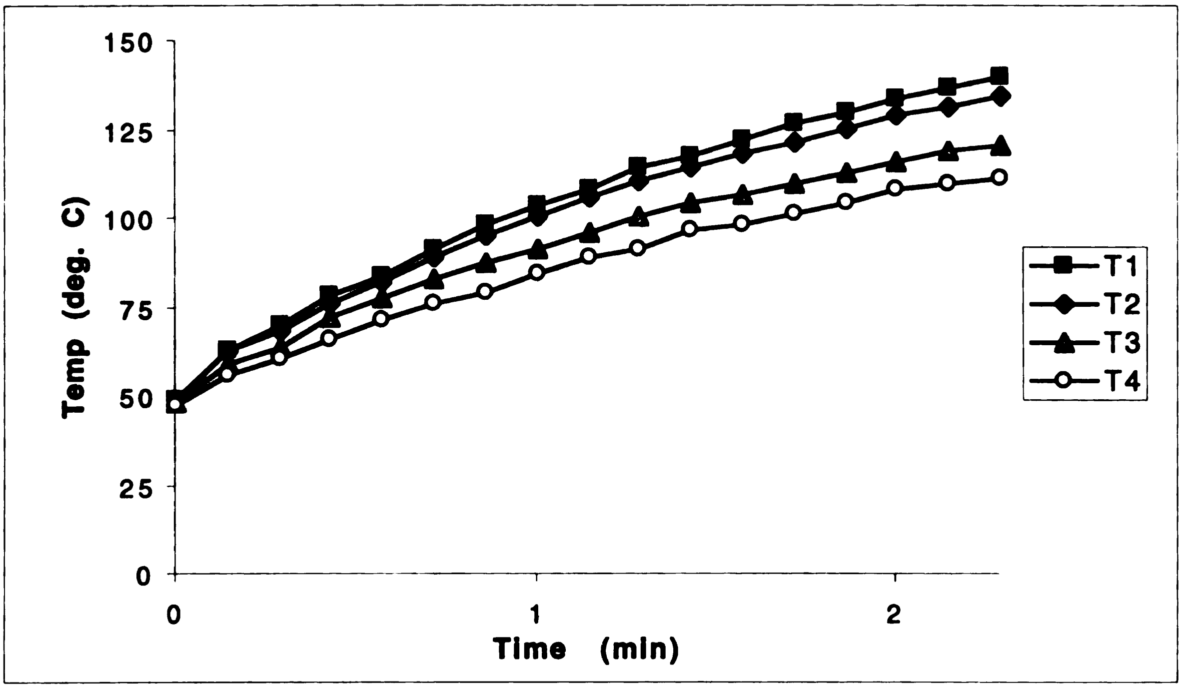

Figure 6.9 - Heating Profile: F2 (Mode 3)

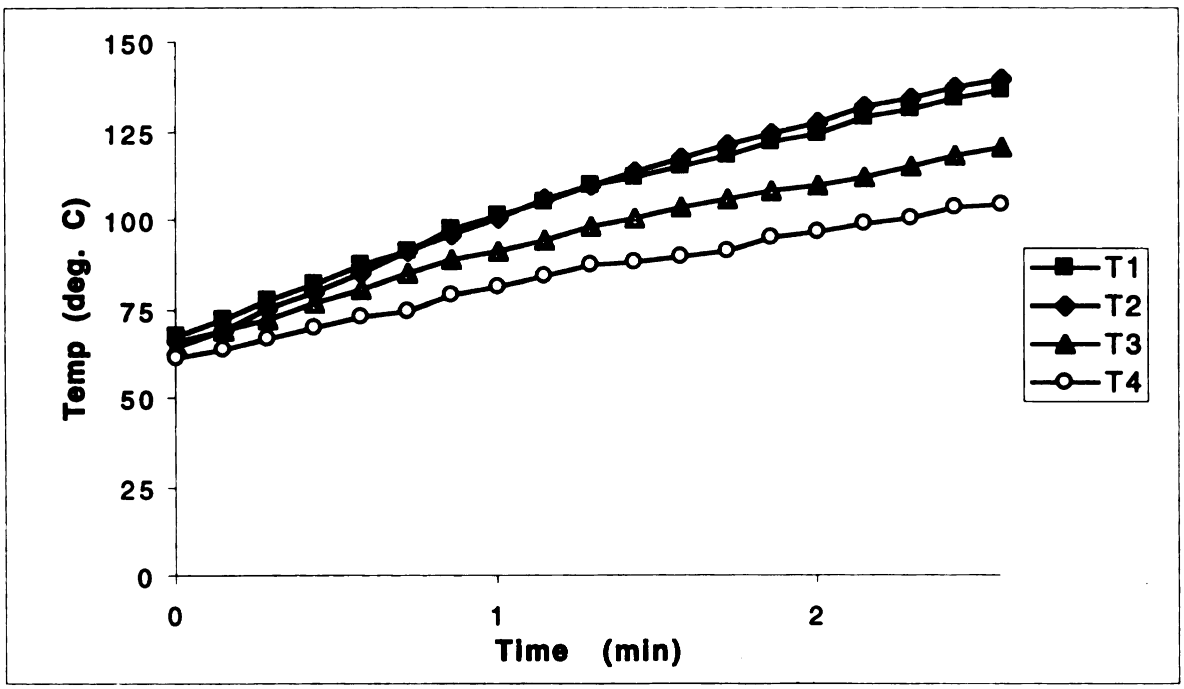

Figure 6.10 - Heating Profile: F2 (Mode 4)

6.2.4.3 Series F3

Four modes were selected for use with the 2" square samples in the F3-series experiments. The heating profiles for modes 1 through 4 are given below in figures 6.11 through 6.14.

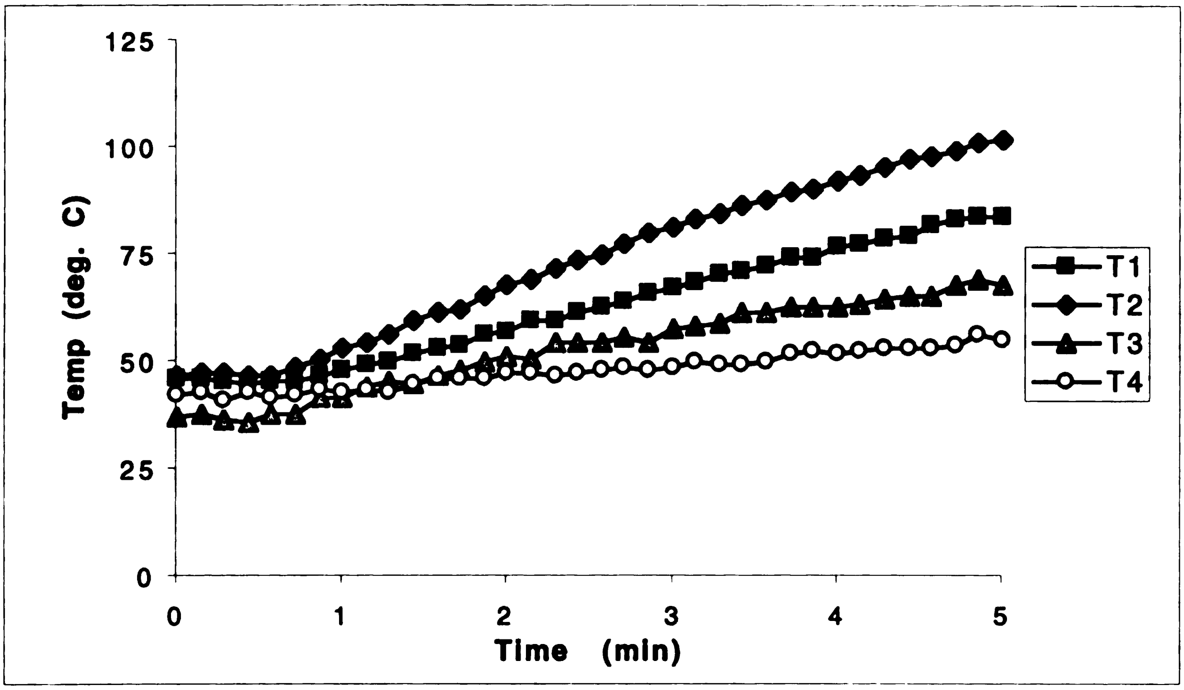

Figure 6.11 - Heating Profile: F3 (Mode 1)

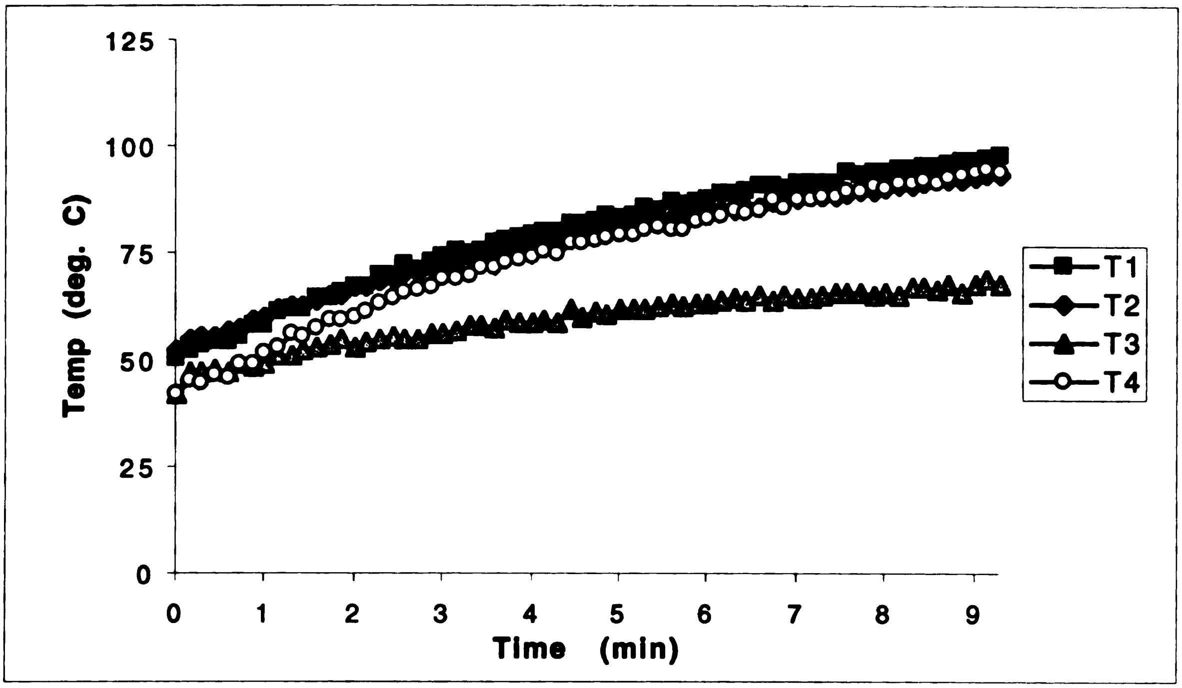

Figure 6.12 - Heating Profile: F3 (Mode 2)

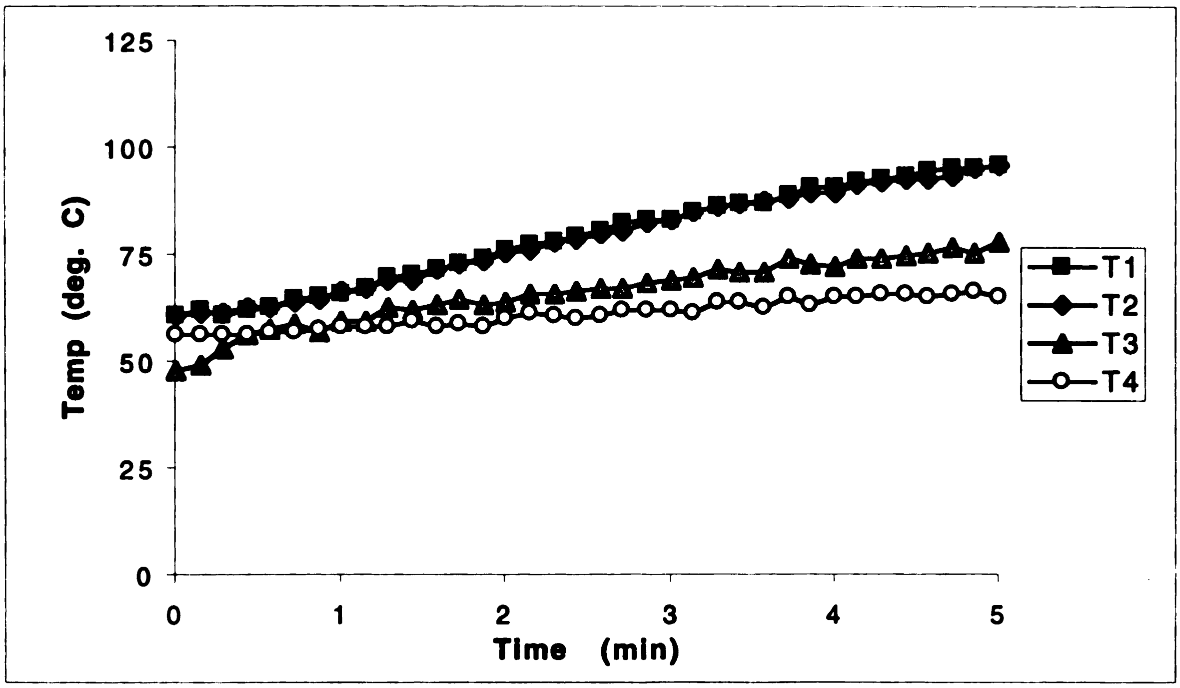

Figure 6.13 - Heating Profile: F3 (Mode 3)

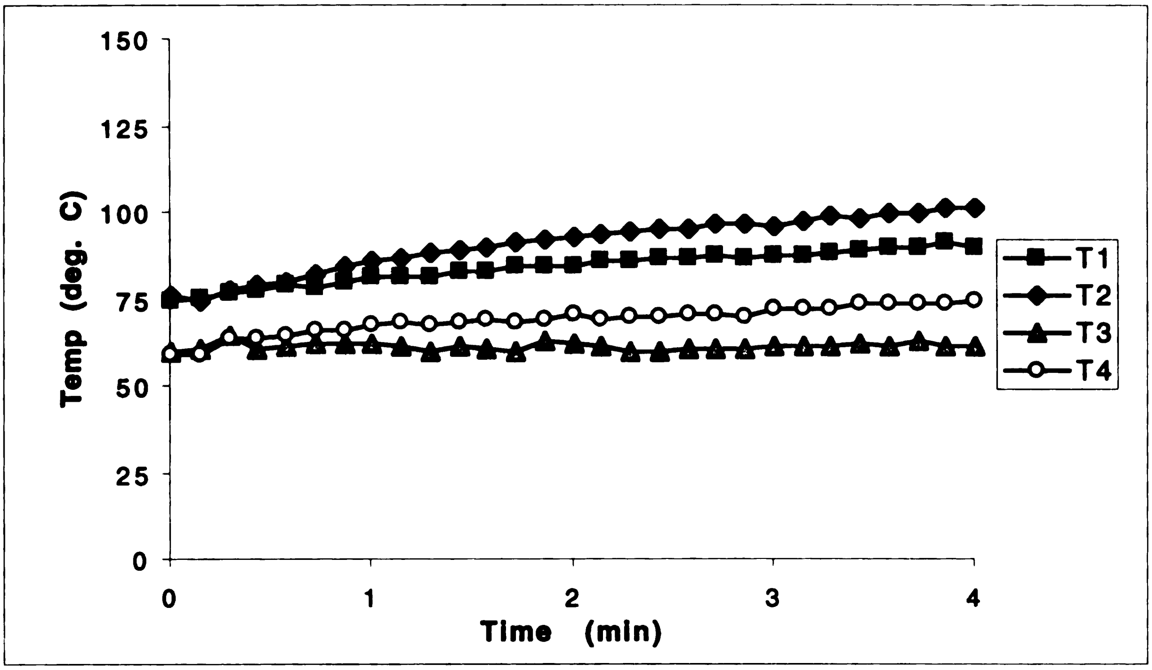

Figure 6.14 - Heating Profile: F3 (Mode 4)

6.2.4.4 Series F4

Three modes were selected for use with the 2" square samples in the F4-series experiments. The heating profiles for modes 1 through 3 are given below in figures 6.15 through 6.17.

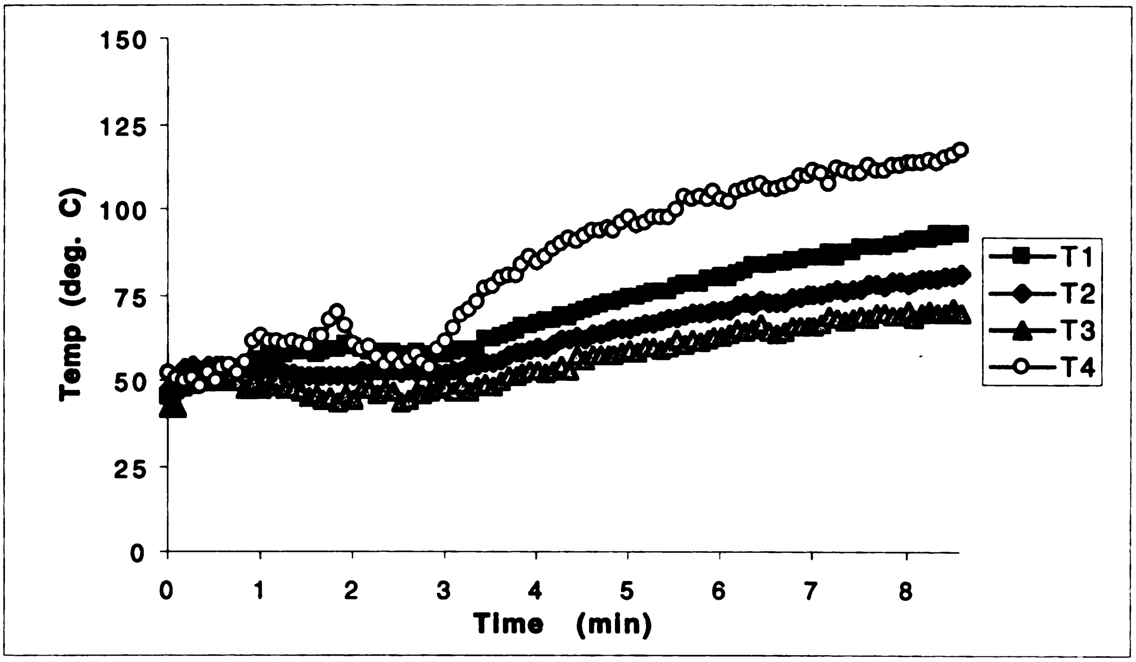

Figure 6.15 - Heating Profile: F4 (Mode 1)

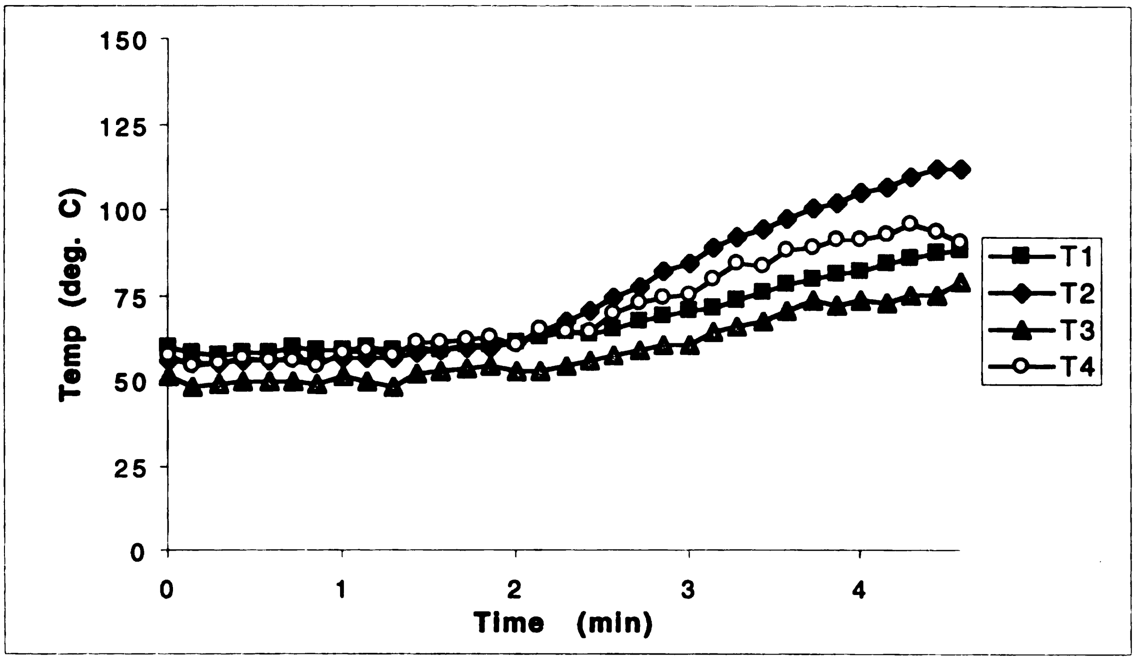

Figure 6.16- Heating Profile: F4 (Mode 2)

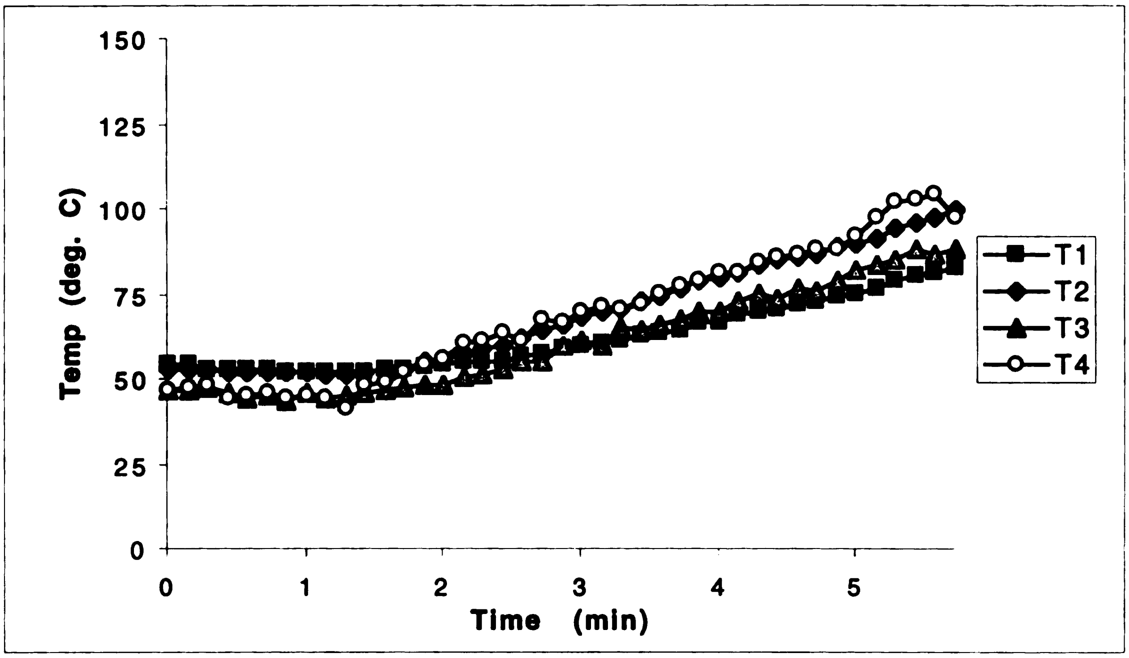

Figure 6.17 - Heating Profile: F4 (Mode 3)

6.2.5 Mode Maps

Information about the applicability of each selected processing mode is contained in a series of IPCA structures referred to as **mode maps**. These structures are specific to the geometry, placement and composition of the composite layup, hence there are distinct mode maps for each experimental series F1 through F4. The maps were manually generated for the initial experimental series F1. For subsequent experimental series F2 through F4, the maps were generated automatically by the IPCA mode map generation utility described in appendix C.

Each mode map is used with an associated qualitative part state, which determines how the map is used to determine the ideal processing mode for a sample which is in a given qualitative state. For example, if a particular sample is cold in the center, some modes would be more appropriate candidates to warm the center portion than others (it would be desirable to choose a mode which heated the center in order to reduce the thermal gradient across the sample.) The mode map provides this information to IPCA, given a suitable description of the qualitative state of the sample.

For the experimental series F1 through F3, a 16-valued qualitative part state was used. The sixteen qualitative part states are shown along the y-axis of the mode maps for the F1 through F3 series. The part states result from 2^4 possible binary states of the 4 temperature probe inputs when compared with the average sample temperature. For example, part state 2 indicates that T4 is cold while T1-T3 are hot. The sixteen possible part states are depicted in table 6.2 below.

For the experimental series F4, a 12-valued qualitative part state was used. The twelve qualitative part states are shown along the y-axis of the F4-series mode map. The part states result from 4x3 possible combinations of the hottest and coldest of 4 temperature probes. There are 4 choices for the hottest probe, then 3 remaining choices for the coldest probe, yielding 4x3=12 qualitative part states. For example, part state 1 indicates that T1 is the hottest probe and T2 is the coldest. The twelve possible part states are given in table 6.3 below.

re

ex

su

Part State	T1	T2	T3	T4
1	Hot	Hot	Hot	Hot
2	Hot	Hot	Hot	Cold
3	Hot	Hot	Cold	Hot
4	Hot	Hot	Cold	Cold
5	Hot	Cold	Hot	Hot
6	Hot	Cold	Hot	Cold
7	Hot	Cold	Cold	Hot
8	Hot	Cold	Cold	Cold
9	Cold	Hot	Hot	Hot
10	Cold	Hot	Hot	Cold
11	Cold	Hot	Cold	Hot
12	Cold	Hot	Cold	Cold
13	Cold	Cold	Hot	Hot
14	Cold	Cold	Hot	Cold
15	Cold	Cold	Cold	Hot
16	Cold	Cold	Cold	Cold

Table 6.2 - Qualitative Part States for Series F1-F3

The qualitative part state model was altered for series F4 on the basis of experimental results from the F1 through F3 series, which indicated that the two measured temperature extremes should be weighted more heavily than the two more moderate temperature measurements in determining the state of the material.

6.2.5

T

6.18

avail

ness

axis i

8) wh

8.

Part State	Hot probe	Cold probe
1	T1	T2
2	T1	T3
3	T1	T4
4	T2	T1
5	T2	T3
6	T2	T4
7	T3	T1
8	T3	T2
9	T3	T4
10	T4	T1
11	T4	T2
12	T4	T3

Table 6.3 - Qualitative Part States for Series F4

6.2.5.1 Series F1

The mode maps for experimental series F1 are shown in figure 6.18 and 6.19. Figure 6.18 depicts the initial, hand-generated surface before mode adaptation occurred. The available modes 1 through 5 are shown along the x-axis. The z-axis shows appropriateness ratings as a function of mode and qualitative part state. Higher ratings along the z-axis indicate greater applicability. For example, there is a peak in the map at the point (S4, 8) which indicates that mode 4 is highly useful when the sample is in qualitative part state 8.

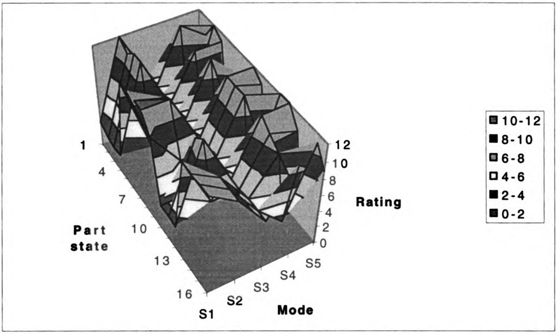

Figure 6.18 - Mode Map Surface: F1 (initial)

As the F1 experimental series proceeded, the mode map gradually changed due to the influence of the adaptive mode switching methods employed. The adaptive methods changed the applicability measures for the sample in response to actual processing experience. The resulting "tuned" map is shown in figure 6.19.

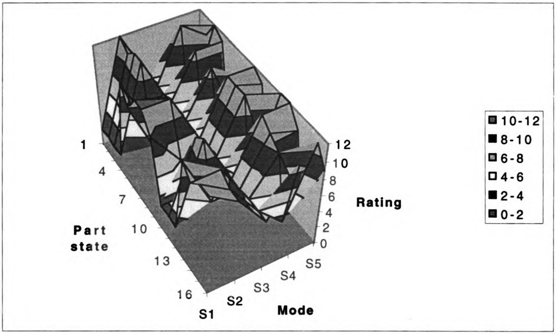

Figure 6.19 - Mode Map Surface: F1 (final)

6.2.5.2 Series F2

The mode maps for experimental series F2 are given in figures 6.20 and 6.21. The initial mode map shown in figure 6.20 was generated with the IPCA mode map generation utility from the heating profiles previously presented. Mode adaptation was again used to "fine-tune" the map across the set of F2 experiments. The final map (after adaptation) is shown in figure 6.21.

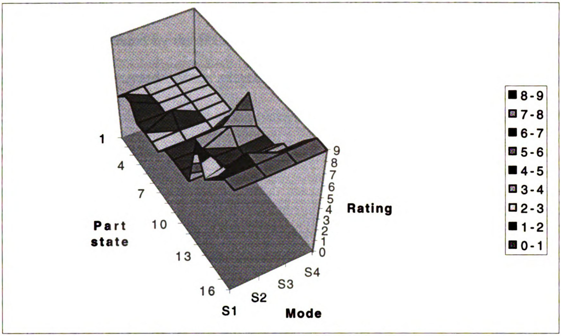

Figure 6.20 - Mode Map Surface: F2 (initial)

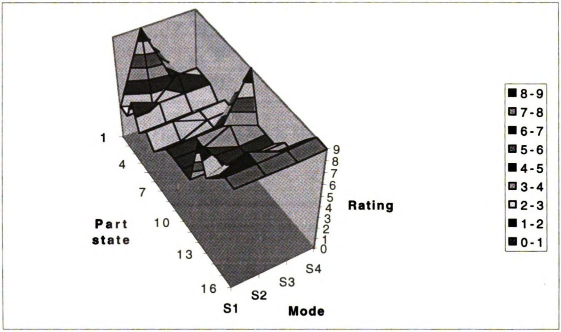

Figure 6.21 - Mode Map Surface: F2 (final)

6.2.5.3 Series F3

The mode map for experimental series F3 is given in figure 6.22. The mode map was automatically generated by the IPCA mode map generation utility based on heating profiles previously presented. Mode adaptation was not used during series F3, so the map remained static and did not adjust with processing experience. Just as with series F1 and F2, the 16-state qualitative part model was used.

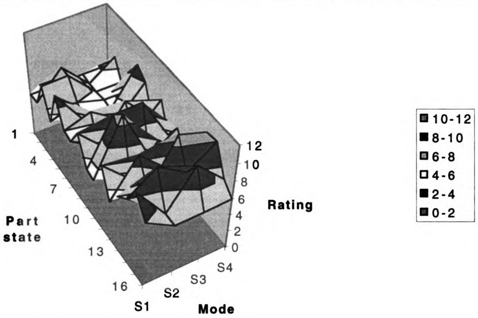

Figure 6.22 - Mode Map Surface: F3

6.2.5.4 Series F4

The mode map for experimental series F4 is given in figure 6.23. Mode adaptation was not used during experimental series F4, but the 12-value qualitative part state model was used. The F4 mode map was automatically generated by the IPCA mode map generation utility described in appendix C. The mode map is based on the heating profiles for the 2" square 12-ply DGEBPA sample described in the previous section.

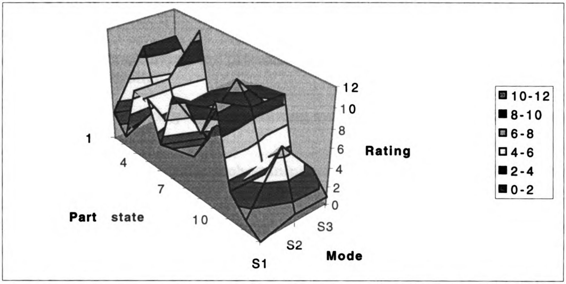

Figure 6.23 - Mode Map Surface: F4

6.2.6 Mechanical Limitations

Mechanical limitations prevent instantaneous switching between processing modes since the cavity and probe tuning mechanisms must be physically moved to new positions. This is unfortunate from the perspective of microwave-based materials processing, since lower gradients could be achieved if the switching delay could be reduced or eliminated. The limitations of the cavity and probe tuning mechanisms and their impact on the experimental series F1 through F4 are explained in this section.

The delay in switching between modes (the **mode switch** time) has several components:

- **Delay in transmitting mode-change directives to the External Controller**
- **Delay in receiving and processing mode-change directives at the External Controller**
- **Delay in moving the cavity-tuning mechanism to an approximate new location**
- **Delay in moving the probe-tuning mechanism to an approximate new location**
- **Delay in tuning the requested mode at new cavity and probe locations**

The first two components of the mode switch time are almost negligible. The delay in transmitting the mode-change directive to the External Controller depends on the end-to-end speed of the network link established between IPCA and the External Controller, but is generally far less than 1 second. Similarly, the time required for the External Controller to receive and process the mode-change directive is generally less than 1 second, unless the External Controller is busy with another task at the time.

The significant components of the mode switch time result from the need to move the cavity and probe tuning mechanisms. This portion of the overall delay is called the **seek time**. The seek time between two modes is directly proportional to the distance that the cavity and probe tuning mechanisms are required to travel. Also, mode seeking takes longer when the cavity-tuning mechanism is moved upward (against the force of gravity) than when it is moved downward because the stepper-motors have limited torque and tend to "miss" occasional steps when heavily loaded.

The other significant portion of the overall mode switch time results from the delay in tuning the precise mode location once the cavity and probe mechanisms are located within the correct general vicinity. Since the dielectric constant of the composite sample changes during curing (and the mode locations thus shift), the modes cannot be simply selected by establishing particular static cavity and probe positions. The External Controller must locate (lock) the mode within a range of cavity and probe positions. The **lock time** can vary, depending on the precision of the mode lock specified. The precision is specified within the External Controller as a fraction of incident power which must be achieved before the tuning attempt is aborted. Specifying a more precise mode lock will result in a longer average lock time.

The average mode switch times for the F1 through F3 series were measured and the results are shown in figures 6.24 through 6.26.

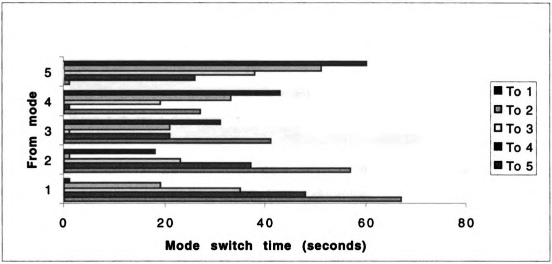

Figure 6.24 - Average Mode Switch Times: F1

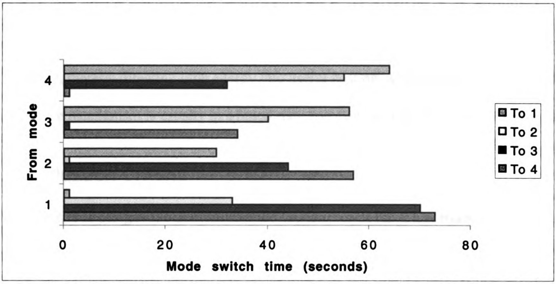

Figure 6.25 - Average Mode Switch Times: F2

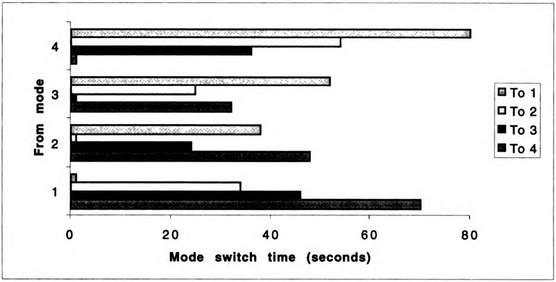

Figure 6.26 - Average Mode Switch Times: F3

In all F-series experiments, an artificial **hold time** was used to obtain a trade-off between heating rate and thermal gradient. The hold time is defined as the minimum time between successfully tuning a particular mode and permitting a request for a different mode. If this time is small, the mechanism switches between modes often, thus maintaining a small gradient. However, under these conditions the cavity is not frequently in a tuned state, thereby directing little energy into useful heating. Conversely, if the hold time is large, the mechanism is usually in a tuned state and directs more energy into the sample.

6.3 Evaluation Criteria

The C1 series is designed to demonstrate the capabilities of the IPCA Blueprint Generator and Plan Assembler implementations through the use of selected fabrication models. The criteria used to evaluate this series include completeness and correctness of the generated process control plans, including the appropriateness of the selected plan fragments.

Series C2 demonstrates the breadth of domain coverage in the Blueprint Generator and fragment library. Blueprints generated under the C2 series are examined for fragment and goal content, demonstrating that a range of materials and fabrication scenarios are covered. Evaluation includes demonstrating variant fragment content in the generated blueprints and showing that the referenced fragments are contained within the fragment library.

Series C3 is used to demonstrate process control plan behavior under various simulated conditions. Evaluation criteria include appropriateness of response to simulated conditions and reaction time.

Series F1 is used to show the appropriateness of the prototype implementation by effectively controlling a real-world process. This series demonstrates that control behavior improves as fragments are suitably modified, either through adaptive methods or through direct modification. Evaluation criteria for series F1 include subjective quality measures on the fabricated structures and minimization of recorded thermal gradients for the 3" square layup. Subjective quality measures are based on a visual analysis of the completed structure. Any charring, incomplete curing, delamination, voids, warping or other surface faults are noted in the analysis and used to evaluate effectiveness of the control approach. The criteria for indirectly evaluating the quality of the cure involves an examination of the thermal profile. Heating rates, extent of thermal gradient and power delivered to the structure are used as evaluation criteria. Improvement is measured by confirming a reduction in thermal gradients during fabrication and a reduction in the defects on the completed structures.

Series F2 demonstrates that a second structure can be fabricated and is used to demonstrate the trade-off between heating rapidly and maintaining a minimal thermal gradient. Due to the low mass of the 1" square part used in series F2, much greater heating rates can be obtained than with the 3" square part used in series F1. The limitations of the mechanical tuning mechanisms are explored and a case is presented for scaling to higher power microwave sources and larger samples. The benefits of using a variable-frequency power supply in place of manual mode tuning are presented.

Series F3 demonstrates coverage for a third structural geometry: the 12-ply 2" square. The experiments in this series are designed to demonstrate rapid heating with slightly larger parts. Heating rates are not as great as in series F2, but controllability is greater due to the larger size of the sample in comparison with the cavity diameter.

Series F4 shows how complimentary mode selection and mode-switching policy influence the quality of processing. The policy used for these runs involves selecting modes based only on the two most extreme (hottest and coldest) sample temperatures. Unfortunately, the mode sets used in the F-series experiments were not fully complimentary with regard to their heating patterns. Ideally, one would prefer to utilize modes which compliment each other to the largest possible degree in order to minimize thermal gradients. The F4 series shows that a high-quality mode set can considerably reduce thermal gradients during processing.

7.

Th

tal

7.

7.

Th

su

un

he

ate

sur

CHAPTER VII

7.1 Experimental Results and Discussion

The experimental results are summarized in this section. Results are presented using tables and graphs. A brief discussion of results is presented with each experimental series.

7.2 C-Series

7.2.1 Series C1

The C1 series demonstrates that the IPCA blueprint generator and fragment library has sufficient breadth to generate process control plans for fabricating composite structures under various conditions. These conditions include the generation of plans for gathering heating profiles for candidate modes, fabricating composite materials without pre-generated mode maps, fixed-time curing and variations for high and low mass samples. An summary of experimental series C1 is given in table 7.1.

Experiment	Conditions	Process control plan / fragment content	Fragment notes
C1.1	Normal plan, DGEBPA, Carbon fiber, 3" square, 4 temp. probes, Balance gradi- ent and heating rate	MDInitializationPolicyA MDModeControlPolicyB MDEmergencyPolicyA MDPowerControlPolicyB MDPhaseChangePolicyB MinProcessingTemp MaxProcesssingTemp RecommendedProcessingPower HoldTime (before/after warmup) ...	Standard Reactive Standard Multi-level Event-based 120 deg. C 140 deg. C 45 watts 60/30 sec.
C1.2	Heating profiles DGEBPA 3" square 5 modes	MDModeControlPolicyD ProcessingTime ModeUsed RecommendedProcessingPower	Heating 50 min. 5 40 watts
C1.3	No mode maps	MDModeControlPolicyC HoldTime	Reactive- Adaptive 180 sec.
C1.4	Fixed-time	MDPhaseChangePolicyA	Time-based
C1.5	Low mass	RecommendedProcessingPower	30 watts
C1.6	High mass	RecommendedProcessingPower	60 watts

Table 7.1 - Summary of Experimental Series C1

Experiment C1.1 and C1.4 confirmed that IPCA can generate complete process control plans. The plans were exercised against the simulated control unit to verify appropriate fragments responses and timing. The C1.2 experiment demonstrated that IPCA can generate process control plans designed to automatically gather mode (heating profile) data. This scenario was used to collect the mode map information for experimental series F2 and F3. The C1.2 plan sequentially places the cavity into each mode from the candidate set, applies a preset power level (based on the mass of the sample), allows the material to either heat to a minimum temperature (120 deg. C for DGEBPA-based samples) or for ten minutes, and turns off the power to allow sample cooling. The procedure is repeated until all modes have been tested.

Experiment C1.3 involves producing a process control plan for a sample with no available mode maps. A default set of maps is created, with all elements at an intermediate and equal rating. A reactive-adaptive mode switch method is included which randomly chooses modes (initially) and modifies the maps as processing proceeds.

Experiments C1.5 and C1.6 demonstrate that the blueprint generator can appropriately scale recommended processing power with sample mass. The high-mass experiment C1.6 involved doubling the normal 30g. mass of a hypothetical 12-ply 3" DGEBPA sample to 60g by doubling the number of plies to 24. It was expected (and verified) that the Blueprint Generator would modify the recommended processing power. Recommended processing power (used after initial material heating is complete) was scaled upward from 45 to 60 watts during experiment C1.6. Similarly, the low-mass experiment C1.5 halved the specified part mass to 15g. (6-ply) and the recommended processing power was scaled downward from 45 to 30 watts.

7.2.2 Series C2

The C2 series demonstrates that the IPCA Blueprint Generator and fragment library have sufficient coverage to produce process control plans which are suitable for processing a variety of sample sizes and compositions. The experiments C2.1 through C2.3 were conducted before processing began in the laboratory for the F1 through F3 series. These experiments show that the Blueprint Generator adjusts power levels in response to sample size and mass. The recommended processing power for C1.2 was 45 watts, for C1.3 it was 30 watts and for C2.3 it was 20 watts.

Experiment	Conditions	Fragment content	Fragment notes
C2.1	DGEBPA 3" square	MinProcessingTemp TargetProcessingTemp MaxProcessingTemp ThermalDegradationTemp RecommendedProcessingPower	120 130 140 160 45 watts
C2.2	DGEBPA 2" square	RecommendedProcesingPower	40 watts
C2.3	DGEBPA 1" square	RecommendedProcessingPower	35 watts
C2.4	PolyesterGlass 3" square	MinProcessingTemp TargetProcessingTemp MaxProcessingTemp ThermalDegradationTemp RecommendedProcessingPower	180 190 200 250 100 watts
C2.5	PolyesterGlass 2" square	RecommendedProcessingPower	80 watts
C2.6	PolyesterGlass 1" square	RecommendedProcessingPower	60 watts

Table 7.2 - Summary of Experimental Series C2

Experiments C2.4 through C2.6 show that the Blueprint Generator can set appropriate processing and thermal degradation temperatures for various matrix materials. In C2.1 through C2.3, the target processing temperature for DGEBPA was set to 130 deg. C and the thermal degradation temperature was set to 160 deg. C. The processing temperatures for experiments C2.4 through C2.6 were set at 180 deg. C, while the power levels were set to 100, 80 and 60 watts. The thermal degradation temperature was set to 250 deg. C.

7.2.3 Series C3

The C3 series demonstrates that IPCA-generated process control plans can effectively respond to a variety of unexpected inputs due to sensor failure, loss of control, mechanical failure and thermal runaway.

Experiment	Simulated condition	IPCA reaction
C3.1	Overtemp: TDeg temp. exceeded	Shut off power next execution cycle; continue executing process control plan; return power to 80% after cooling below desired processing temperature
C3.2	Stuck Lc	Stop execution of process control plan next cycle; shut-down signalling actuator or sensor failure
C3.3	Stuck Lp	Stop execution of process control plan next cycle; shut-down signalling actuator or sensor failure
C3.4	Low Pi	Stop execution of process control plan next cycle; shut-down assuming sensor or power supply failure
C3.5	High Pi	Stop execution of process control plan next cycle; shut-down assuming sensor or equipment failure
C3.6	High Pr	Stop execution of process control plan next cycle; shut-down assuming sensor or equipment failure

Table 7.3 - Summary of Experimental Series C3

During experiment C3.1, the simulated sample temperature was suddenly increased beyond the thermal degradation limit. The emergency sponsor activated on the next execution cycle and shutdown the power completely.

Experiments C3.2 and C3.3 demonstrated that IPCA-generated plans would not continue to apply control signals which could burn out control electronics or motors. The Lc and Lp positioning potentiometer signals were frozen during a simulation. The control signals called for a change of mode, but the mechanisms did not "move." The emergency sponsor was activated and the processing run was shut down.

For experiment C3.4 and C3.5, the Pi sensor reading was artificially maintained at 0 or 100, thus indicating sensor failure or power supply failure. In C3.5, the reflected power sensor reading was maintained at 80% of incident power, indicating a dangerous safety condition in which excessive power was being reflected back from the cavity. In both experiments, the emergency sponsor was activated on the next control cycle, the equipment was shut down and the execution of the plan was terminated.

7.3 F-Series

The F-series experiments were conducted in the laboratory, using the single-mode microwave cavity already described. The results of selected experiments from the F-series are presented in this section.

7.3.1 Series F1

Series F1 demonstrated that IPCA can produce process control plans which effectively control the fabrication of 3" 12-ply unidirectional DGEBPA samples. The experiments which formed the F1 series were the first real-world trials for the IPCA implementation. The average thermal gradients were gradually reduced over time, due to the use of the adaptive mode-switching methods. The standard deviation of the thermal gradient was also reduced with processing experience. All the samples from this series were cured for 90 minutes (with the exception of F1.1). Processing run F1.8 was performed on a 24-ply sample. Sensor and actuator values were not recorded during processing run F1.8, thus statistics are not available for that particular experiment.

Experiment	Time (min)	Hold time (sec)	T_{max} (C)	$Grad_{max}$ (C)	$Grad_{avg}$ (C)	$Grad_{stddev}$
F1.1	32	60	60.61	22.02	6.77	3.85
F1.2	90	60	155.96	76.27	51.78	16.01
F1.3	90	60	150.01	58.64	42.90	11.37
F1.4	90	60	135.53	58.44	28.78	8.42
F1.5	90	60	157.61	31.78	20.27	5.54
F1.6	90	90	152.22	69.92	51.17	6.59
F1.7	90	90	130.48	60.93	39.72	4.99
F1.8	90	90	n/a	n/a	n/a	n/a

Table 7.4 - Summary of Experimental Series F1

The early (figure 7.1) and late (figure 7.2) thermal profiles clearly show the reduced thermal gradients during processing, mainly attributed to mode map adaptation. The average thermal gradient was reduced from over 51 degrees C in experiment F1.2 to under 21 degrees C in experiment F1.5. The gradients for experiments F1.6 and F1.7 were somewhat greater due to an increase in the hold time from 60 to 90 seconds. The longer hold time encouraged faster heating in the latter experiments, but increased the average and maximum gradients, and this is clearly reflected in the data of table 7.4.

Each F-series thermal profile depicts the temperatures at the four probe locations T1-T4 in degrees Celsius (C), the incident power level Pi (measured in watts), and the current processing mode. The mode number has been scaled by a factor of 10 to assist with overall clarity and readability, so mode 1 is graphed with a value of 10, mode 2 with a value of 20, and mode 5 shows up as 50.

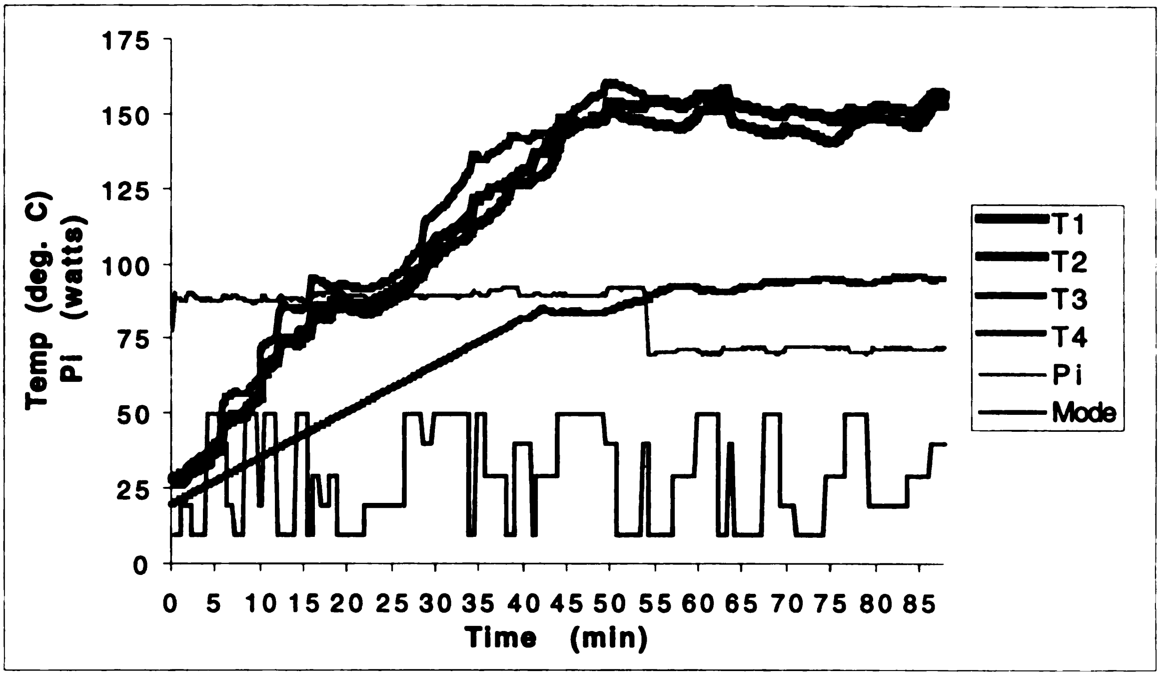

Figure 7.1 - Thermal Profile (F1.2)

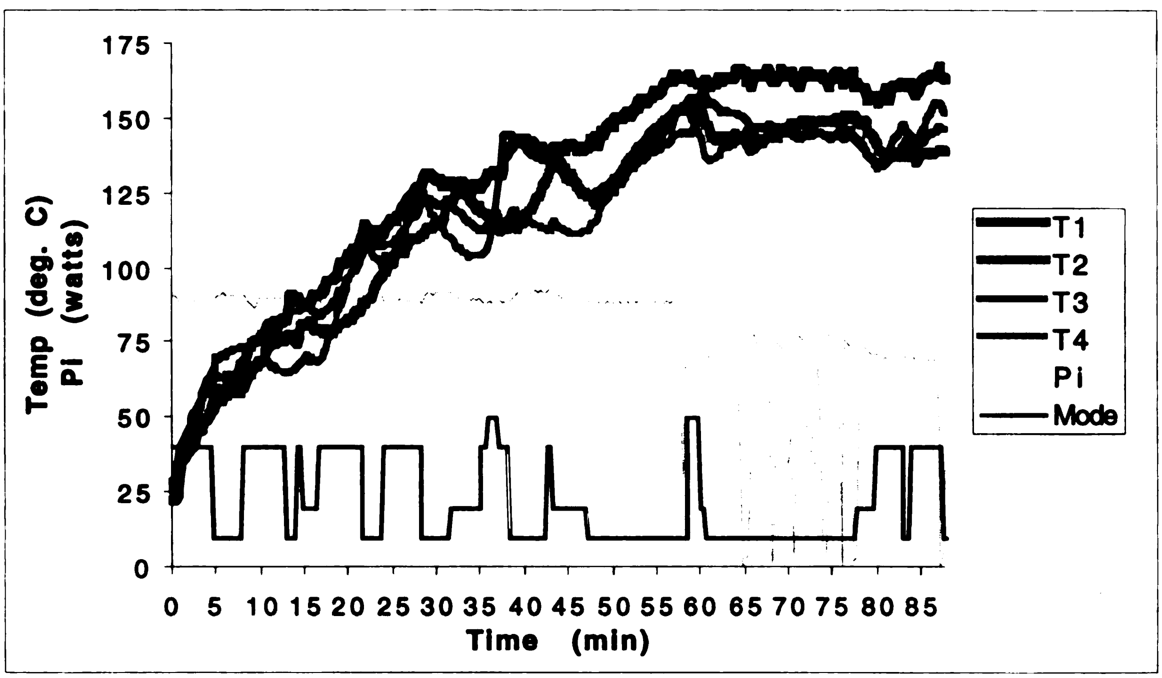

Figure 7.2 - Thermal Profile (F1.5)

Since reduced processing gradients result in higher quality composite structures with superior mechanical properties, it is reasonable to state that the quality of the resultant parts should have increased as the thermal gradients were minimized during subsequent processing runs. Although material testing was not performed on the samples resulting from the IPCA research, it can be said that lower thermal gradients yield higher-quality composite structures. The question of how much the mechanical properties of the resulting materials improved is a question beyond the scope of this research.

Therefore, during the analysis of subsequent processing runs and in comparison with previous results [Fellows 91], the thermal gradient will be taken as a quantitative indicator of resultant part quality.

Figure 7.3 depicts the parts resulting from experiments F1.2 through F1.8, left to right, top to bottom. The part from F1.2 is charred on the sides. This effect was reduced in F1.3 and F1.4, as the average gradients fell from 51.78 deg. C to 42.90 and 28.78 deg. C, respectively. The least charring occurred in F1.5, with an average thermal gradient of just 20.27 deg. C. Note the uniform, dark surface and lack of discoloration in the F1.5 image.

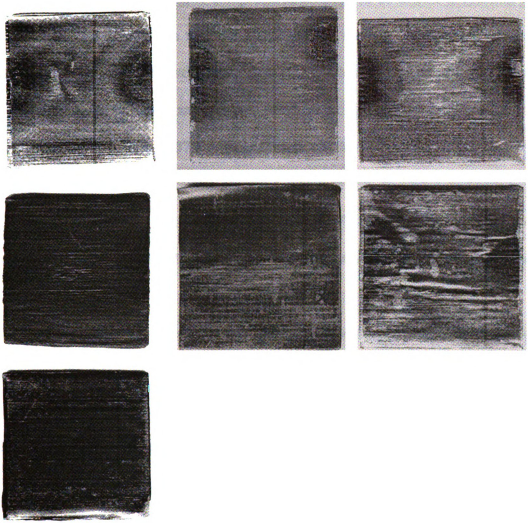

Figure 7.3 - Surface Scans (F1.2- F1.8)

The samples from F1.6 and F1.7 are not charred, but they show some evidence of surface irregularities, perhaps as a result of the more rapid material heating and slower mode-switching response. Sample F1.8 also showed a uniform appearance and lack of charring, demonstrating that variances in material thickness do not necessarily interfere with IPCA's ability to obtain a high-quality cure.

7.3.2 Series F2

Series F2 demonstrated that smaller composite structures can be heated to reaction temperature rapidly (in comparison with the 3 inch squares used in experimental series F1) and that there is a trade-off between heating quickly and heating evenly. One-inch square 12-ply unidirectional DGEBPA samples were used for all F2 series experiments. Average gradients ranging from 4.76 to 37.3 degrees C were obtained. Heating rates from 0.90 to 45.97 degrees C per minute were also observed, which greatly exceeded the heating rates obtained in the F1 series. Experiments F2.8 through F2.10 used a proportional power control method while all other F2-series experiments used a fixed power as listed in table 7.5.

A clear relationship between hold time, thermal gradient and heating rate may be constructed from the data in table 7.5, and the precise details of this trade-off will be explored in section 7.4 concerning the quantitative analysis of experimental results.

The central qualitative conclusion drawn from the F2 series experiments is that longer hold times create faster sample heating but larger thermal gradients. This is because shorter hold times cause the cavity to spend more time in an detuned state (since it is switching between modes more frequently) and that this mechanical thrashing behavior causes the sample to absorb less (focused) microwave energy. If the hold time is very short, IPCA is permitted to request mode changes very frequently, thus preventing the cavity from remaining in a tuned state for a long period of time.

Experiment	Time (min)	Hold time (sec)	Power (watts)	T_{max} (C)	Heating rate (C/min)	$Grad_{avg}$ (C)
F2.1	8	60	20	126.95	12.68	37.30
F2.2	6.5	30	30	120.00	20.46	17.17
F2.3	8	15	30	67.81	0.90	4.76
F2.4	12	20	40	107.97	4.28	9.70
F2.5	1.5	30	50	120.29	45.97	22.16
F2.6	2.5	30	40	129.68	16.95	32.28
F2.7	8.5	20	50	131.03	11.61	33.06
F2.8	2	30	Prop	109.24	35.16	21.71
F2.9	7.5	30	Prop	82.24	1.63	12.09
F2.10	8.5	45	Prop	117.20	8.03	20.71

Table 7.5 - Summary of Experimental Series F2

A demonstration of rapid heating during experiment F2.1 can be seen in figure 7.4. The heating rate from experiment F2.1 was very good (compared to the F1 series results) at a rate of 12.68 deg. C per minute, but the temperatures diverged to a 38.2 deg. C maximum gradient. The hold time was 60 seconds, allowing time for only 4 mode changes during the processing run. The inability to switch between modes more rapidly had a positive effect on the heating rate, but had a negative effect on the thermal gradient.

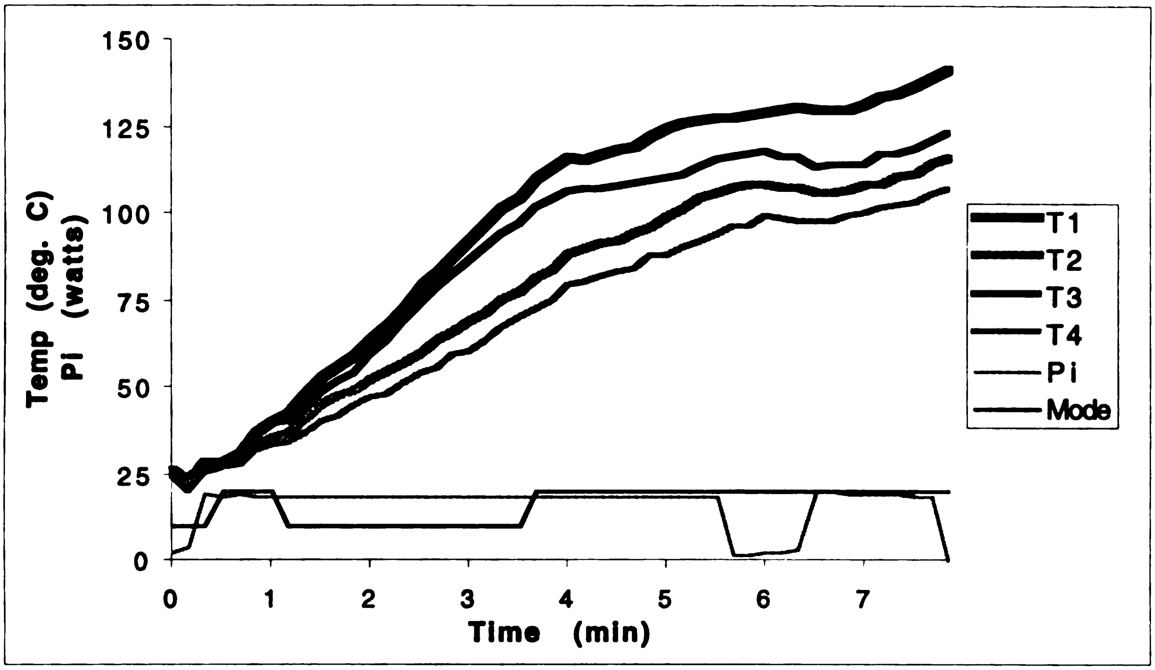

Figure 7.4 - Thermal Profile (F2.1)

Conversely, during experiment F2.3 (shown in figure 7.5) the average gradient was only 4.75 deg. C, but the average heating rate was less than 1 degree C per minute. The hold time for F2.3 was 15 seconds. This experiment showed that switching between modes very quickly leads to a very low average heating rate, especially when a low power level (30 watts) was used.

Several other combinations of hold time and power level were explored in order to provide data with which to characterize this time-gradient trade-off.

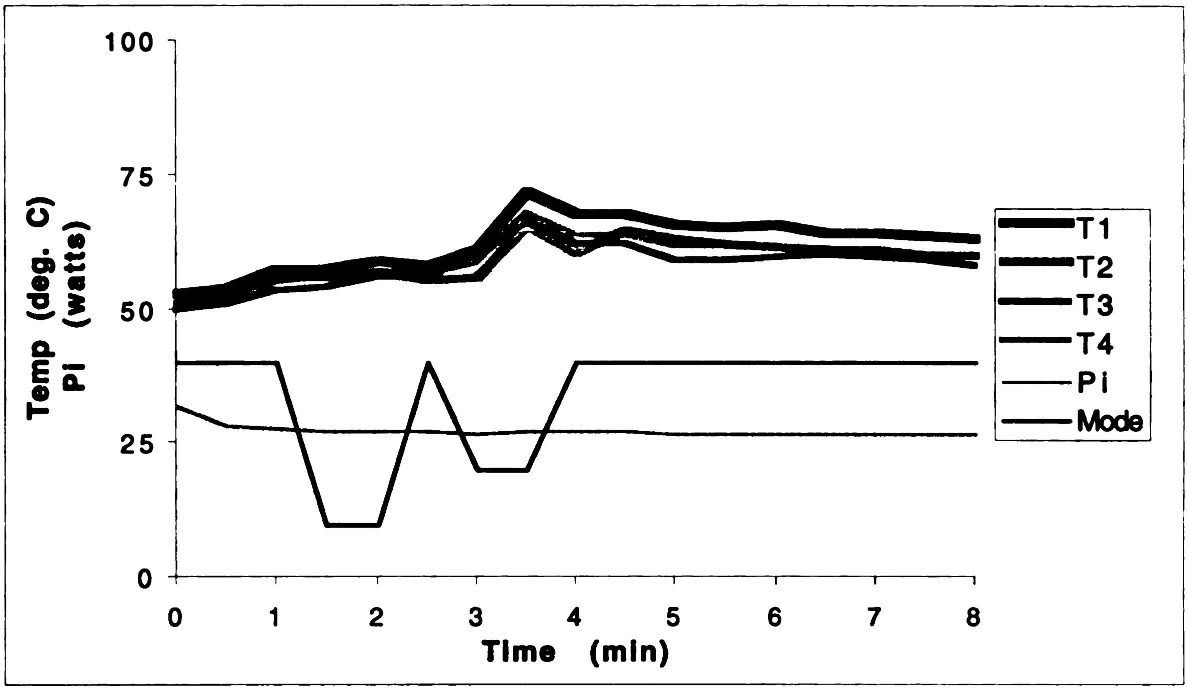

Figure 7.5 - Thermal Profile (F2.3)

A reasonable compromise between heating rate and gradient was realized in experiment F2.4 (figure 7.6). The average gradient was only 9.7 deg C., yet the sample heated effectively at a rate of 4.28 degrees C per minute. The hold time was 20 seconds and the power used was 40 watts. Twelve mode switch operations occurred during the 12 minute fabrication run, which produced both rapid material heating and low thermal gradient.

The F2 series demonstrated the trade-off between thermal gradient and hold time, and it also demonstrated the feasibility of processing a second material geometry. Thus, the F2 series showed that IPCA could successfully process composite structures with different geometries.

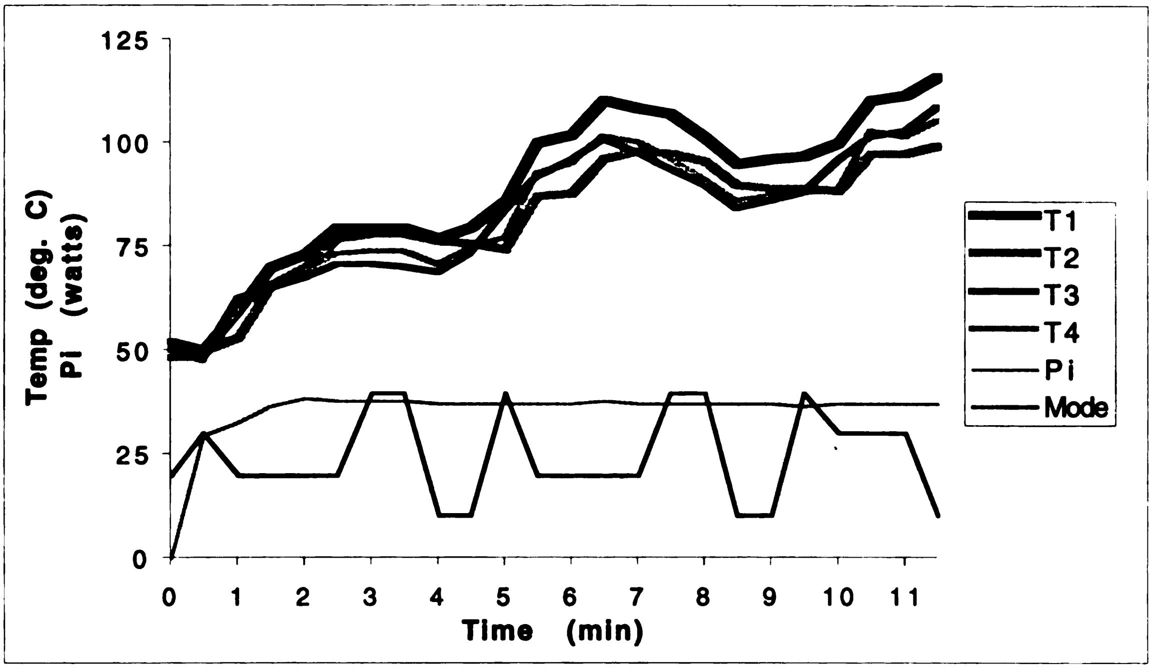

Figure 7.6 - Thermal Profile (F2.4)

The samples resulting from fabrication experiments F2.1 through F2.8 are shown in figure 7.7. The same sample was used for the three proportional power control experiments F2.8 through F2.10, and that sample is shown in figure 7.7 as F2.8. This time-saving practice was made possible by the fact that short processing times were employed and the experiments were carried out in succession, after allowing for the material to cool.

In contrast to the F1 series experiments, none of the samples from series F2 were charred during processing. This result demonstrates a marked improvement in the effectiveness of the power control fragment over the behavior observed during series F1. However, since the F2 series experiments were not intended to process the samples for extended periods of time, the curing was incomplete. Thus, the lack of charring is also partially explained by a shorter processing time.

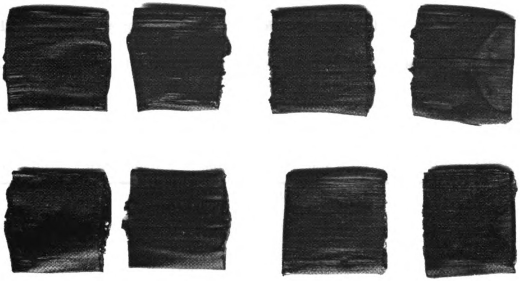

Figure 7.7 - Surface Scans (F2.1 - F2.8)

7.3.3 Series F3

The F3 series demonstrates that a compromise can be achieved between the rapid but somewhat uneven heating obtained in series F2 and the slow but uniform heating obtained in series F1. The series also demonstrated that a third structural geometry could be processed.

A series of 2" DGEBPA square 12-ply samples was used for series F3. The selection of this planar geometry was driven by the availability of a mold and by heat transfer calculations which showed that a 2" square mass would heat effectively in the cavity.

The experimental summary for series F3 is shown in table 7.6. A hold time of 60 seconds was used for experiments F3.1 and F3.2. The hold time selected for experiment F3.3

was 30 seconds. The average heating rate for the F3 series was 5.70 degrees C per minute, while the average thermal gradient was 25.25 deg. C. As expected, the heating rates and gradients were intermediate between those observed during the F1 and F2-series fabrication experiments. A third structural geometry was successfully processed.

Experiment	Time (min)	T_{max} (C)	$Grad_{max}$ (C)	$Grad_{avg}$ (C)	$Grad_{stddev}$
F3.1	19	58.42	45.55	22.72	13.86
F3.2	9	113.09	56.05	30.08	14.02
F3.3	11.5	124.23	37.29	22.96	7.73

Table 7.6 - Summary of Experimental Series F3

The thermal profile for experiment F3.2 is shown in figure 7.8, and was chosen to illustrate a problem which prompted a change in the approaches to mode map generation and mode selection used in series F4. Upon consultation with domain experts, it was determined that the thermal behavior obtained in experimental series F1 through F3 demonstrated little or no "crossing" behavior. More exactly, the thermal profiles tended to confirm that one or more temperature probe locations consistently remained above the average sample temperature, while others were always below average. This is undesirable, even if the thermal gradient is low, because a higher-quality cure would result from an equal temperature-time integral for each temperature probe location.

Referring to experiment F3.2, for example (see figure 7.8), it can be easily observed that T2 and T4 never reach a temperature at which reaction initiation can begin. For the DGEBPA material used in this experiment, that temperature, called the "onset of reaction temperature," would be well-above 100 deg. C. Additionally, the temperature-time product would ideally be equal throughout the sample for the duration of the processing run, helping to ensure optimal mechanical properties in the resulting material. This practice would also help ensure that samples were cured more uniformly.

These observations were used as a basis for modification to the mode-mapping and mode-switching algorithms. The modified algorithms were subsequently used for experimental series F4, and are described in the next section.

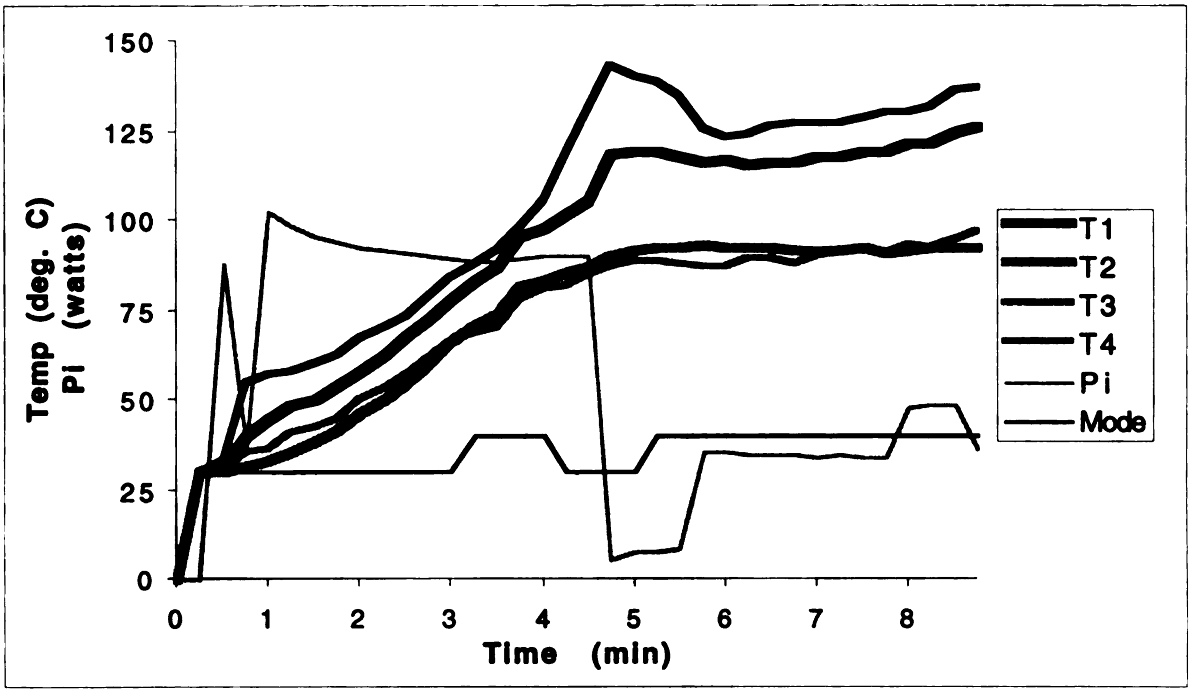

Figure 7.8 - Thermal Profile (F3.2)

The experimental series F1 through F3 utilized a mode-switching fragment designed to strictly minimize the thermal gradient across the sample. Thus, the mode-switching fragment was designed to select a new mode whenever the gradient exceeded a preset threshold for the material and the difference between the temperature readings could be reduced (subject, of course, to the hold time constraint).

The mode-switching fragment used for series F3 was modified for the F4 series experiments. The F4 mode switching fragment was designed to consider only the highest and lowest temperature probe readings in evaluating the qualitative part state. The modified mode switching fragment then selects the processing mode which will minimize the thermal gradient and also ensure maximum crossing behavior between the probe measurements. This approach was to encourage equal temperature-time integrals throughout material processing, but it also dictated a change in the manner in which mode maps are created and used. These details are discussed in the next section describing experimental series F4.

The scanned images of the resulting composite samples for the F3 series are shown in figure 7.9. All samples were incompletely cured, due to the short processing times. No charring was noted in any of the samples.

Figure 7.9 - Surface Scans (F3.1 - F3.3)

7.3.4 Series F4

The F4 series was designed to demonstrate that an alternate mode-switching policy could improve resultant part quality and as an experiment to determine whether thermal gradients could be reduced further from the levels obtained in the F3 series experiments. An alternate mode-switching strategy was employed. Unlike series F1 through F3, the mode selection strategy used in series F4 estimates the qualitative part state based on only the highest and lowest reported probe temperatures. This approach required the use of new mode maps, described in section 6.2.5.

The experimental summary for series F4 is shown in table 7.7. A 60/30 second hold time was used. The hold time was maintained at 60 seconds (to encourage rapid heating) until the sample reached a minimum processing temperature. When the minimum processing temperature was reached, the hold time was reduced to 30 seconds by the phase change plan fragment included in the F4-series process control plans.

Experiment	Time (min)	T_{max} (C)	$Grad_{max}$ (C)	$Grad_{avg}$ (C)	$Grad_{stddev}$
F4.1	60	131.31	49.80	17.41	6.84
F4.2	60	136.82	38.52	16.30	7.16
F4.3	40	137.22	50.29	12.82	8.16
F4.4	40	138.67	33.20	9.47	4.50

Table 7.7 - Summary of Experimental Series F4

The fabrication results from the F4 series showed a significant improvement in terms of both thermal gradients and heating rates over the series F1 through F3. For example, the thermal profile for F4.1 is shown in figure 7.10. The temperature profiles for each temperature probe show excellent crossing behavior, resulting in a nearly uniform average temperature at each probe location.

Experimental results confirm that the modified F4 mode switching algorithm decreased average thermal gradients in comparison with the gradients observed during series F3. The average thermal gradient for experiment F4.4 was only 9.47 degrees, a 59% reduction from the smallest gradients obtained in the F3 series. All temperature probes were maintained above the temperature required for the onset of reaction (the minimum processing temperature for the DGEBPA material), yet well below the thermal degradation limit. Fabrication run F4.4 (shown in figure 7.11) gave the highest-quality results obtained in this research.

In comparison with the results of previous research, the thermal gradients have been reduced [Fellows 91] by approximately two-thirds. Quantitative comparisons with previous research results are provided in section 7.4.

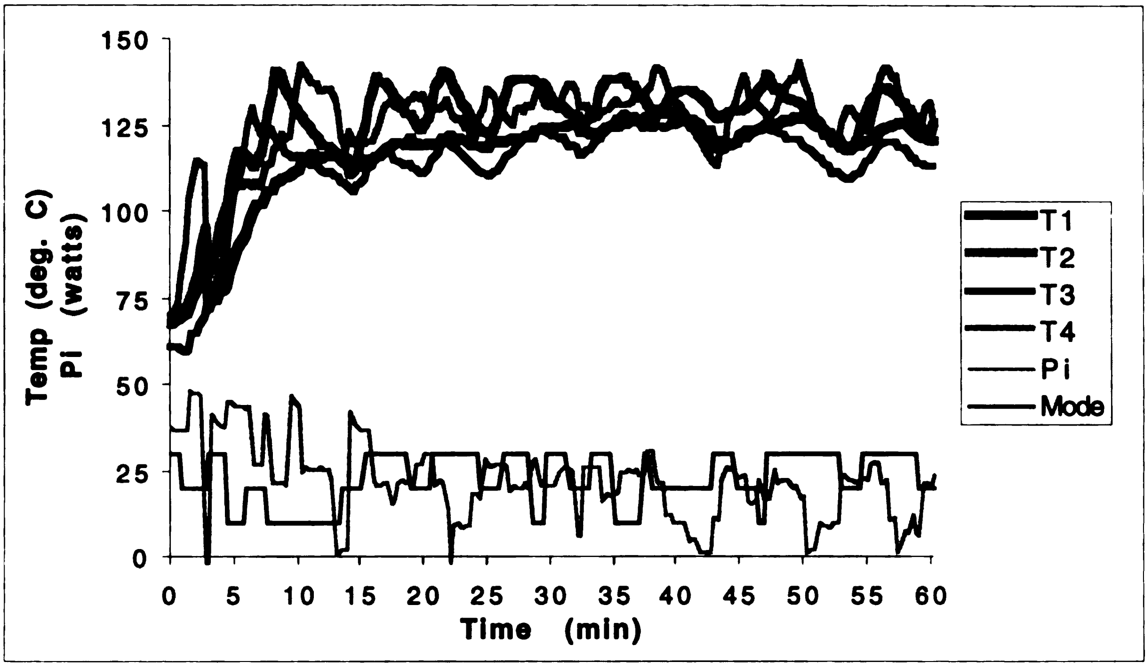

Figure 7.10 - Thermal Profile (F4.1)

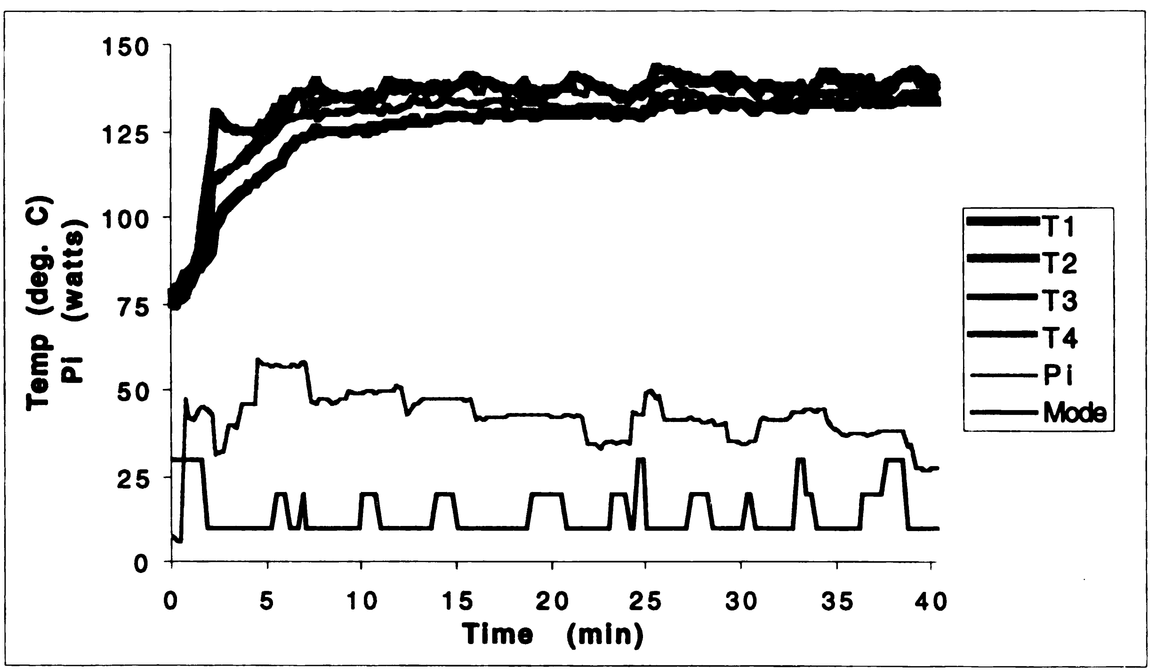

Figure 7.11 - Thermal Profile (F4.4)

The samples from experimental series F4 are shown in figure 7.12. Each sample was uniformly cured with no evidence of charring or large voids. These samples represent the highest-quality materials produced during any of the experimental fabrication runs.

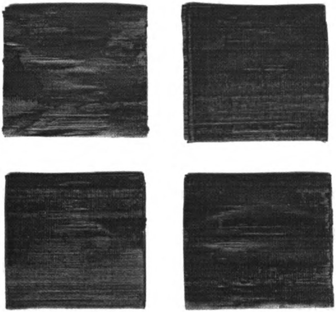

Figure 7.12 - Surface Scans (F4.1 - F4.4)

7.4 Analysis

The C1 series demonstrated that IPCA is capable of producing a variety of plan types for the fabrication and heating of composite materials. The capabilities of the Blueprint Generator and Plan Assembler were effectively demonstrated in this series.

The breadth of the Blueprint Generator and fragment library were demonstrated in the C2 series. The series showed that two materials were supported (DGEBPA/Carbon and Polyester/Glass) and that processing parameters were set appropriately for each material type. The series also demonstrated that power levels were appropriately set for various sample configurations (3" square, 2" square and 1" square unidirectional layups).

Series C3 demonstrated effective control under various hypothetical failure or hazardous situations during processing. An appropriate action was taken in each of the C3 experiments to prevent damage to processing equipment and surroundings.

Experimental series F1 showed that IPCA can fabricate high-quality composite materials by effectively controlling the processing hardware. The performance improved with processing experience, demonstrating the effectiveness of the adaptive mode-switching fragments. Gradients were eventually brought below 22 degrees, which is a 26% reduction over previous automated control approaches exhibiting a 30 degree gradient [Fellows 91]. Although the F1.5 gradient was low, it could have been made lower with more effort. As reflected in the reduced charring of the samples from F1.2 through F1.5, improving fragment control algorithms (in this case, power control) has a direct payoff in resultant cure quality. Additional refinements were made which eventually led to the F4 series experiments and lower thermal gradients.

Series F2 demonstrated the trade-off between thermal gradient and average heating rate. As shown in figure 7.13, there is a direct correlation between observed heating rate, hold time and observed thermal gradient.

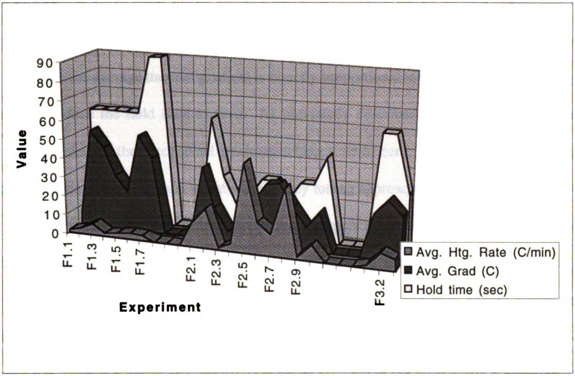

Figure 7.13 - Hold Time, Gradient and Heating Rate Relationships

The correlation between hold time and gradient can be seen from figure 7.13. As hold time increases, so does the average gradient. Similarly, the average heating rate is directly proportional to the hold time and inversely proportional to the average thermal gradient. As hold time increases, so does the heating rate. Also evident from the figure is the reduction in thermal gradients with processing experience due to the action of the adaptive mode-switching fragments.

The correlation coefficient between the average heating rates and average thermal gradients was 0.78 across the entire set of F-series experiments, indicating a very strong correlation between the two data series. The correlation coefficients between the selected hold time and average observed thermal gradients were 0.62 across the entire F-series, which also indicates a strong correlation between hold time and average observed thermal gradients.

The F2 series demonstrates an inadequacy of the mechanical mode-switching apparatus. Since the average thermal gradient is directly proportional to the hold time, it is desirable to reduce the hold time to zero. As previously described, mechanical limitations prohibit reducing the mechanical mode switching time to zero. These important results point out the desirability of a variable frequency tuning approach free of mechanical limitations. If the frequency of the microwave power supply could be varied under the control of IPCA, the hold time could be reduced to essentially zero, thus affording the opportunity to establish new lower limits on thermal gradients during processing. Modes would be selected not by moving metal shorting plates and probes, but by shifting power supply frequencies. This is a promising direction for future research in the microwave fabrication of composite materials.

Series F3 demonstrated that IPCA can process parts of a third geometry (2" square). The heating rates and average thermal gradients observed during the F3 series were intermediate between those observed during the F1 (3" square) and F2 (1" square) series, thus demonstrating that heating rate is dependant on part mass, given limited available processing power. This result demonstrates the scalability of the fabrication technology in terms of the ability to achieve a desired heating rate.

Discarding the median temperature probe measurements and basing the qualitative part state only on the highest and lowest temperatures resulted in superior thermal gradients, compared with series F1 through F3. Once again, the quality of the IPCA control method fragments were shown to have a direct effect on the quality of the resultant fabrication run.

Series F4 demonstrated that thermal gradients can be reduced to very low levels when high-quality complimentary modes are selected for processing and appropriate qualitative part state models are used. The qualitative part state model used in the F4 series proved very useful in reducing gradients below 10 degrees C. This level of performance represents approximately a 66% reduction over previous fabrication methods based on mode-switching and comparable materials [Fellows 91], which produced gradients on the order of 30 deg. C using two complimentary modes.

The original mode switching work [Fellows 91] was not based on motor-driven tuning mechanisms, resulting a longer locking time since the actuators were moved manually. Some of the factors contributing to the decrease in thermal gradients observed in the IPCA research include the availability of an increased number of characterized modes from which to select, a reliance on motor-driven actuators which facilitate more rapid mode switching and the employment of more sophisticated adaptive mode-switching algorithms.

The IPCA-supported fabrication process is also much easier to manage than previous approaches, since IPCA automates the tasks of capturing heating profiles, mode map generation and fabrication control. This level of automation also tends to increase the availability of characterized processing modes at fabrication time, since thermal profiles and mode maps do not have to be generated manually.

Collectively, the C-series and F-series experiments demonstrated the viability, applicability and extensibility of the IPCA architecture. The experiments showed that IPCA can effectively control processes during rapid material heating for a variety of part geometries and can also minimize thermal gradients during fabrication through the selection of appropriate processing parameters and process control strategies.

CHAPTER VIII

8.1 Conclusions

The IPCA architecture has made contributions in two major research areas. In the area of knowledge-based systems and task-specific architectures, the following contributions have been identified:

- **A flexible, task-specific intelligent control architecture.**
- **An extension to task-specific approaches to routine design & planning.**
- **IPCA provides a unique, comprehensive planning and control architecture which combines many of the elements of design, process control and planning research.**
- **IPCA provides a unique automated design, assembly and execution monitoring mechanism.**
- **IPCA provides an extensible mechanism for combining planning and control techniques into a single, unified approach.**
- **IPCA extends the DSPL language for routine design with sophisticated planning and control capabilities.**

Within the domain of composite materials manufacturing, the following contributions have been made:

- **A structured framework for representing knowledge regarding composite materials processing.**

- **An improvement over conventional control techniques for managing single-mode microwave polymeric composites processing.**

- **A prototype for supporting the automated synthesis and execution of process control plans.**

- **A facility has been provided for managing knowledge regarding composite material fabrication technologies.**

- **Reactive planning and monitoring facilities have been provided to enable the production of high-quality composite structures using mode-switching techniques.**

- **The limitations of mechanical mode switching have been investigated. Variable-frequency tuning is proposed as a promising future research area.**

- **Considerable reduction of thermal gradients over previous automatic control approaches.**

- **Demonstration of feasibility and applicability of the IPCA approach to the control of high-speed polymer processing.**

The IPCA architecture provides a framework for capturing knowledge about composite materials fabrication technologies. The Blueprint Generator contains knowledge about many processes and their potential applicability. The system can help to choose appropriate fabrication technologies given a list of processing requirements. Additionally, the IPCA implementation provides the first intelligent, automated real-time reactive monitoring facility for the microwave fabrication of composite structures.

The limitations of the IPCA architecture include dependance upon the routine design framework on which the Blueprint Generator is based and the sequential nature of plan and control method selection. The IPCA architecture is limited to use in applications for which the DSPL-based Blueprint Generator can produce a suitable plan blueprint. This class of applications extends beyond composite material fabrication into the more general area of process control. Additionally, the architecture may not be well-suited to controlling simultaneous processes with few interacting goals. The IPCA plan execution monitoring facility is designed for process control applications in which goals may interact or conflict and a central, cyclic evaluation mechanism is more efficient than multiple control plans executing in parallel.

8.2 Future Work

The IPCA architecture and prototype could be extended in the following areas:

- **Addition of IPCA support for MDSPL multiple blueprint generation.**
- **Interactive replanning capability between the Plan Execution Monitor and Plan Generator.**
- **Investigation of variable-frequency tuning and the associated performance limitations.**
- **Incorporation of increasingly advanced control method and planning fragments.**
- **Enhancing the depth and breadth of materials and processing technology coverage.**

APPENDIX A
Glossary

Abstract parameter

One of four types of parameters which may be included in a process control plan. Abstract parameters have their values updated by the Observation Unit on each execution cycle, immediately prior to Execution Unit fragment selection. See chapter 5.

Blueprint

The non-executable design artifact produced by the Blueprint Generator. The blueprint contains a specification of the control method fragments, plan fragments and goals which are deemed useful in controlling a specific process. Blueprints are used by the Plan Assembler to create executable process control plans.

Blueprint Generator

A central component of the IPCA architecture. The Blueprint Generator is implemented as a DSPL problem solver. The task of the Blueprint Generator is to produce blueprints from fabrication models. See chapter 5.

Composite material

A material composed of two of more constituents which do not co-mingle at the molecular level. Composite materials generally consist of a matrix and a reinforcement. See appendix E.

Constraint

A DSPL construct, used to determine whether a portion of an overall design was successful. Constraints may be interspersed between plans, tasks or steps in a DSPL implementation to control problem-solving. See chapter 4.

Control Interface Unit

An important portion of the IPCA architecture used to coordinate communications with external control systems, sensors and actuators. The Observation Unit relies on services provided by the Control Interface Unit to acquire sensor data. The Plan Execution Unit also relies on the Control Interface Unit for directing output to appropriate physical actuators. See chapter 5.

Control method fragment

One of two types of fragments included in the fragment library. Control method fragments may directly influence output parameters but cannot alter processing goals. Many such fragments are typically included in a single process control plan. See chapters 2, 4 and 5.

Crossply

A particular type of layup in which the fibers in alternating layers are non-parallel. A crossply layup may be used to encourage more isotropic material properties than those which might be achieved through the use of a unidirectional layup. See chapter 3 and appendix E.

Delamination

A generally undesirable failure state of a composite material in which the layers (plies) in the composite peel apart. See chapter 3.

DGEBPA

Diglycidyl esther of bisphenol A, a polymeric composite matrix material. This material was used for the fabrication experiments reported in chapter 7. See chapter 3 and appendix E.

DSPL

Design structures and plans language, a GT-based approach to routine design and planning problems. A DSPL-based tool was used to implement the Blueprint Generator. See chapter 4.

EM field

An alternating vector field consisting of electric and magnetic components. An EM field is maintained in the cavity during material processing to deliver energy into the material. See chapter 3, 6.

Execution Unit

The execution unit is responsible for coordinating fragment selection and execution within the IPCA architecture. See chapter 5.

External Controller

A low-level feedback control system used to manage mode tuning and data gathering operations. IPCA interfaces with this external system through services provided by the Control Interface Unit. See chapter 3 and appendix D.

Fabrication model

Used as input to direct the blueprint generation process, the fabrication model contains specific information regarding the material to be processed and the availability and configuration of processing equipment. See chapter 4.

Fiber

One type of reinforcement commonly used to strengthen polymeric composites. Often, fibers are composed of carbon, graphite or glass. See also, reinforcement. See appendix E.

Fragment

A generalized, reusable control method or planning algorithm which can be assembled with other fragments to create many distinct control strategies when combined with appropriate goals specified by the Blueprint Generator. The Plan Assembler assembles fragments into executable process control plans. See also plan fragment, control method fragment. See chapter 4.

Fragment library

The library of fragments which is managed by the Fragment Librarian.

Fragment Librarian

An important portion of the IPCA architecture responsible for the storage, retrieval and maintenance of plan and control method fragments. See chapter 5.

Generic Task approach

A task-specific approach to knowledge-based systems used as a basis for the IPCA architecture. See chapters 2 and 4.

Gradient (thermal)

A temperature differential which is undesirable during composite material processing. Large gradients result in uneven heating and reduced properties in the finished composite structure. See chapter 3.

GT approach

See generic task approach.

Input parameter

One of four types of parameters managed by the Observation Unit. Input parameters are used to represent sensor values obtained during real-time execution monitoring. Input parameters are updated by the Observation Unit (through communication with the Control Interface Unit) at the start of every execution cycle. See chapter 5.

IPCA

Intelligent Process Control Architecture, the intelligent control architecture which is described in this document. See chapters 1, 2 and 4.

Knowledge-level

A distinct problem-solving level proposed by Alan Newell, advocating an implementation-independent description of complex knowledge-based problem-solving tasks. See chapter 4.

Labview

A visual programming language and development environment by National Instruments. Designed to support instrument control, data collection and general information processing. Used to implement the external control system. See chapter 3.

Layup

A method for preparing composite structures in which layers of material are (usually) hand-laid into a specific form (such as a mold) in preparation for processing. See chapter 3, appendix E.

Matchgroup table

A table of two or more columns and one or more rows, used by sponsors to indicate the utility of a sponsored item under various conditions. In the IPCA architecture, matchgroup tables may contain wildcards in addition to the usual binary values. The last column in every matchgroup table contains ratings which are returned with a corresponding row matches. See chapter 4.

Matrix

One of the constituents of a composite material. The function of the matrix is to transfer loads to the reinforcement and to protect the reinforcement from environmental damage. Matrices are frequently polymeric, metal or ceramic. See appendix E.

Microwave-based composite material processing

An alternate to conventional autoclave-based composite material fabrication which has the potential to offer advantages in terms of speed, material properties, energy efficiency and controllability. See chapter 3.

Mode

A tuned, stable EM standing wave pattern within the single-mode cavity. Each mode typically heats a given composite structure in a unique manner. By changing modes, thermal gradients may be effectively controlled. See chapter 3, 6.

Observation Unit

The portion of the IPCA architecture which is responsible for monitoring sensors and updating parameters based on the resultant values. See chapter 5.

Output parameter

One of four parameters managed by the Observation Unit. The values of all output parameters are conveyed by the Observation Unit to the appropriate actuators through the services of the Control Interface Unit at the completion of each execution cycle. See chapter 5.

Plan (DSPL)

A DSPL construct used to specify the tasks which carry out a particular portion of the overall design. Each specialist contains a list of plans which may be used to carry out a portion of the overall design. See chapter 4.

Plan Assembler

The portion of the IPCA architecture responsible for assembling selected fragments into complete, executable process control plans. The Plan Assembler requests specified fragments from the Fragment Librarian, sponsors them and integrates them into a hierarchical sponsor-selector structure. See chapter 5.

Plan Execution Monitor

The IPCA architectural component responsible for process control plan execution and real-time process monitoring. The Plan Execution Monitor is comprised of the Execution Unit, Observation Unit and Control Interface Unit. See chapters 4 and 5.

Plan fragment

One of two types of fragments managed by the Fragment Librarian. Plan fragments may modify processing goals, but they do not directly control actuators or other control outputs. See also, fragment. See chapters 2, 4 and 5.

Plan generator

The plan generation subsystem of the IPCA architecture, consisting of the Blueprint Generator and Plan Assembler. See chapter 5.

Ply

A single layer in a multi-layered composite structure. See chapter 3 and appendix E.

Process control plan

The executable artifact produced by the Plan Assembler. The process control plan contains fragments, processing goals and sponsor-selector structures used by the Execution Unit to choose appropriate fragments for execution in real-time. See chapters 4 and 5.

Reactive planning

An approach to knowledge-based planning which attempts to address the problems associated with plan execution in dynamic environments. Reactive planners usually support interleaved, incremental planning and plan execution in an attempt to provide rapid response to unexpected events. See chapter 2.

Reinforcement

One of two composite material constituents. Reinforcements give composite materials strength. Common reinforcements include fibers, particulates and whiskers. See appendix E.

Selector

One of the components of a sponsor-selector mechanism. Selectors are used to group and select among related sponsors. See chapter 4.

Single-mode tuned cavity

A particular type of waveguide (cavity) in which a variety of standing EM field patterns may be produced and maintained. The interior volume of the prototype cavity (described in chapter 3) is variable, thus supporting the mode tuning process. See chapter 3, 6.

SmallTalk

A high-level object-oriented programming environment used to implement the IPCA prototype.

Specialist (DSPL)

A DSPL construct used to coordinate problem-solving for a portion of the overall design task. Each specialist contains a set of plans from which to select to accomplish the design process for a particular subcomponent and is charged with the responsibility for successfully designing that portion of the overall artifact. See chapter 4.

Sponsor-selector system

A mechanism used to support knowledge-based selection by using sponsors, selectors and matchgroup tables. Used by IPCA within the Blueprint Generator and in the process control plan. See chapter 4.

Sponsor

One of the components of a sponsor-selector mechanism. Sponsors rate the applicability of associated selection items (typically fragments) through the use of matchgroup tables. See also sponsor-selector system. See chapter 4.

Static planning

An approach to knowledge-based planning suited for non-dynamic environments. See chapter 2.

Step (DSPL)

The lowest-level action in DSPL, involving setting a single design parameter to a particular value. See chapter 4.

Task (DSPL)

A collection of steps and interspersed constraints used to support the design of a particular subcomponent of an overall design. See chapter 4.

Thermoset composite

A composite with a thermosetting (not thermoplastic) matrix. Thermosetting materials undergo crosslinking reactions which are generally exothermic and non-reversible. See appendix E.

Variable parameter

One of four types of parameters supported by the Observation Unit. Variable parameter values are calculated during each execution cycle from methods included in the process control plan. See chapter 5.

Void

An undesirable "bubble" of trapped water vapor within a cured composite material which generally serves to adversely effect mechanical properties. See appendix E.

APPENDIX B
IPCA User's Manual

The IPCA software prototype is described in this section. An overview of the features of each architectural component implementation is presented in a chronological progression, beginning with launching the system and proceeding through blueprint generation, assembly and execution.

The IPCA software implementation will run on any platform supported by ParcPlace VisualWorks™ release 1.0 and Objectworks\Smalltalk™ release 4.1 [ParcPlace 92]. Supported operating systems include Apple Macintosh™, Microsoft Windows™, IBM RS/6000™, and UNIX™. Although the Macintosh™ platform was used to generate all figures and examples, the use and appearance of the software is nearly identical across all platforms. Throughout this section, references to visual interface components (windows, buttons, selection items, menus) are given in ***italic bold***.

Launching IPCA

The IPCA prototype is launched from the Objectworks\Smalltalk™ environment by using the ***Finder*** tool as depicted in figure B.1. To launch IPCA, the class ***AC Launcher*** is selected and the ***Start*** option is selected. This creates an instance of the ***AC Launcher*** class, shown in figure B.2. This visual component is called the ***IPCA Launcher***.

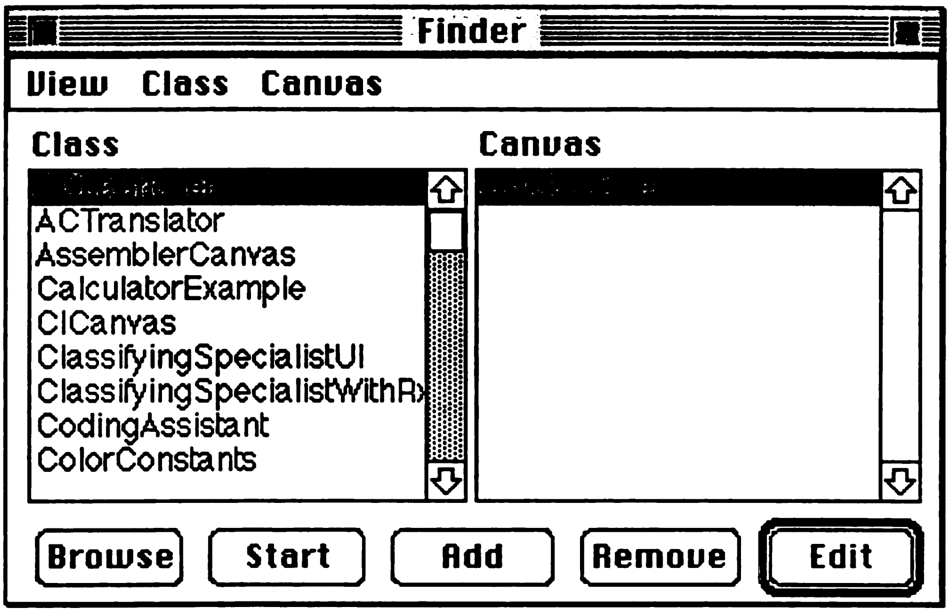

Figure B.1 - Starting IPCA using the VisualWorks Finder

The ***IPCA launcher*** interface is used to activate the Blueprint Generator, Plan Assembler and Plan Execution Monitor. The launcher serves simply as a tool to facilitate the creation of an instance of a selected software component. Normally, the components are launched in sequence to generate, assemble and execute process control plans. However, the launcher allows flexibility since multiple components of a particular type may be launched simultaneously and in any order. Closing the launcher interface causes the entire IPCA software system to terminate, as the launcher is the top-most component of the IPCA user interface.

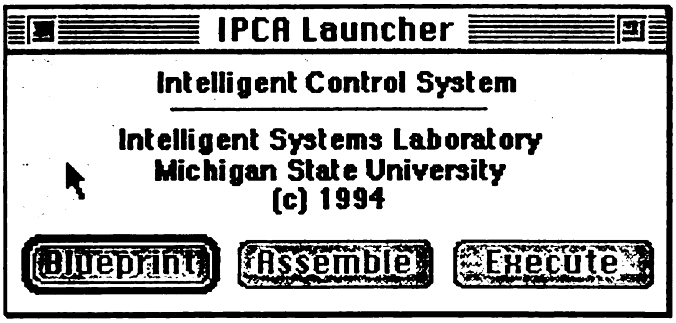

Figure B.2 - Launcher Interface

Blueprint Generation

Selecting the ***Blueprint*** option causes the top-level ***Blueprint Generator Interface*** shown in figure B.3 to appear (the interface is titled ***RD Interface*** because some of the visual components of IPCA are shared with other SmallTalk-based tools which preceded it, namely, the DSPL-based Routine Designer).

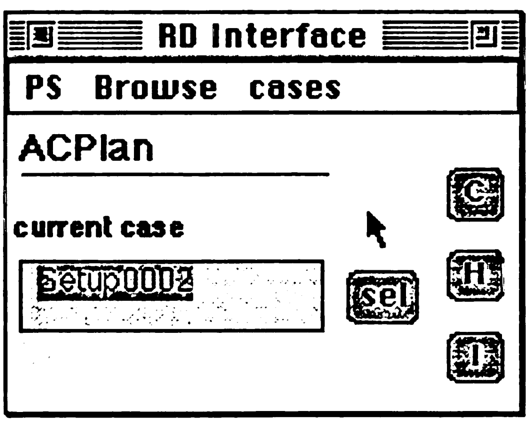

Figure B.3 - Blueprint Generator Interface

The drop-down *PS* menu shown in figure B.4 reveals options for running the Blueprint Generator.

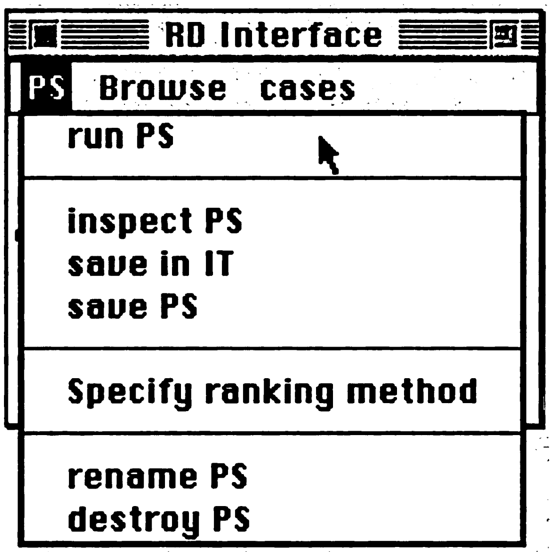

Figure B.4 - Blueprint Generator Problem-Solver (PS) Menu

The ***Browse*** menu allows access and modification to the structure of the DSPL problem-solver as shown in figure B.5. Options are included to support adding, deleting, renaming and moving specialists, plans, tasks, steps and constraints.

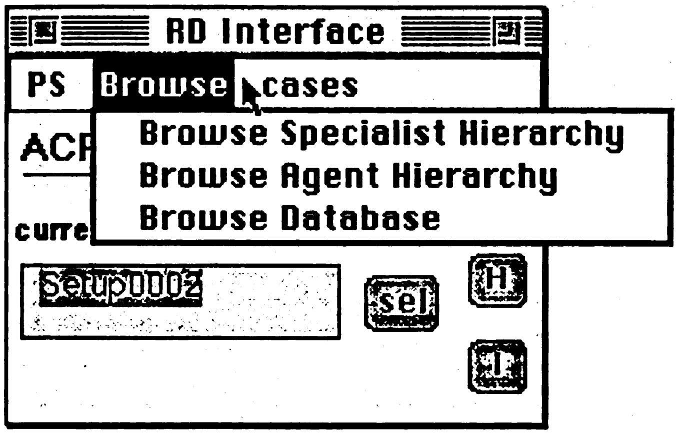

Figure B.5 - Blueprint Generator Browse Menu

The Blueprint Generator ***Cases*** menu, shown in figure B.6, allows access to IPCA fabrication models through the use of a sophisticated database interface. It is through these services that the fabrication models are maintained. The fabrication models are referred to as cases throughout the ***Blueprint Generator*** Interface.

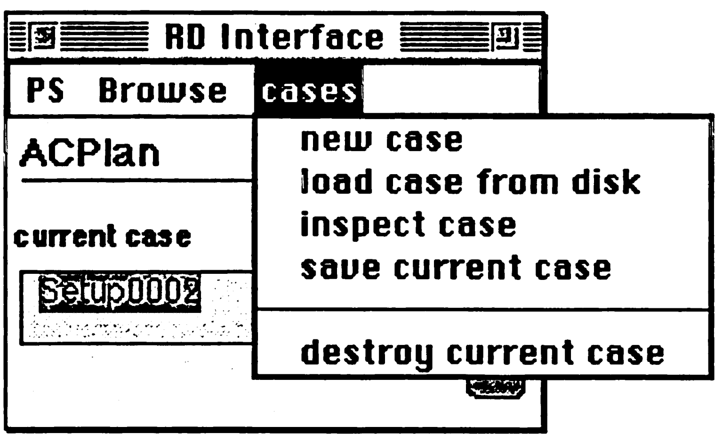

Figure B.6 - Blueprint Generator Cases Menu

The process of blueprint generation begins with selecting the ***run PS*** option from the

PS menu. Immediately after selecting the ***run PS*** option, the menu shown in figure B.7 appears. The routine design tool upon which the Blueprint Generator is based can support single or multiple design [Kamel 94]. The prototype Blueprint Generator is designed to work only with the ***Single design*** option, so this option should always be selected from the resulting menu.

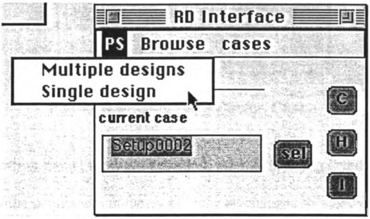

Figure B.7 - Specifying Design Type

Upon selecting the ***Single design*** option, the menu shown in figure B.8 appears. From this menu, it is possible to specify the case (fabrication model) which will be used for blueprint generation. If the desired fabrication model (***Setup0002*** in this example) is displayed, the ***Run Current Case*** option may be selected, otherwise the ***Run New Case*** option is selected and a fabrication model is specified.

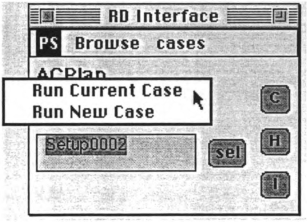

Figure B.8 - Specifying Design Cases

Once a fabrication model has been specified, the model is displayed as in figure B.9. Values for any of the parameters in the model may be changed by selecting the parameter of interest. This gives an opportunity to change the inputs to the Blueprint Generator without altering the stored copy of the fabrication model. The blueprint generation process is started by selecting the ***Proceed*** option.

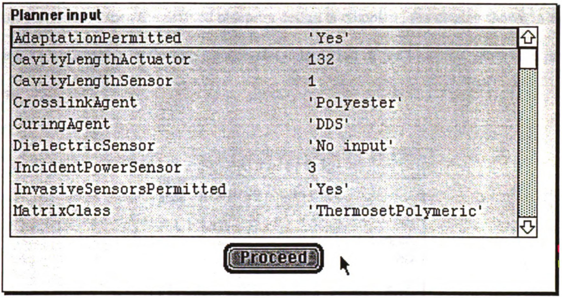

Figure B.9 - Reviewing the Selected Fabrication Model

Some time will elapse while the Blueprint Generator completes the rough blueprint design. Generally, the rough design takes less than a minute to complete on most platforms. When the rough design pass is complete, the display shown in figure B.10 is presented.

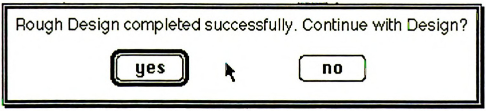

Figure B.10 - Blueprint Generator Rough Design Complete

Selecting the ***yes*** option will cause the Blueprint Generator to continue with the detailed design. Selecting ***no*** will cause the Blueprint Generator to abort the blueprint generation process. After selecting ***yes***, another brief pause will be observed as the Blueprint Generator generates the detailed blueprint. This process takes about a minute on most platforms. When the detailed blueprint design is complete, the display shown in figure B.11 will appear, indicating a successful completion of the blueprint generation process.

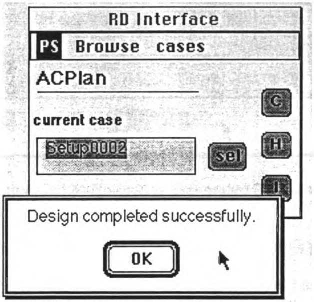

Figure B.11 - Blueprint Generator Design Complete

Selecting ***OK*** will cause the display in figure B.12 to appear.

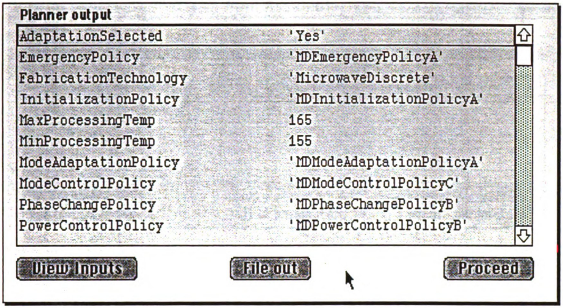

Figure B.12 - Reviewing Generated Blueprints

The blueprint can then be reviewed. The steps which caused a particular value to be set may be viewed by selecting a particular key/value pair from the list. Selecting ***File out*** causes the completed blueprint to be written to disk. The ***Proceed*** option may then be used to return to the top-level ***Blueprint Generator*** Interface shown in figure B.3. The ***Blueprint Generator*** Interface may then be closed and the ***assemble*** option selected from the IPCA launcher.

Assembling Process Control Plans

Selecting the ***assemble*** option causes the top-level user interface for the Plan Assembler to appear. The ***Assemble*** option invokes the assembly process, and the ***Library*** option invokes the Fragment Librarian interface shown in figure B.14.

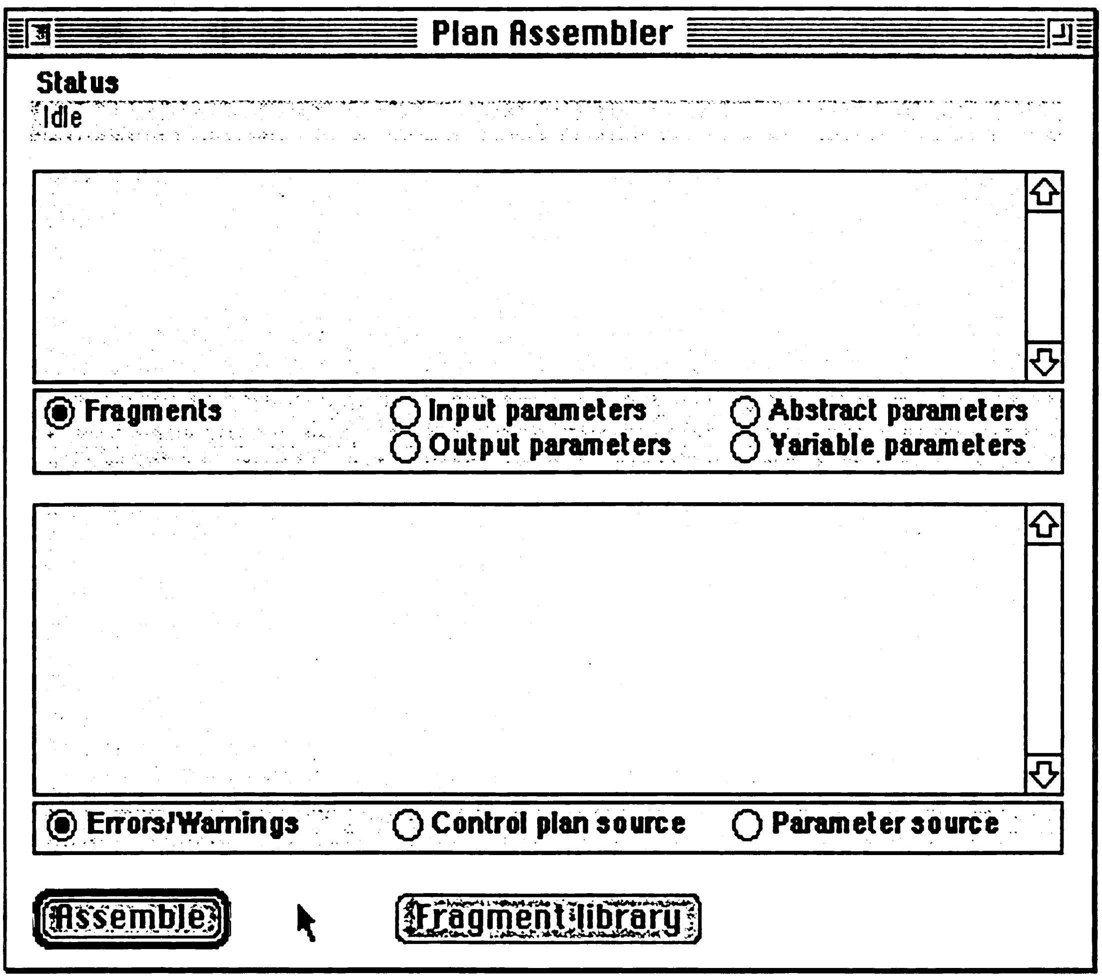

Figure B.13 - Plan Assembler Interface

In figure B.14, the element ***Plan-Source*** of the fragment ***Master*** in the library ***ModeSetsSetup0002*** is displayed. This element contains source code for the selected fragment, which can be edited and placed into the fragment library. All fragment changes are either incorporated or rejected when closing the ***Plan Fragment Librarian*** interface. To accept all changes and make them permanent, close the ***Plan Fragment Librarian*** interface. To abort changes and leave the fragment library untouched, choose the ***Abort*** option.

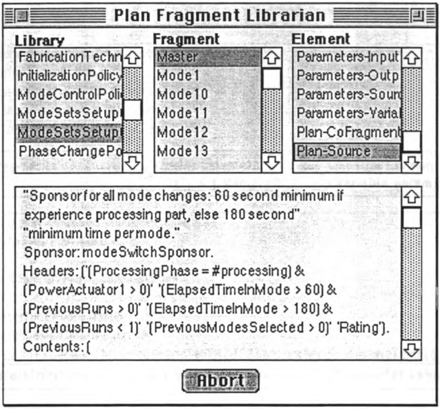

Figure B.14 - Fragment Librarian Interface

The ***Assemble*** option (shown in figure B.13) causes the current blueprint to be assembled. After assembly, the ***Plan Assembler*** interface appears as in figure B.15, showing information about the assembled process control plan, including lists of fragments and parameters which were included in the plan. The resulting process control plan and parameter source can be viewed in the lower scrollable display area, along with information about any errors or warnings produced during plan assembly.

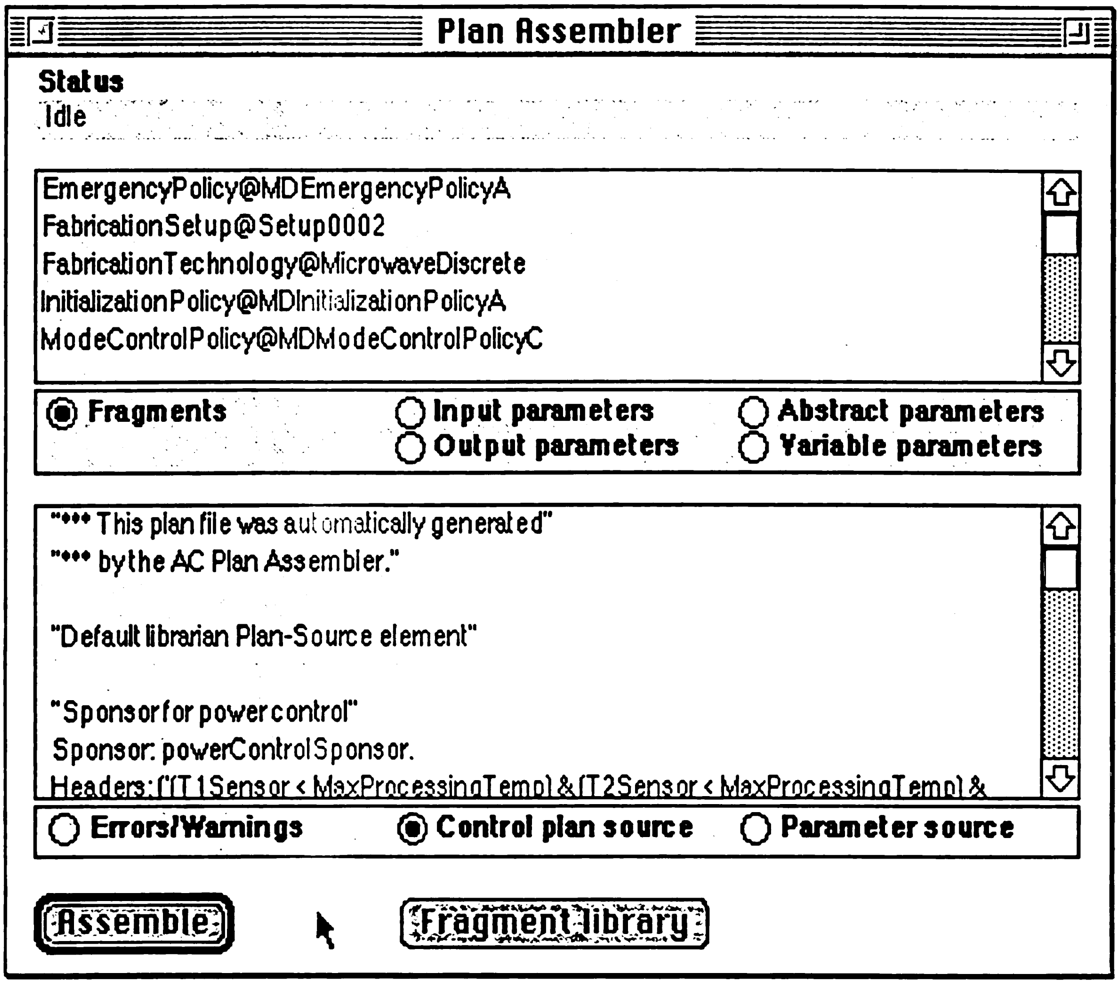

Figure B.15 - Completed Plan Assembly

Plan Execution

After plan assembly is complete, the ***Plan Assembler*** interface is closed and the ***Execute*** option is selected from the ***IPCA Launcher,*** causing the ***Execution Monitor*** interface shown in figure B.16 to appear. The Execution Unit, Observation Unit, and the Control Interface Unit are simultaneously launched from the ***Execution Monitor*** interface. Selecting the ***Go!*** option causes the Execution Unit, Observation Unit and Control Interface Unit to launch immediately using default process control plan files. By using the ***Open*** and ***Launch*** options, other process control plans can be selected instead of the default (newest) instances.

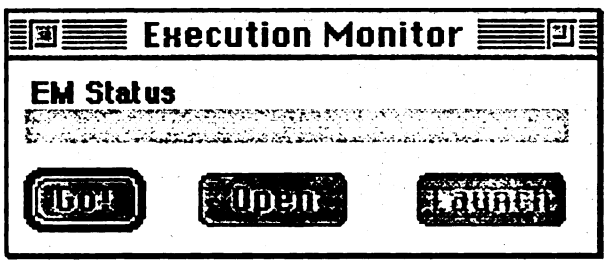

Figure B.16 - Execution Monitor Interface

Figure B.17 shows the ***Execution Unit*** interface. The number of execution cycles are displayed, along with the cumulative number of sponsor evaluations and average execution cycle period. The current (winning) sponsor and bid are displayed. The ***Step*** option is used to single-step execution while the ***Resume*** and ***Suspend*** options are used to start and stop execution, respectively.

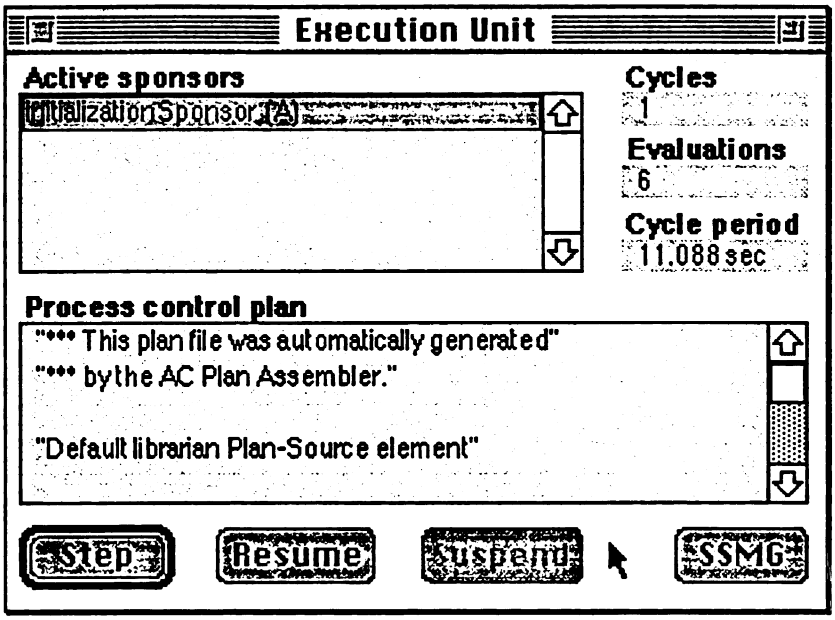

Figure B.17 - Execution Unit Interface

The *SSMG* option invokes the ***Sponsor-Selector Match Group*** browser shown in figure B.18, so that the structure of the process control plan may be examined. The current selector is displayed, together with a list of sponsors grouped under the selector, the match group and plan or control method fragment. The >> and << controls are used to browse through different levels in the process control plan hierarchy.

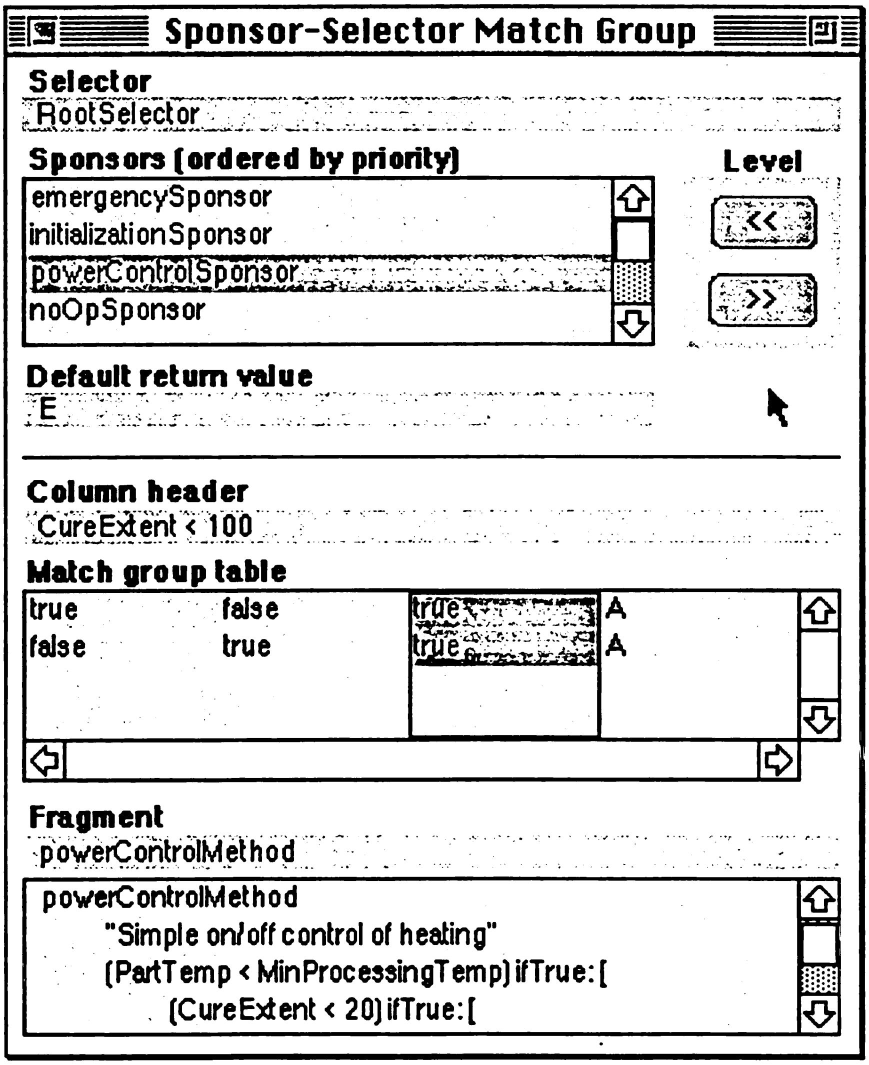

Figure B.18 - Sponsor-Selector Match Group (SSMG) Browser

Figure B.19 shows the ***Observation Unit*** interface. Input, output, abstract and variable parameters are displayed with their most recent values and corresponding timestamps.

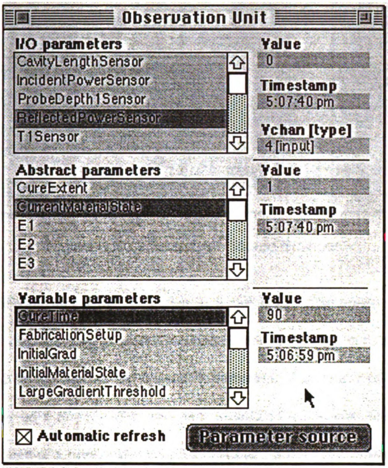

Figure B.19 - Observation Unit Interface

The content of the parameter file can be displayed by selecting the ***Parameter source*** option, as shown in figure B.20.

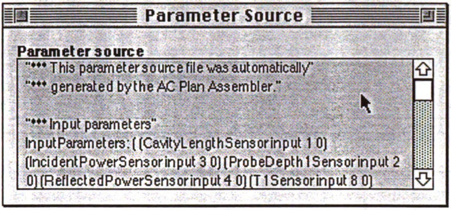

Figure B.20 - Parameter Source Browser

The ***Control Interface Unit*** is shown in figure B.21. The upper scrollable area, labeled ***Source***, displays a list of sensed physical input channels and their most recent values, as reported by the ***Source*** for sampled information. The source of sensed process information may be a ***Simulated*** control unit (described in the next section), the ***Live*** External Controller, or a previously-recorded data ***Log***. Similarly, the ***Sink*** area displays the most recent output channel designators and values, which can be directed at the ***Simulated*** control unit, the ***Live*** control system, or ***Both***. The checkboxes labeled ***Graphics*** and ***Log CI activity to:*** are used to display the part graphics interface shown in figure B.23 and to log all Control Interface Unit activity to a specified file, respectively.

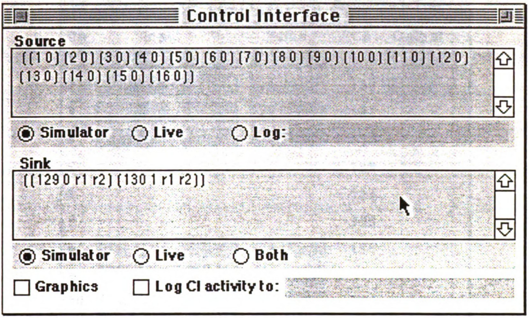

Figure B.21 - Control Interface Unit Implementation

The ***Simulated CU*** (control unit), shown in figure B.22, is used to provide a limited facility for testing process control plans and fragments without the need for physical control system hardware. The interface appears when ***Simulator*** is selected as the ***Source*** for the Control Interface Unit. The controls may be used to simulate any sensor values required to test the behavior of any process control plan of interest.

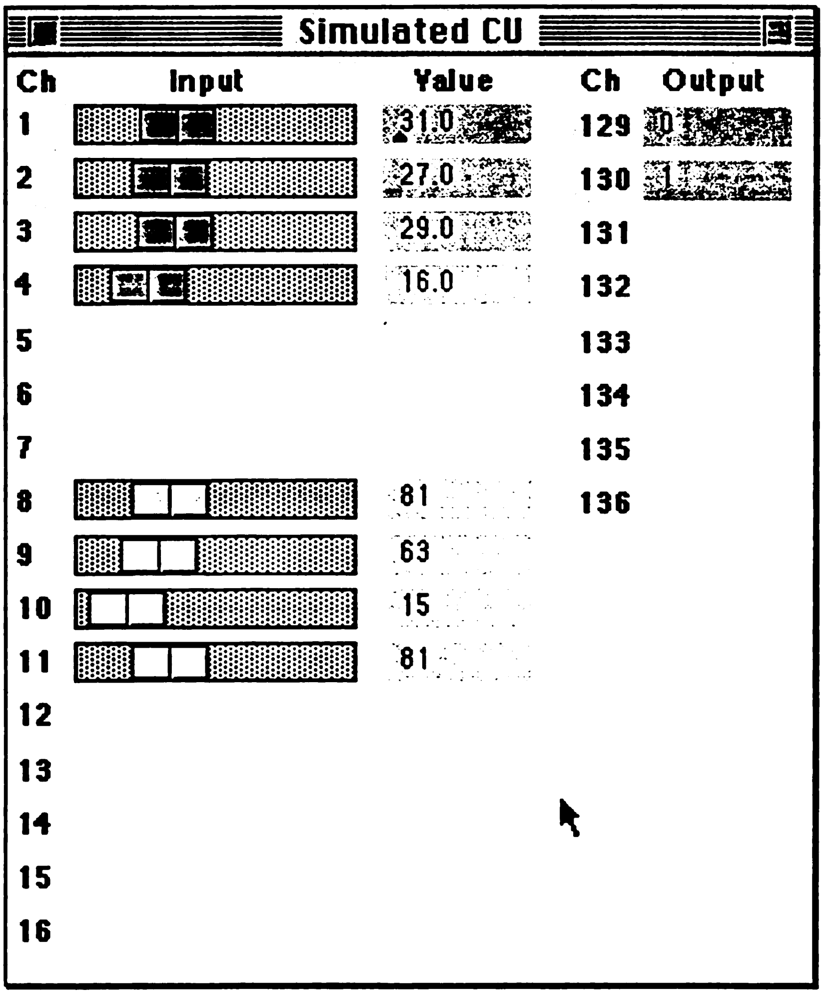

Figure B.22 - Simulated Control Unit (CU) Interface

During a simulation, process inputs are obtained directly from the simulated control unit. By selecting ***Simulator*** as the ***Sink*** for the Control Interface Unit, output parameter values may be monitored as well, rather than being sent to the External Controller. Only the virtual channels which are utilized by the process control plan are displayed. Also, note that the simulated control unit is not a simulation of the process, but a tool intended for debugging process control plans and for permitting plan execution in the absence of processing hardware.

The ***Part model graphics*** display is shown in figure B.23. This display appears whenever the ***Graphics*** box on the ***Control Interface*** is checked. The display is intended to graphically illustrate the current state of the controlled process. In the upper portion of the display, the positions of the upper shorting plate and probe are shown, together with a bar-graph indicator for selected power level and a pair of simulated analog meters for monitoring incident and reflected power. In the lower portion of the display, the temperature of the part is shown though the use of color. Absolute and relative temperatures are displayed, with blue indicating the coldest regions and red indicating the warmest.

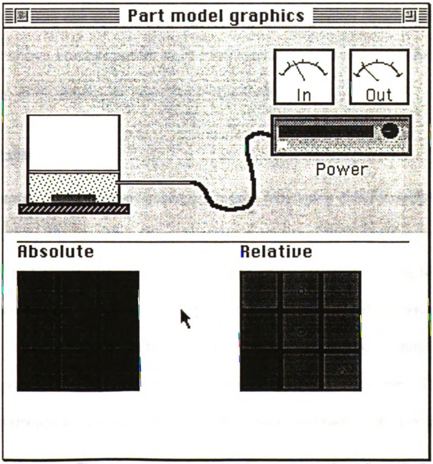

Figure B.23 - Part Graphics Interface

APPENDIX C
IPCA Software Utilities

The IPCA implementation is supported by a set of software utilities. A log translator utility is used to translate logged Control Interface Unit files into a format suitable for importing by major mathematical, statistical analysis, visualization or spreadsheet software packages. A mode-map generation utility is used to produce mode-map fragments from heating profile data sets collected through mode characterization (heating) experiments. This utility allows heating profiles to be automatically incorporated into the IPCA fragment library. Both utilities are described in this section.

Log Translator

The log translator utility is used to prepare logged fabrication data for analysis. Control Interface Unit data is recorded during fabrication whenever the ***Log CI activity*** option is selected on the ***Control Interface*** window. Each Control Interface Unit log file contains a complete summary of the input and output parameters for a specific fabrication experiment. The values of all input and output parameters are logged at the completion of each execution cycle, giving a partial record of the fabrication run. The values of abstract and variable parameters are not logged in the current implementation of the Control Interface Unit.

The log translator utility extracts sensor information from a specified Control Interface Unit log file and translates it into tab-delimited ASCII format suitable for import to a number of statistical analysis, mathematical visualization, spreadsheet or database packages for further analysis. The graphs included in chapter 6 (showing temperature profiles during specific fabrication experiments) were produced by importing translated Control Interface Unit log files into Microsoft Excel™.

The ***Log Translator*** interface shown in figure C.1 appears when the Visualworks\Smalltalk™ ***Finder*** is used to ***Start*** an instance of the ***AC Translator*** class. Selecting the ***Translate*** option causes the dialog in figure C.2 to appear.

Figure C.1 - Log Translator Interface

The name and location of the file to be translated must be supplied. The translator utility may take a minute or more to translate the file into tab-delimited ASCII format. A dialog showing a ***Translation Complete*** message then appears. Selecting ***OK*** causes the interface to return to the state shown in figure C.1. Another file may then be translated, or the interface may be closed.

Figure C.2 - Log Translator File Specification

The translated file name is based on the original log file name, but includes the suffix ".xl". In the example shown, a file named "SamplePlan.log1.xl" would be produced, leaving the original file untouched. The translated file appears on the same volume and in the same folder as the original.

Mode Map Generator

The mode map generator utility is used to produce mode-switching control fragments from previously logged heating information (typically, log files produced by the Execution Monitor). The utility accepts tab-delimited (translated) input with any number of columns and rows. Each row typically contains information regarding temperature and other sensor values at a particular moment. The series of rows thus reveals a history of sensor values across some timespan.

The mode map generator functions by performing an analysis of the heating characteristics of each mode included in the logged heating information. A set of temperature differentials is first computed for each mode included in the logged heating information.

Each differential is computed by subtracting the initial temperature at each probe (upon first entering the mode) from the final temperature for the same probe. The differentials are then sorted, producing an ordered list. For example, the list might be (T1 T3 T4 T2) for a particular mode, say, mode 1, indicating that T1 heated the most, followed by T3, T4 and T2.

The list is then used in a matrix multiplication operation involving a representation of all possible qualitative part states to create the resulting mode 1 map. The multiplication weights each of the matches between the qualitative states and the position of each selected probe in the list. For example, if a part is in the qualitative state (T2 T4 T3 T1), T2 is hottest and T1 is coolest. This would be an ideal match for the differential observed above, since T1 would be heated most and T2 would be heated least. The matchgroup table row for this mode and qualitative state would be filled in with a very high rating. However, the rating in the state row (T1 T3 T4 T2) would be quite low, because the mode heats T1 the most and T2 the least. Therefore, the row for this mode and qualitative state would be filled in with a very low rating.

The user interface works almost identically to that of the log translator. The column numbers corresponding to the logged temperature information must be provided as shown in figure C.3. If four temperatures were monitored, T5-T8 must be set to zero. For example, figure C.3 shows that information on T1 is to be taken from column 3 in the input, T2 is found in column 4 and so on. The fixed/variable radio buttons are used to specify whether a map for a particular mode is to be generated (all the data belongs to a particular "fixed" mode) or whether the data includes information gathered from multiple (variable) modes. If "variable" is selected, the column which specifies the mode number must be included (column 7 in this example). If "fixed" is selected, the mode to which all the temperature data pertains must be specified instead of a column number, so that the appropriate numbered map may be produced.

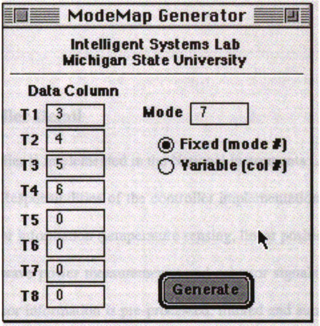

Figure C.3 - Mode Map Generator Interface

After specifying data columns and mode type/column/number, the ***Generate*** button is used to initiate the mode generation process. A file dialog box immediately appears and a filename is specified for the data file. The data is read, processed and the resulting mode maps are written to a file with the same name, in the same folder, but with a ".maps" extension. The file contains the complete fragments for the appropriate mode(s). The resulting mode map fragments may then be pasted into the fragment library.

APPENDIX D
External Controller and IPCA Interface Protocol

External Controller Detail

The External Controller is implemented in the National Instruments LabVIEW™ development environment. Responsibilities of the controller implementation include monitoring and processing sensor information (temperature sensing, linear position sensing, incident and reflected microwave power measurements) and actuator signals (power control and mode control). Sensor information is pre-processed, filtered and summarized before it is reported to IPCA. The External Controller also provides mode-tuning services. Any processing modes requested by the IPCA system are tuned, maintained and locked. The controller attempts to maintain a tuning lock on any requested mode by seeking a predetermined probe and cavity shorting plate position and then monitoring and minimizing reflected power in a predetermined linear region of operation. The tuning lock is accomplished and maintained through incremental probe and cavity length adjustments designed to find the precise location of the mode and maintain it through changes in dielectric properties of the composite. The External Controller interface is shown in figure D.1.

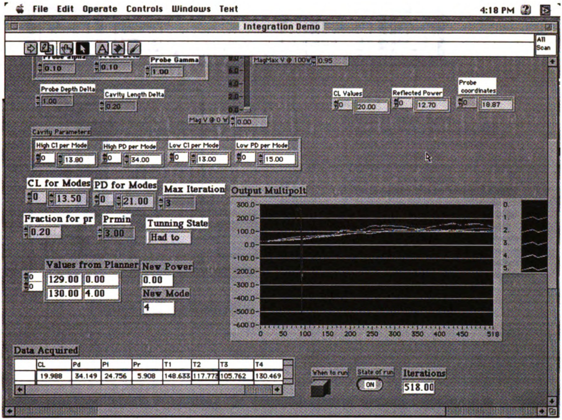

Figure D.1 - External Controller Interface

From figure D.2, it is clear that two computer systems are used for operating the cavity and controlling the fabrication process. One machine runs the intelligent control system implementation while the external control system software executes on the second machine. The two machines communicate through the use of a network cable.

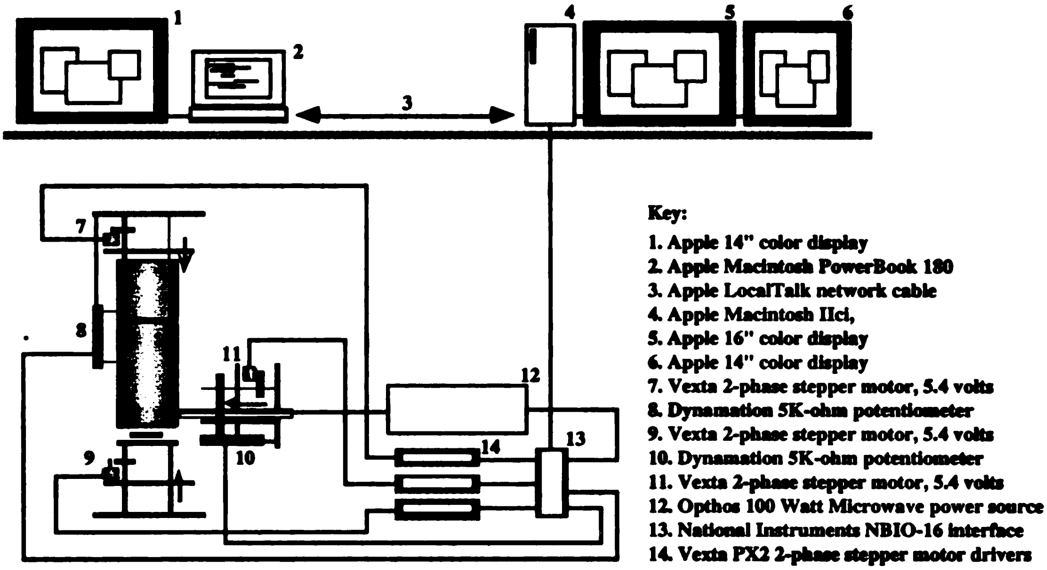

Figure D.2 - Microwave Equipment and Configuration

The machine which is running the external control system software has a National Instruments NBIO-16™ I/O interface installed [National xx]. This interface provides a number of analog and digital input and output channels which are used for monitoring input power, reflected power and cavity and probe positions. The interface is also used to control the sample loading, cavity tuning and probe tuning mechanisms and the microwave power level.

A shared RAM disk is maintained on one of the machines. A RAM disk was chosen because it was desirable to make shared access as fast as possible. The shared area is used to pass messages between the two systems. The implementations on both machines post and check for messages on the shared disk to facilitate communication between the Plan Execution Monitor and the External Controller. This implementation has advantages of simplicity, network independence and machine independence.

The power system wiring is shown in figure D.3. The equipment on the left side of the figure is used during fabrication runs. The output of the 100W, 2.45GHz fixed-frequency power source is connected to a coaxial directional coupler which allows incident power measurement (at -20dB) while allowing the remaining power to flow directly through the device unattenuated. Power then flows through a microwave circulator, a solenoid-actuated coaxial switch and into the microwave coupling probe. Power reflected back from the cavity (not coupled into the sample) flows from the probe, back through the switch, into the circulator and finally into a 50 Ohm coaxial power resistor where it is dissipated as heat. The power flowing into the resistor is monitored by a second power sensor.

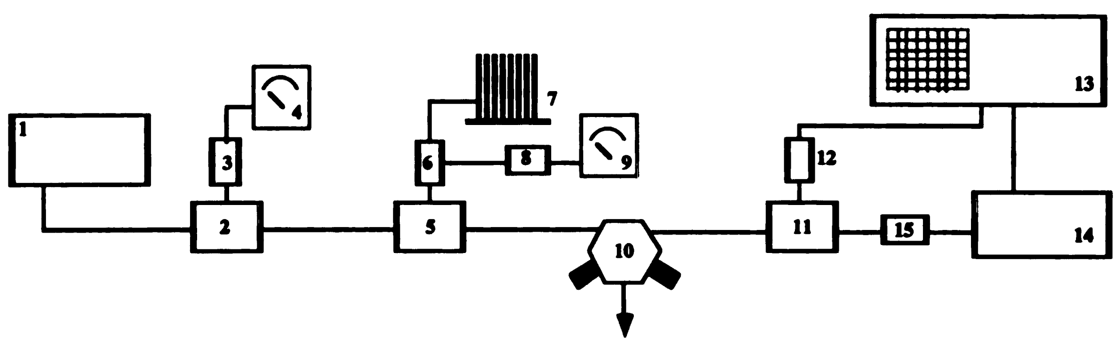

Key:

1. Othos Instruments 100W 2.45 GhZ power source
2. Narda coaxial directional coupler, -20dB/0dB
3. HP power sensor, 3.5W average power
4. HP power meter (incident)
5. Ferrite circulator
6. Narda coaxial directional coupler, -30dB/0dB
7. Termaline coaxial resistor, 50 ohm, 500W
8. HP power sensor, 300mW average power
9. HP power meter (reflected)
10. Transco RF SPDT ransmission switch
11. Western Microwave directional coupler
12. Narda crystal detector .01-18GhZ
13. Tektronix oscilloscope
14. HP sweep oscillator
15. Marconi F6217-47

Figure D.3 - Experimental Microwave Power Delivery Wiring

The right-most portion of figure D.2 contains equipment which is used for cavity tuning operations. This equipment is not used during processing. The equipment includes a sweep oscillator used to generate low-power swept-frequency EM signals in a narrow band around a center frequency of 2.45GHz. This signal is used in conjunction with a reflected signal measurement from the probe to drive an oscilloscope in X-Y mode. When configured correctly, the oscilloscope shows reflected power as a function of frequency and can be used with either a loaded or unloaded cavity to locate modes. When a mode is located, the reflected power is minimized and this is readily observable on the oscilloscope. Additional detail on the experimental setup can be found in [Adegbite 95].

IPCA External Controller Interface Protocol

This section describes the communications interchange protocol between the IPCA prototype and the External Controller. The overall software system can be hosted on a variety of computing platforms on one or more machines connected by means of any Apple™-supported network including LocalTalk™, Ethernet™ and AppleTalk Remote Access™. The information interchange activities between the two systems are facilitated through the use of standard AppleShare™ software and a set of shared files. This approach eliminates the need for platform-specific interfaces to Visualworks\Smalltalk™ and National Instruments LabVIEW™. All that is needed to implement the protocol is the ability to open, read, write and delete ordinary ASCII files.

Example Platforms

The overall system may run on a variety of computing platforms on one or more machines. The only restrictions are that IPCA and the External Controller both have access to the shared files used for information interchange.

Some example platforms are given below:

- **Two Macintosh™ systems connected with a LocalTalk™ cable.**
- **Two Macintosh™ systems connected with Ethernet™.**
- **Two Macintosh™ systems connected with AppleTalk Remote Access™ (using a modem).**
- **One Macintosh™ and one PC connected with Ethernet™.**
- **One 68040-based or higher Macintosh™ with A/UX™ running IPCA and the External Controller concurrently.**

Shared Files

The protocol utilizes a set of files located on a shared disk as a sort of "shared memory" for information interchange between the processes (the Plan Execution Monitor and the External Controller are each considered a process). The files are in common ASCII (plain text) format. A shared RAM disk is used to host the files for increased performance. All files must reside within the same folder on the same shared disk.

Any process attempting to read from or write to one of the shared files should observe certain restrictions to ensure that no process makes it impossible for another process to gain access to the data. In general, processes must not monopolize a shared file by holding it open for reading or writing for more than 1 second. After a successful attempt to read/ write, the process should not immediately attempt to read or write the file again without waiting for a courtesy interval of at least 1 second. Read and write file operations should conclude within 1 second from the time the file is first opened. A process attempting to check on the existence of a file (polling) should not attempt to access the file more frequently than twice per second.

Files "SourceData" and "SourceDataAvailable"

The SourceData file contains sampled (sensed) process data which needs to be sent from the External Controller to the Plan Execution Monitor. The External Controller writes to this file to supply data to the Plan Execution Monitor. When valid SourceData is available for the Plan Execution Monitor, the External Controller writes a zero-length file named SourceDataAvailable as a flag to indicate that the SourceData file is available to the Plan Execution Monitor. The Plan Execution Monitor politely polls (no more often than every 0.5 second) for the existence of the SourceData file, so the External Controller may be unable to open it for writing at all times on the first attempt. Similarly, the Plan Execution Monitor will not always be able to count on opening the SourceDataAvailable file for reading (on the first attempt) because the External Controller may have it locked while trying to create it. Since there is some (minimal) contention for these shared files, both the External Controller and the Plan Execution Monitor must exercise care not to lock each other out for extended periods of time. Thus, both processes should observe the timing constraints given in the section above.

Note: The Plan Execution Monitor deletes the SourceDataAvailable file as soon as the SourceData file is read. The External Controller is then free to write to both files upon detecting that the SourceDataAvailable file no longer exists.

File Format

Each line in the SourceData file consists of a quadruple of values in the following format:

<channel #> <value> <reserved1> <reserved2>

The <channel #> field is an integer in the range [1..128]. If the channel # is outside this range, the remaining data on the line will be ignored. Thus, the protocol allows for up to 128 input channels. The channels can be mapped in any way to the physical sensors, with or without gaps in the sequence. One tuple appears in the file for each sensed channel. The <value> field contains a floating point number which represents the sensed value for the channel. The <reserved1> and <reserved2> fields should be treated as strings. All fields (including the reserved fields) must contain no spaces and are restricted to no more than 31 bytes in length.

One or more spaces are used to separate fields and a standard CR character is used to terminate each line. Each line is restricted to be no more than 128 characters in length. The entire file can contain no more than 128 lines. The maximum size for this file is therefore 16K bytes. Typically, the file will be much smaller.

The SourceData file could be accessed with the following 'C' statements:

```
FILE *inFile;
int channelNumber;
float value;
char res1[32], res2[32];

inFile = fopen("SharedVolumeName:PathToSharedFile:SourceData", "r");
if (inFile != NULL)
fscanf(inFile, "%i %f %s %s", &channelNumber, &value, res1, res2);
```

Files "SinkData" and "SinkDataAvailable"

The SinkData file contains control actions to be sent from the Plan Execution Monitor to the External Controller. The External Controller reads this file and the Plan Execution Monitor writes to it. When the Plan Execution Monitor finishes writing the file, it creates the SinkDataAvailable file to indicate that the SinkData file is available. The External Controller pends on the existence of the SinkDataAvailable file. When the SinkDataAvailable file is found to exist, the External Controller opens the SinkData file, reads its contents, closes the file and immediately deletes the SinkDataAvailable file.

Each line in the SinkData file consists of a 4-tuple as above. The only difference is that the <channel #> field is restricted to be an integer in the range [129..256]. Any line with a <channel #> field value outside this range should be discarded.

Note: The External Controller must delete the SinkDataAvailable file as soon as the SinkData file is read. The Plan Execution Monitor is then free to write to both files upon detecting that the SinkDataAvailable file no longer exists.

Protocol Example

The interchange protocol is adequate for information interchange, but is very simple.

A key feature of the protocol is that the processes are restricted to advance in lock-step. Therefore, if no actions are indicated or no new sensed data is available, <u>a zero-length file must still be written</u> by the External Controller (to SourceData, in step 5 below) or by the Plan Execution Monitor (SinkData, in step 8 below):

1. A shared disk with an empty folder is created to host the four communications files.
2. The Plan Execution Monitor is started up, initializes itself and it waits for the SourceDataAvailable file to appear.
3. The External Controller is started up and it initializes itself.
4. The External Controller senses data from the process (temperature, power, position, etc.).
5. The External Controller writes the data to the SourceData file, closes it and the creates SourceDataAvailable. If no data is available to the External Controller, SourceData simply ends up having a length of zero bytes. The External Controller then pends on the existence of SinkDataAvailable, performing whatever control actions are appropriate between poll attempts. Polling rate is restricted to no more than 2 polls per second (the External Controller should have plenty of other things to do in the several seconds it may take to get SinkDataAvailable).
6. The Plan Execution Monitor (which has been waiting for SourceDataAvailable) sees the SourceDataAvailable file, opens the SourceData file, reads it, closes it and then deletes SourceDataAvailable.
7. The Plan Execution Monitor determines appropriate actions to take given the process-

ing plan/goals and the current sensed information.
8. The Plan Execution Monitor writes appropriate control actions to the SinkData file, closes it and creates SinkDataAvailable. If no actions are indicated, the SinkData file will have a size of zero bytes.
9 The External Controller sees SinkDataAvailable, opens SinkData, reads it, closes it and then deletes SinkDataAvailable. The External Controller takes the actions indicated.
10. Goto 4.

Sample Data Files

Example SourceData and SinkData files are given below. Some indication of the variety of data which could appear in the reserved fields is shown. In practice, SourceData and SinkData files are unlikely to display this much variety of appearance in the reserved fields.

SourceData:

```
1 25.0 res1 res2
2 35.24 ThisIsOKToo 0.0
3 45.2 0 0
4 52.8 A B
14 21 OKToHaveGapsInChannelNumbers !
15 34.2 0 0
```

SinkData:

```
129 10.5 res1 res2
130 15 0.0 0
150 367.21343 0 0
0 0.0 ThisIsIgnoredBecause TheChannelNumberIsOutOfRange
```

APPENDIX E
Domain Background

Composite Materials

Engineers and scientists have sought to enhance, perfect and extend the basic structural materials for centuries. History has recorded entire eras named in recognition of the materials which predominated, including The Stone Age, The Iron Age and The Bronze Age. A continuous flow of structural materials were discovered and produced, each offering unique advantages over previous materials in terms of strength, stiffness, weight, corrosion resistance and other physical properties. Of course, the research and development of new materials is presently more important than ever before. Many applications await the development of high-strength, low-weight materials and there is tremendous promise in the research, application and development of modern advanced composites to serve these increasingly demanding needs.

The meaning of the term composite has evolved over many centuries. The initial common use of the term simply referred to any combination of two or more dissimilar materials with properties superior to either of the starting materials. Bone, wood and reinforced concrete are some examples of composites under this initial definition. Generally, composite materials retain the macroscopic properties and individuality of their constituents.

Metallic alloys, for example, are not considered to be composites because the starting materials co-mingle at the molecular level to produce an isotropic material. Physical properties in an isotropic material are insensitive to the direction or orientation of the material. Composites are generally anisotropic materials (their physical properties differ with respect to their spatial orientation) and are also largely inhomogeneous [Strong 89].

As research and development efforts produced an increasing number and variety of composites, the need for a refinement in terminology became evident and a distinction was made between naturally occurring composite structures (such as bone) and synthetic composite materials (such as reinforced plastics). Generally, when one speaks of composites in a modern context, one refers to synthetic (man-made) composite materials which incorporate a reinforcement in a continuous matrix. The role of the reinforcements in these materials is to accept loads transferred to them by the matrix. The reinforcements give a composite material its strength. The matrix serves a dual role -- load transfer to the reinforcement and protection and isolation of the reinforcement from its surrounding environment. Some composites are arranged as a laminar structure with layers of reinforcements held in a matrix binder.

Included among the advantages of composite materials are high strength, design flexibility, corrosion resistance and excellent physical properties such as tensile modulus and stiffness. The major disadvantages to composite material use are material, design and fabrication costs. The wide variety of applications for composites includes aerospace components and structures, aircraft components, automotive parts, marine and sports equipment, circuit boards and artificial human joints.

Reinforcements

One may characterize composite materials along a number of dimensions such as the type of reinforcement or matrix used or the physical construction of the composite. Three classes of reinforcements are common -- fibers, particulates, and whiskers (which are sometimes alternately described as a subclass of particulates). A partial taxonomy of reinforcements is presented in figure E.1.

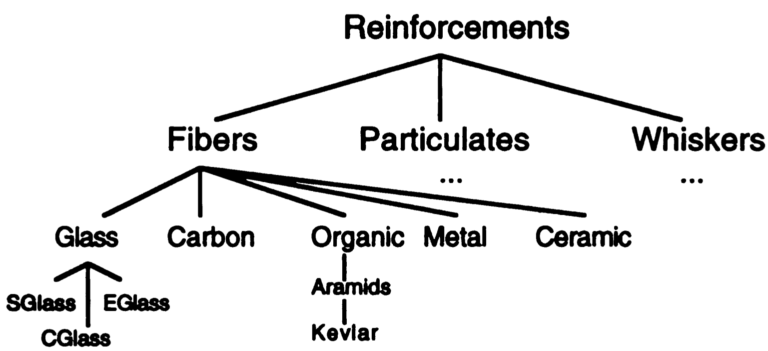

Figure E.1 - Partial Taxonomy of Reinforcement Types

Fibrous reinforcements are the most common of the three reinforcement types and have a single comparatively long axis which is capable of supporting a significant tensile load. Fibers are typically near-circular in cross section and dimensionally large when compared to particulates and whiskers. Fibrous reinforcements are available in wide variety of forms including continuous fibers, various types of knitted fabrics, weaves and braids. Woven or braided materials need not be homogeneous with respect to fiber composition or diameter. Some of the most common fibers are E-Glass, S-Glass, carbon/graphite and aramids. Carbon/graphite fibers have very high strength and tensile modulus, and are commonly used as reinforcements for polymeric composites. The interphase region between fiber and matrix is critical to the properties of the composite. It is in this region that loads are transferred from the matrix to the fibers, and so the bond strength between the fiber and matrix is of critical importance.

Particulate reinforcements have no preferred orientation and their shape is less important than with fibrous reinforcements. The particles may be of a wide variety of shapes, sizes and compositions including ceramic, metal and glass. Examples of particulate composite materials include concrete and particle board.

Whiskers are single crystalline particles with an elongated shape and a preferred orientation. Their physical properties are generally superior to particles but they must be carefully oriented within the composite material matrix, presenting an additional challenge during fabrication.

Matrix Materials

Three major types of matrix materials are common -- polymers, metals and ceramics. A partial taxonomic classification of matrix materials is presented in figure E.2.

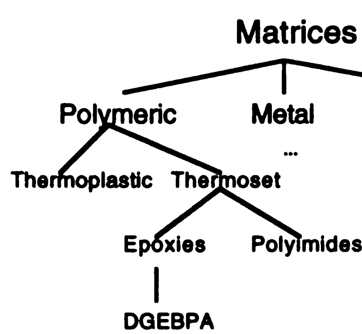

Figure E.2 - Partial Taxonomy of Matrix Types

Fillers such as chalk, clay or glass are often added to matrix materials in an effort to reduce costs. Fillers are distinguished from reinforcements because their primary function is to reduce cost and not to provide load-carrying capabilities within the material. Polymeric matrix materials (resins) are the most common of the matrix types and may be further classified as either thermoset or thermoplastic. There are tens of thousands of patented polymers and polymer systems. Polymer and copolymer structures may be linear, random, blocked, branched or cross-linked [Strong 89].

Thermoset materials are cured under heat and pressure. As a thermoset polymeric composite undergoes the curing process, a non reversible chemical crosslinking reaction occurs as illustrated in figure 3.3, which is adapted from [Strong 89].

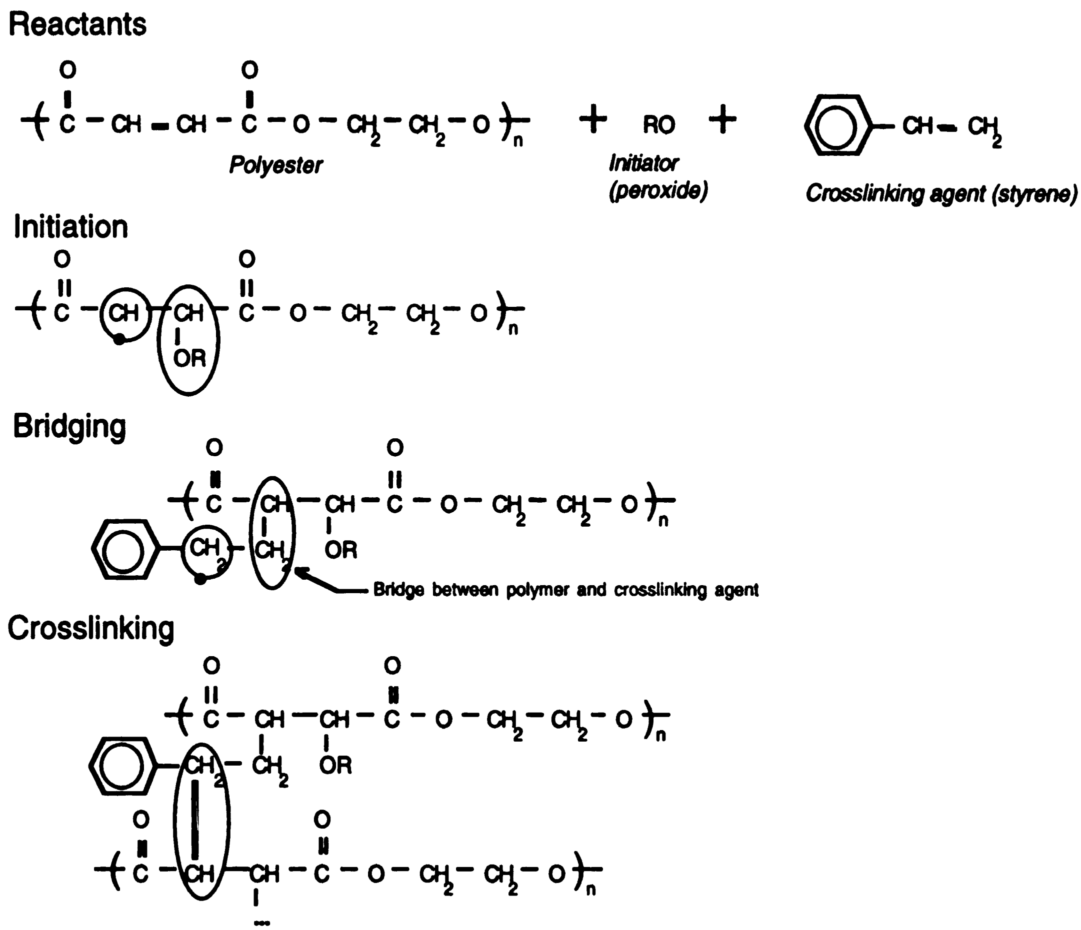

Figure E.3 - Chemical Crosslinking Behavior of Thermoset Composite

During the initial step in the reaction, an initiator breaks the carbon double bond on the polymer. This creates a bonding site for the bridging step. During bridging, the crosslinking agent bonds to this site. In the process, another bonding site is created on the crosslinking agent, which is then free to break another carbon double bond on another polymer chain. The polymers thus become crosslinked after many applications of this procedure. Crosslinking reactions are typically exothermic (they release heat) and the resulting materials cannot be reformed by re-heating. The greater the number of crosslinks formed in the material, the greater the viscosity and, ultimately, the stiffer the material becomes.

For demanding applications, epoxy/amine curing systems offer better performance at higher costs. Amine groups are commonly used as hardeners which encourage greater crosslinking density during the cure by reacting with epoxy groups. Among the common epoxy products is diglycidyl esther of bisphenol-A (DGEBPA).

In contrast, thermoplastic polymeric materials such as polyoxymethylene (POM) and polymethyl methacrylate (PMMA) are processed by heating the solid polymer until a liquid melt phase is reached in which the viscosity of the polymer is relatively low. In this liquid melt phase, the polymer has a low average molecular weight. A reactive cure cycle is not required for thermoplastic materials, which helps to reduce manufacturing complexity and cost. As the long-chained polymers subsequently gel, viscosity increases with average molecular weight of the repeated polymer groups and the average degree of polymerization. The polymer chains generally become quite long as the polymer solidifies, and very large interconnected structures are formed which offer excellent resistance to creep failure. The polymer chains may be physically intertwined, but are not crosslinked to one another chemically through a crosslinking agent (as is the case with thermoset materials). Molecular weights range from 5,000 to 100,000 for typical polymers. Catalysts and initiators may be used to encourage mass polymerization. Some conventional thermoplastic polymers are nylon, polycarbonate, polyethylene, polystyrene, polyester and acrylics. Thermoplastic materials may be re-heated and cooled repeatedly, offering the potential for convenient recycling. Fiberglass reinforcements are commonly used with thermoplastic materials in low-cost applications. The structure of a typical high-performance thermoplastic (PEEK) is shown in figure E.4.

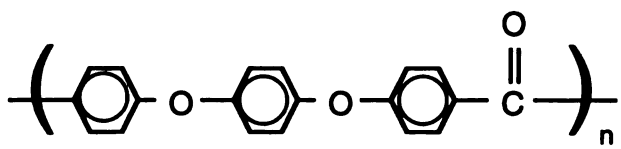

PolyEtherEtherKetone (PEEK)

Figure E.4 - Chemical Structure of High-Performance Thermoplastic

Thermoplastic materials offer advantages in their shorter manufacturing times, lower equipment investment costs and higher impact and creep resistance, but have some inferior structural properties when compared with thermoset materials. However, advances have been made with thermoplastics, and many high-temperature and high-strength advanced thermoplastics are becoming available to help narrow this traditional performance gap. Thermoset materials offer generally excellent properties in terms of strength and stiffness, and may have higher temperature capabilities. Polyamides are used in very high-temperature polymeric applications up to temperatures of approximately 370° C.

Metal matrices such as aluminum, titanium, magnesium or copper are generally used for very high-strength, medium-temperature applications. Metal matrix composites offer greater thermal stability and superior physical properties over those observed in pure metals [Harris 86]. Particles and whiskers are commonly used as reinforcements due to difficulties in wetting-out fibrous reinforcements during fabrication. Metal matrix composites are often isotropic materials like their metal counterparts.

Ceramic matrices are used primarily for very high-temperature applications. Generally, these are crystalline, non-organic, ionic-bonded, ordered structures which are capable of withstanding very high temperatures (1100-2200° C). Ceramic matrix composites offer high dimensional stability and high resistance to environmental attack, but are very brittle and are prone to catastrophic failure modes (especially cracking). They are not generally used as structural parts for these reasons.

Layup and Processing Technologies

A number of methods exist for the fabrication of composite materials, including manual lay-up, prepregging, filament winding, pultrusion, matched-die molding and resin-transfer molding [Srivatsan 91].

Manual lay-up involves layering mats within a mold by hand. One simple method is to layer the mat and cover with resin. This method obviously has the potential for poor wetting of the reinforcement (and therefore poor bonding) but offers a simple, low-cost method for non-critical applications. Mats are also available which are pre-impregnated with resin. These materials are called prepregs and they generally achieve superior wet-out of the reinforcing fiber, usually with the help of a technique called vacuum-bagging. Manual lay-up methods have several problem areas including lack of uniformity, voids, and shrinkage due to uneven resin application. Only one side of the material is finished (smooth) if matched dies are not utilized. Accordingly, male molds are used to yield smooth inner surfaces (such as a hot tub) while female molds result in finished outer surfaces (such as a boat hull).

Filament winding is a technique in which a continuous tape or fiber is wound on a mandrel (a rotating mold mounted on a lathe) to form the resulting part. Dry filament winding uses a prepreg tow as its winding medium, while in wet filament winding the continuous fiber is dipped (impregnated) in a resin bath just prior to winding. With wet winding, the viscosity of the resin is quite important -- dripping may occur if the resin becomes too fluid, while poor wetting may result if the resin is too viscous.

Pultrusion is a continuous process involving impregnation, shaping and curing steps. Pultrusion offers high throughput and minimal material waste, but part cross-sectional area must be largely constant. Automated equipment for pulling, winding and cutting operations is typically employed.

Matched-die molding may utilize sheet molding compound (SMC) or bulk molding compound (BMC) together with male/female mold pairs. Material is placed in one mold half and pressed into the second mold half. Pressure and heat are then maintained to cure the material, which displays two finished surfaces.

Resin-transfer molding (RTM) may be used for greater manufacturing productivity. In this technique, a mold is pre-loaded with the reinforcement and closed up. Resin is then injected under pressure (actually, a vacuum is usually drawn within the mold to assist in creating a pressure differential) and the part is cured. RTM has an advantage over manual lay-up in terms of throughput, and the process may be automated more easily as well. There are several critical factors, however, including design of the mold, viscosity of the resin and potential movement of the reinforcement during injection, which can make RTM impractical in many applications.

Layup

A number of methods exist for the fabrication of composite materials, including manual lay-up, prepregging, filament winding, pultrusion, matched-die molding and resin-transfer molding [Srivatsan 91].

Manual lay-up involves layering mats within a mold by hand. One simple method is to layer the mat and cover with resin. This method obviously has the potential for poor wetting of the reinforcement (and therefore poor bonding) but offers a simple, low-cost method for non-critical applications. Mats are also available which are pre-impregnated with resin. These materials are called prepregs and they generally achieve superior wet-out of the reinforcing fiber, usually with the help of a technique called vacuum-bagging. Manual lay-up methods have several problem areas including lack of uniformity, voids, and shrinkage due to uneven resin application. Only one side of the material is finished (smooth) if matched dies are not utilized. Accordingly, male molds are used to yield smooth inner surfaces (such as a hot tub) while female molds result in finished outer surfaces (such as a boat hull).

Filament winding is a technique in which a continuous tape or fiber is wound on a mandrel (a rotating mold mounted on a lathe) to form the resulting part. Dry filament winding uses a prepreg tow as its winding medium, while in wet filament winding the continuous fiber is dipped (impregnated) in a resin bath just prior to winding. With wet winding, the viscosity of the resin is quite important -- dripping may occur if the resin becomes too fluid, while poor wetting may result if the resin is too viscous.

Pultrusion is a continuous process involving impregnation, shaping and curing steps. Pultrusion offers high throughput and minimal material waste, but part cross-sectional area must be largely constant. Automated equipment for pulling, winding and cutting operations is typically employed.

Matched-die molding may utilize sheet molding compound (SMC) or bulk molding compound (BMC) together with male/female mold pairs. Material is placed in one mold half and pressed into the second mold half. Pressure and heat are then maintained to cure the material, which displays two finished surfaces.

Resin-transfer molding (RTM) may be used for greater manufacturing productivity. In this technique, a mold is pre-loaded with the reinforcement and closed up. Resin is then injected under pressure (actually, a vacuum is usually drawn within the mold to assist in creating a pressure differential) and the part is cured. RTM has an advantage over manual lay-up in terms of throughput, and the process may be automated more easily as well. There are several critical factors, however, including design of the mold, viscosity of the resin and potential movement of the reinforcement during injection, which can make RTM impractical in many applications.

APPENDIX F
Extending the Capabilities of IPCA: A Case Study

Extending the IPCA Framework

IPCA is designed to allow straightforward incorporation of new planning and control method fragments into the existing problem-solving framework. It is possible to extend the available set of IPCA control algorithms and planning approaches, to provide coverage for additional fabrication technologies, new materials and so on. The IPCA architecture serves as an extensible framework for representing knowledge about composite materials processing.

Each extension potentially impacts IPCA in two areas:

Knowledge about the applicability of the extension is placed into an appropriate agent within the DSPL-based Blueprint Generator

Procedural algorithms covering new control and planning methods are added to the fragment library

Details regarding these extensions are reviewed in the sections below. Some common types of extensions are shown in table F.1.

Type of Extension	Impact on Blueprint Generator	Impact on Fragment Library	Notes
Add new control method	Add knowledge about method applicability	Add procedural body of method to fragment library	Method name must be unique in library hierarchy
Add new planning method	Add knowledge of when method is useful	Add procedural body of method to fragment library	Method name must be unique in library hierarchy
Add new composite material	Add knowledge of material properties	None required	Processing temps, times, characteristics of material
Add new fabrication technology	Add knowledge about when technology is indicated; add DSPL agents to design blueprints	Add any new fragments required by the technology	Design DSPL structure and add to existing IPCA framework

Table F.1 - Common IPCA Extensions

The final section in this appendix details the procedure for adding a new type of mode-switching control method fragment to IPCA.

Blueprint Generator Extensions

Extensions to the DSPL framework of the Blueprint Generator are made by inserting appropriate design agents into the existing IPCA DSPL hierarchy. This task is accomplished using the IPCA Blueprint Generator interface (see appendix B) to browse the agent hierarchy and modify it. To accommodate a new fabrication technology, for example, it is necessary to insert a new specialist for the technology in an appropriate position under the fabrication technology specialist. Similarly, to add coverage for new matrix materials, the design agents must be modified to provide for setting appropriate processing temperatures for the materials involved.

Whenever new planing or control method fragments are added to the fragment library, knowledge about the applicability of the fragments must be included in the constraints and steps within the existing IPCA DSPL design hierarchy so the new fragments are specified for use under the proper conditions.

Fragment Library Extensions and Changes

Extensions to the IPCA fragment library are made by using the Fragment Librarian tool (see appendix B). The Fragment Librarian is capable of adding, deleting and renaming plan and control method fragments. Using the command key and selecting an item in the *Library*, *Fragment* or *Element* lists causes a pop-up menu to appear which includes options to add, delete and rename items in the lists. The lists are hierarchical, meaning that each *Library* entry can have many *Fragment* entries, and each *Fragment* entry can have many *Element* entries. Each *Element*, in turn, contains text which can be edited.

During plan assembly, items in the *Library* and *Fragment* lists are matched against the key/value pairs in the blueprint being assembled. For example, if the pair *(ModeControlPolicy ModeControlPolicyC)* is included in the blueprint, this means that the Blueprint Generator previously set the value of the DSPL variable *ModeControlPolicy* to the value *ModeControlPolicyC*. The Plan Assembler will attempt to match the key *ModeControlPolicy* against entries in the *Library* list. If a match is found, the assembler next attempts to match the value *ModeControlPolicyC* against items in the *Fragment* list. If a match is found, the *Elements* from the matched fragment are included in the process control plan.

Each fragment contains the following elements:

Parameters-Abstract: A list of the abstract parameters used by the fragment.
Example: (PartTemp PartGrad ElapsedProcessingTime).

Parameters-Input: A list of the input parameters used by the fragment.
Example: (T1Sensor T2Sensor T3Sensor T4Sensor).

Parameters-Output: A list of the output parameters used by the fragment.
Example: (ModeActuator1 PowerActuator1).

Parameters-Source: The SmallTalk source code for the executable portion of any abstract parameters required by the fragment. This code is evaluated at each execution cycle for all abstract parameters. See fragment library for examples of format.

Parameters-Variable: A list of the variable parameters used by the fragment.
Example: (ProcessingGoal InitializationComplete).

Plan-CoFragments: A list of fragments which need to be included with this fragment. Usually this list is empty, but it can contain references to other dependant fragments.
Example: (ModeControlPolicy@ModeControlPolicyC)

Plan-Source: The SmallTalk source code for the executable portion of the fragment. See fragment library for examples of format.

Case Study - Adding a new Fragment Type to IPCA

This section contains a case study detailing how a new control method fragment was added to IPCA. The fragment supports a new mode-switching behavior which attempts to equalize the average temperature of all probes. The new fragment was used for experimental series F4.

Previous mode switching methods (used in experimental series F1 through F3) attempted to minimize thermal gradients by selecting processing modes that would instantantanously minimize the temperature differential between the probes. This approach resulted in the observation that some probes were consistently warm while others were consistently cool. This situation is undesirable during fabrication due to the effects of incomplete chemical crosslinking in some of the cooler areas of the sample.

The new fragment overcomes these problems by switching between modes using a different mode switching strategy. The new strategy attempts to achieve and maintain equal average temperatures for all probes by basing mode switches on historical temperature data for each probe. For additional detail, see chapter 6.

The following steps were performed to design the new fragment and integrate it into IPCA:

1. The algorithm for the new mode switching fragment was designed and coded. The input, output, variable and abstract parameters used by the new fragment were included in the element lists using the fragment librarian.
2. The SetModeControlPolicy step under the MDProcessStructure specialist in the Blueprint Generator DSPL hierarchy was modified to recommend the new fragment as the standard mode switching policy for all reactive mode switching applications.
3. A process control plan was generated and tested against the simulated control unit to verify proper behavior and debug the new fragment.
4. The process control plan was tested in the laboratory in experimental series F4.

BIBLIOGRAPHY

F.L. Abrams, T.L. Lagnese, S.R. LeClair, C.W. Lee, P.H. Garrett, and J.B. Park, Qualitative Process Automation for Autoclave Curing of Composites, Air Force Wright Aeronautical Laboratories Technical Report 97-4083.

F.L. Abrams, An Expert System Process Controller for Advanced Composites, Proceedings; Detroit Re Tec '87 (Society of Plastics Engineers, Detroit, MI).

L. Acar, and U. Ozguner, Design of Hierarchically Distributed Expert Controllers for Large-Scale Systems, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 12-17.

V.O. Adegbite, Dissertation for the degree of Doctor of Philosophy, Michigan State University, East Lansing, MI, 1995.

J.S. Albus, A Theory of Intelligent Systems, Fifth IEEE International Symposium on Intelligent Control (Philadelphia, Sept. 5-7, 1990) pp. 866-875.

J.S. Albus, Lumia, R., and McCain, H, Hierarchical Control of Intelligent Machines Applied to Space Station Telerobots, 1987 IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, January 19-20 1987), pp. 20-26.

R. Alterman, An Adaptive Planner, Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI '86), pp. 65-69.

J.A. Ambros-Ingerson, and S. Steel, Integrating Planning, Execution and Monitoring, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI '88), pp. 83-88.

P.J. Antsaklis, K.M. Passino, and S.J. Wang, Autonomous Control Systems: Architecture and Fundamental Issues, Proceedings of the 1988 American Control Conference, (Atlanta, Georgia, June 15-17, 1988), American Automatic Control Council, pp. 602-607.

K.E. Arzen, Knowledge-Based Control Systems - Aspects on the Unification of Conventional Control Systems and Knowledge-Based Systems, Proceedings of the 1989 American Control Conference, (Pittsburgh, Pennsylvania, June 21-23, 1989), American Automatic Control Council, pp. 2233-2238.

J. Asmussen, et. al., Single-Mode or Controlled Multimode Microwave Cavity Applicators for Precision Material Processing, Review of Scientific Instruments, Vol. 58, No. 8, 1987, pp. 1477-1486.

K.J. Astrom, B. Wittenmark, Computer-Controlled Systems, Prentice-Hall, Englewood Cliffs, New Jersey, Second edition, 1990.

M. Aynsley, D. Peel, and A.J. Morris, A Real-Time Knowledge-Based System for Fermentation Control, Proceedings of the 1989 American Control Conference, (Pittsburgh, Pennsylvania, June 21-23, 1989), American Automatic Control Council, pp. 2239-2244.

V.V. Badami, P. Nielson, and J.B. Comly, An Intelligent Controller for Process Automation, Proceedings of the Fifth IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, September 5-7 1990), pp. 1076-1081.

R. Barletta, and D. Hennessy, Case Adaptation in Autoclave Layout Design, Lockheed AI Center O/96-20 B/259.

M.R. Basila Jr., and A. Cinar, A Model-Object Based Supervisory Expert System for Fault Tolerant Chemical Reactor Control, Proceedings of the 1989 American Control Conference, (Pittsburgh, Pennsylvania, June 21-23, 1989), American Automatic Control Council, pp. 1348-1353.

T.E. Bihari, T.M. Walliser, and M.R. Patterson, Controlling the Adaptive Suspension Vehicle, IEEE Computer, Vol. 22, No. 6 (June 1989), pp. 59-65.

M. Boddy, and T. Dean, Solving Time-Dependent Planning Problems, Proceedings of the 11th IJCAI (1989), pp. 979-984

S.P. Boyd and C.H. Barratt, Linear Controller Design - Limits of Performance, Prentice Hall, Englewood Cliffs, New Jersey, 1991.

D.C. Brillhart, and S.F. Wible, An Expert System for Real-Time Process Characterization and Control, Proceedings of the Seventh IEEE/CHMT International Electronic Manufacturing Technology Symposium (San Francisco, California, September 25-27, 1989), pp. 76-83.

W.L. Brogan, Modern Control Theory, Prentice Hall, Englewood Cliffs, New Jersey, Third edition, 1991.

R.A. Brooks, A Robust Layered Control System for a Mobile Robot, IEEE Journal of Robotics and Automation, Vol. 2, No. 1, March 1986, pp. 14-23.

C.A. Broverman, and W.B. Croft, Reasoning About Exceptions During Plan Execution Monitoring, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI- '88), pp. 190-195.

D.C. Brown and B. Chandrasekaran, Design Problem Solving - Knowledge Structures and Control Strategies, Morgan Kaufmann, San Mateo, California, 1989.

F.C. Campbell, A.R. Mallow, F.R. Muncaster, E.L. Boman, and G.A. Blase, Computer-Aided Curing of Composites, Air Force Wright Aeronautical Laboratories Technical Report 86-4060.

S. Chand, On-line Self-Monitoring Tuner for Proportional Integral Derivative Controllers, 1991 IEEE Decision and Control, Vol. 2, pp. 1905-1906.

B. Chandrasekaran, R. Bhatnagar, and D.D. Sharma, Real-Time Disturbance Control, Communications of the ACM, August 1991, pp.32-47.

B. Chandrasekaran, Design Problem Solving: A Task Analysis, AI Magazine, Winter, 1990, pp. 59-71.

B. Chandrasekaran, et. al., An Approach to Routine Planning, International Journal of Man-Machine Studies, 1986.

B. Chandrasekaran, Generic Tasks in Knowledge-Based Reasoning: High-Level Building Blocks for Expert System Design, IEEE Expert, Vol. 1, No. 3, Fall 1986, pp. 23-30.

B. Chandrasekaran, J. Josephson, A. Keuneke, and D. Herman, Building Routine-Planning Systems and Explaining their Behavior, International Journal of Man-Machine Studies, Vol. 30, 1986, pp. 377-398.

D. Chapman, Planning for Conjunctive Goals, Artificial Intelligence, Vol. 32, 1987, pp. 333-377.

D. Chapman, Penguins Can Make Cake, AI Magazine, Winter 1989, pp. 45-50.
J.C. Chautard, and C. Honnorat, PEX: A Reactive Procedure Based Decision Maker, 1991 IEEE Conference on AI Applications, pp. 368-371.

M. Cibrario, and J. Levine, Saddle-Node Bifurcation Control with Application to Thermal Runaway of Continuous Stirred Tank Reactors, 1991 IEEE Decision and Control, Vol. 2, pp. 1551-1552.

T.B. Co, J.G. Williams, and Z. Chen, Control Studies on the Advanced Composites Manufacturing Process, Composites Manufacturing, December 1990, pp. 253-259.

N. Coleman, An Emulation/Simulation for Intelligent Controls, Proceedings of the Workshop on Software Tools for Distributed Intelligent Control Systems, (July 17-19, 1990, Pacifica, California), pp. 37-52.

R.D. Coyne, M.A. Roseman, A.D. Radford, M. Balachandran and J.S. Gero, Knowledge-Based Design Systems, Addison-Wesley, Reading, Massachusetts, 1990.

J.R. Crosscope, and R.D. Bonnell, An Integrated Intelligent Controller Employing both Conceptual and Procedural Knowledge, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 416-422.

J.B. Cruz, and A.R. Stubberud, Knowledge-Based Approach to Multiple Control Coordination in Complex Systems, Proceedings of the 1987 IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, January 19-20, 1987), IEEE Computer Society Press, pp. 50-53.

B. D'Ambrosio, M.R. Fehling, P. Raulefs, and B.M. Wilber, Real-Time Process Management for Materials Composition in Chemical Manufacturing, IEEE Expert, Vol. 2, No. 2, Summer 1987, pp. 80-93.

B. Dasarathy, Timing Constraints of Real-Time Systems: Constructs for Expressing them, Methods for Validating Them, IEEE Transactions on Software Engineering, Vol. 11, No. 1, January 1985, pp. 80-86.

T. Dean, and M. Boddy, An Analysis of Time-Dependent Planning, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI '88), pp. 49-54.

T. Dean, and M. Boddy, Incremental Causal Reasoning, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI '88), pp. 196-201.

T. Dean, and G. Siegle, An Approach to Reasoning About Continuous Change for Applications in Planning, Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI '90), pp. 132-137.

D. Decker, J. Sticklen, et al., An Adaptive Intelligent Control Architecture for the Microwave Fabrication of Polymeric Thermoset Composite Materials Utilizing an Empirical Single-Mode Resonant Cavity Scanning Technique, Proceedings of the AIChE 1993 Summer National Meeting, (Seattle, Washington, August, 1993).

D. Decker, J. McDowell, M. Hawley, W. Punch, J. Sticklen, An Intelligent Control Architecture for Processing of Polymer Composites, Proceedings of the 41st Annual Sagamore Army Materials Conference, (Plymouth, Massachusetts, August, 1993

D. Decker, W. Punch, J. McDowell, J. Sticklen, An Intelligent Control Architecture for the Microwave Fabrication of Composite Materials Based on the Generic Task Approach to Knowledge-Based Systems, Proceedings of the AAAI Fall Symposium on Intelligent Systems, (New Orleans, Louisiana, November 3-6 1994).

R. Doraiswami, and W. Liu, An Expert System for Monitoring the Status of Control Systems, Proceedings of the Fifth IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, September 5-7 1990), pp. 457-462.

M. Drummond, and J. Bresina, Planning for Control, Proceedings of the Fifth IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, September 5-7 1990), pp. 657-662.

M. Dudziak, M. Hall, and D. Zaret, IVC: An Intelligent Vehicle Controller with Real-Time Strategic Replanning, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 145-152.

E.H. Durfee, and V.R. Lesser, Incremental Planning to Control a Blackboard-Based Problem Solver, Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI '86), pp. 58-64.

E.H. Durfee, A Cooperative Approach to Planning for Real-Time Control, Proceedings of DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control (1990), pp. 277-283. Morgan Kauffman.

C. Elkan, Incremental Approximate Planning, Proceedings of the Eighth National Conference on Artificial Intelligence, (AAAI '90), MIT Press, Cambridge, Mass, 1990, pp. 145-150.

L.D. Erman, Ed. Intelligent Real-Time Problem Solving, Air Force Office of Scientific Research, Workshop Rep. TTR-ISE-90-101, January 1990.

A. Farsaie, T.R. McKnight, B. Ferren, and C.F. Harrison, Intelligent Controllers for an Autonomous System, Proceedings of the 1987 IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, January 19-20, 1987), IEEE Computer Society Press, pp. 154-157.

L.A. Fellows, Microwave Processing of Polyamides and Polyamide Composites, Thesis (M.S.), Michigan State University, East Lansing, 1991.

J.A. Ferguson, Touring Machines: Autonomous Agents with Attitudes, IEEE Computer, May 1992, pp. 51-55.

R.E. Fikes, and N.J. Nilsson, STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, Artificial Intelligence, Vol. 2, 1971, pp. 189-208.
J. Fiksel, and F. Hayes-Roth, Knowledge Systems for Planning Support, IEEE Expert, Fall 1991, pp.16-23.

R.J. Finocchio, A.J. Blauch, T.S. Lorincz, R. Meth, and B.P. Butz, Implementing a Real-Time Knowledge-Based Controller In a Velocity Controlled, Servo Motor Control System, Proceedings of the Fifth IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, September 5-7 1990), pp. 811-815.

G. Fiorio, F. Lombardi, and A. Villa, Integration of Multi-Hierarchy Control Architectures for Complex Systems, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 71-79.

R.J. Firby, An Investigation into Reactive Planning in Complex Domains, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI -'88), pp. 202-206.

K.D. Forbus, Qualitative Process Theory, Artificial Intelligence, Vol. 24, pp. 85-168.

T.H. Fraichard, and C. Laugier, On-line Reactive Planning for a Non Holonomic Mobile in a Dynamic World, Proceedings of the 1991 IEEE International Conference on Robotics and Automation (Sacramento, California, April 1991), pp. 432-437.

P. Garrett, C.W. Lee, and S.R. LeClair, Qualitative Process Automation vs. Quantitative Process Control, Proceedings of the American Control Conference, Minneapolis, MN, American Automatic Control Council.

Z. Geng, R. Carroll, M. Jamshidi, and R. Kisner, An Adaptive Learning Control Approach, 1991 IEEE Decision and Control, Vol. 2, pp. 1221-1222.

M.P. Georgeff, and F.F. Ingrand, Decision Making in an Embedded Reasoning System, Proceedings of the International Joint Conference on Artificial Intelligence, Morgan Kaufmann, San Mateo, California, 1989, pp. 972-978.

M.P. Georgeff, and A.L. Lansky, Reactive Reasoning and Planning, Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI '87), Seattle, Washington, July 1987, pp. 677-682.

M.L. Ginsberg, Universal Planning: An (Almost) Universally Bad Idea, AI Magazine, Winter 1989, pp. 40-44.

M.L. Ginsberg, Possible Worlds and the Qualification Problem, Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI '87), pp. 212-217.

P. Godefroid, and F. Kabanza, An Efficient Reactive Planner for Synthesizing Reactive Plans, Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI '91), pp. 640-645.

M. Gopal, Modern Control System Theory, John Wiley & Sons, New York, Second edition, 1993.

D.A. Handelman, and R.F. Stengel, Rule-Based Mechanisms of Learning for Intelligent Adaptive Flight Control, Proceedings of the 1988 American Control Conference, (Atlanta, Georgia, June 15-17, 1988), American Automatic Control Council, pp. 208-213.

B. Harris, Engineering Composite Materials, The Institute of Metals, 1986.

M.C. Hawley, and J. Wei, Microwave Processing of Polymer Composites, MRS Spring Meeting, San Francisco, CA, April 1994.

C. Hayes, Using Goal Interactions to Guide Planning, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI '88), pp. 224-228.

B. Hayes-Roth, Architectural Foundations for Real-Time Performance in Intelligent Agents, Journal of Real-time Systems, Vol. 2 (1990) pp. 99-125.

R.F. Herrington, Time-Harmonic Electromagnetic Fields, McGraw-Hill, New York, 1961.

J.P.E. Hodgson, Structures for Intelligent Systems, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 348-353.

L.E. Holloway, C.J. Paul, J.K. Strosnider, and B.H. Krough, Integration of Behavioral Fault-Detection Models and an Intelligent Reactive Scheduler, Proceedings of the Sixth IEEE International Symposium on Intelligent Control (Arlington, Virginia, August 13-15 1991), pp. 134-139.

F. Jahanian, and A. Mok, Safety Analysis of Timing Properties in Real-Time Systems, IEEE Transactions on Software Engineering, Vol. 12, No. 9, September 1986, pp. 890-904.

M. Jamshidi, S. Baugh, D. Barak, and N. Vadiee, A Comparison of an Expert and an Adaptive Fuzzy Control Approach, Proceedings of the 30th IEEE Conference on Decision and Control (Brighton, England, Dec. 1991), Vol. 2, pp. 1907-1908.

J. Jow, M.C. Hawley, and M. Finzel, Microwave Heating and Dielectric Diagnosis Technique in a Single-Mode Resonant Cavity, Review of Scientific Instruments, Vol. 60, No. 1, January 1989, pp. 96-103.

J. Jow, M.C. Hawley, M. Finzel and T. Kern, Dielectric Analysis of Epoxy/Amine Resins using Microwave Cavity Technique, Polymer Engineering and Science, Vol. 28, No. 22, November 1988, pp. 1450-1454.

L.P. Kaelbling, An Architecture for Intelligent Reactive Systems, Proceedings of the 1986 Workshop on Reasoning About Actions and Plans, M.P. Georgeff and A.L. Lansky, eds., 1987, Morgan Kaufmann Publishers, Los Altos, California, pp. 395-410.

S. Kambhampati, A Theory of Plan Modification, Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI '90), pp. 176-182.

A. Kamel, Dissertation for the degree of Doctor of Philosophy, Michigan State University, East Lansing, MI, 1994.

G. Karsai, C. Biegl, S. Padalkar, J. Sztipanovits, and K. Kawamura, Knowledge-Based Approach to Real-Time Supervisory Control, Proceedings of the 1988 American Control Conference, (Atlanta, Georgia, June 15-17, 1988), American Automatic Control Council, pp. 620-625.

H.A. Kautz, and J.F. Allen, Generalized Plan Recognition, Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI '86), pp. 32-37.

H. Kern, and M. Fathi-Torbaghan, Concept for Combining features of Real-time and Expert systems for On-Line Plant Diagnosis, Ninth Annual International Phoenix Conference on Computers and Communications (Scottsdale, Arizona, March 21-23, 1990), pp. 803-807.

C.A. Knoblock, J.D. Tenenberg, and Q. Yang, Characterizing Abstraction Hierarchies for Planning, Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI '91), pp. 692-697.

W. Kohn, Declarative Control Architecture, Communications of the ACM, August 1991, pp. 64-79.

W. Kohn, Declarative Hierarchical Controllers, Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control, (Nov. 1990).

W. Kohn, Advanced Architectures and Methods for Knowledge-Based Planning and Declarative Control, IR and D BSC-021, To be presented as a paper in ISMIS'91, October 1991.

W. Kohn, and T. Skillmann, Hierarchical Control Systems for Autonomous Space Robots, Proceedings of AIAA Conference on Guidance, Navigation and Control, Vol. 1, Minneapolis, MN, Aug. 15-18 1988, pp. 382-390.

T.J. Laffey, P.A. Cox, J.L. Schmidt, S.A. Kao, and J.Y. Read, Real-Time Knowledge-Based Systems, AI Magazine, Vol. 9, No. 1, Spring 1988, pp. 27-45.

T. Laffey, S. Weitzenkamp, J. Read, S. Kao, and J. Schmidt, Intelligent Real-Time Monitoring, Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI - '88, August 1988, Lockheed Artificial Intelligence Center), pp. 72-76.

A.L. Lansky, and D.S. Fogelsong, Localized Representation and Planning Methods for Parallel Domains, Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI '87), pp. 240-245.

A.A. Lazar, J.T. Amenyo, and S. Mazumdar, WIENER: A Distributed Expert System for Dynamic Resource Allocation in Integrated Networks, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 159-164.

S.R. LeClair, F.L. Abrams, and R.F. Matejka, Qualitative Process Automation: Self-Directed Manufacture of Composite Materials, AI EDAM, Vol. 3, No. 2, 1989, pp. 125-136.

S.R. LeClair, and F.L. Abrams, Qualitative Process Automation, International Journal of Computer Integrated Manufacturing, Vol. 2, No. 4, pp. 205-211.

S.R. LeClair, Sensor Fusion: The Application of Artificial Intelligence to Process Control, 1986 Rochester FORTH Conference Proceedings, Rochester, NY, pp. 15-22.

S.R. LeClair, and J.B. Park, Sensor Fusion, Statistical Process Control, edited by Keats and Heuble (Marcel Dekker, New York), NY, pp. 235-267.

C.W. Lee, Composite Cure Process Control by Expert Systems, Department of Chemical Engineering, University of Dayton.

S. Lee, and Y.G. Shin, Multiple Agent Cooperative Problem Solving with Axiom-Based Negotiation, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 469-474.

T.H. Lee, K.W. Lim, and W.C. Lai, Real-Time Multivariable Self-Tuning Controller Using a Feedforward Paradigm with Application to a Coupled Electric-Drive Pilot Plant, IEEE Transactions on Industrial Electronics, August 1991, pp. 237-242.

W.I. Lee, and G.S. Springer, Interaction of Electromagnetic Radiation with Organic Matrix Composites, Journal of Composite Materials, Vol. 18, July 1984.

W.I. Lee, and G.S. Springer, Microwave Curing of Composites, Journal of Composite Materials, Vol. 18, July 1984.

V. Lesser, J. Pavlin, and E. Durfee, Approximate Processing in Real-Time Problem Solving, AI Magazine, Vol. 9, No. 1, Spring 1988, pp. 49-61.

T. Lenz, J.K McDowell, B. Moy, J. Sticklen and M. Hawley, Intelligent Decision Support for Composite Material Design in an Integrated Design Environment, American Society of Composites (ASC) 9th Technical Conference on Composite Materials, University of Delaware, 1994.

T.A. Linden, and J. Glicksman, Contingency Planning for an Autonomous Land Vehicle, Proceedings of the 1987 International Joint Conference on AI, pp. 1047-1054.

A.C. Loos, and G.S. Springer, Curing of Epoxy Matrix Composites, Journal of Composite Materials, Vol. 17, March 1983, pp. 135-169.

D. McDermott, Robot Planning, AI Magazine, Summer 1992, pp. 55-79.

J.K. McDowell, A. Kamel, J. Sticklen and M.C. Hawley, Integrating Material/Part/Process Design for Polymer Composites: A Knowledge-Based Problem Solving Approach, In D. G. Newaz (Ed.), ASC 8th Technical Conference on Composite Materials, Cleveland, OH, American Society of Composites, 1993.

A. Meystel, Planning/Control Architectures for Master Dependent Autonomous Systems with Nonhomogeneous Knowledge Representation, 1987 IEEE International Symposium on Intelligent Control, (January 19-20, 1987, Philadelphia, Pennsylvania), pp. 31-47.

D.J. Minnick, H. Kaufman, and G.W. Neat, Expert Hierarchical Adaptive Controller for Robotic Systems, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 8-11.

J.S.B. Mitchell, D.W. Payton, and D.M. Keirsey, Planning and Reasoning for Autonomous Vehicle Control, International Journal of Intelligent Systems, Vol. 2 (1987), pp. 129-198.

D. Moldovan, W.L. Lee, L. Changhwa, and M. Chung, SNAP Parallel Processing Applied to AI, IEEE Computer, May 1992, pp. 39-49.

F. Morant, P. Albertos, M. Martinez, and J.J. Serrano, Hierarchical Expert System as Supervisory Level in an Adaptive Controller, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 18-25.

B.J. Moy, J.K. McDowell, T. Lenz, J. Sticklen and M.C. Hawley, Expansion of an Intelligent Decision Support System for Process Selection and Process Design in Polymer Composites, Abstract submitted to the 26th International SAMPE Technical Conference, Feb. 1994.

D.J. Musliner, E.H. Durfee, and K.G. Shin, Execution Monitoring and Recovery Planning with Time, 1991 IEEE Conference on AI Applications, pp. 385-388.

D.J. Musliner, J.A. Hendler, A.K. Agrawala, et. al, The Challenges of Real-Time AI, IEEE Computer, January, 1995, pp. 58-66.

S.R. Nann, A. Ray, and S. Kumara, A Decision Support System for Real-Time Control and Monitoring of Dynamical Processes, Proceedings of the 1989 American Control Conference, (Pittsburgh, Pennsylvania, June 21-23, 1989), American Automatic Control Council, pp. 361-365.

National Instruments Corporation, LabVIEW 2 User Manual, September 1991 Edition, Part number 320244-01, 6504 Bridge Point Parkway, Austin, TX.

A. Newell, The Knowledge Level, Artificial Intelligence 18 (1982), pp. 87-127.

A. Nitsche, H. Kern, and J. Janczak, Composites Design Based on Expert Knowledge, Fifth Technical Conference of the American Society for Composites, Technomic Publishing Company, Lancaster, PA, 1990, pp. 382-390.

D.A. Norman, Approaches to the Study of Intelligence, Artificial Intelligence 47 (1991), pp. 327-346.

G.K.H. Pang, An Expert Adaptive Control Scheme in an Intelligent Process Control System, Proceedings of the 1991 IEEE International Symposium on Intelligent Control (Arlington, VA, August 13-25, 1991), pp. 13-17.

W.J. Pardee, and B. Hayes-Roth, Intelligent Real-Time Control of Material Processing, Technical report, Rockwell International Science Center, Palo Alto Laboratory, Palo Alto, CA, 1987.

J. Park, Toward the Development of a Real-Time Expert System, 1986 Rochester FORTH Conference Proceedings, New York, pp. 23-33.

C.J. Paul, A. Acharya, B. Black, and J.K. Strosnider, Reducing Problem-Solving Variance to Improve Predictability, Communications of the ACM, August 1991, pp. 80-93.

C.J. Paul, L.E. Holloway, D. Yan, J.K. Strosnider, and B.H. Krogh, An Intelligent Reactive Monitoring and Scheduling System, IEEE Control Systems, June 1992, pp. 78-86.

D.W. Payton, Internalized Plans: A Representation for Action Resources, Robotics and Autonomous Systems, Vol. 6, No. 1, 1990, pp. 89-103.

D.W. Payton, J.K. Rosenblatt, and D.M. Keirsey, Plan-Guided Reaction, IEEE Transactions on Systems, Man and Cybernetics, Vol. 20, No. 6, 1990, pp. 1370-1382.

D.W. Payton, and T.E. Bihari, Intelligent Real-Time Control of Robotic Vehicles, Communications of the ACM, August 1991, pp. 48-63.

W.F Punch, III, A.K. Goel and D.C. Brown, A Knowledge-Based Selection Mechanism for Control with Applications in Design, Assembly and Planning, accepted for publication in the International Journal of Artificial Intelligence Tools, 1995.

W.F. Punch III, Dissertation for the degree of Doctor of Philosophy, The Ohio State University, 1989.

M. Rao, T.S. Jiang, and J.P. Tsai, From Knowledge Transfer to Knowledge Generation, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 387-392.

M. Rao, T.S. Jiang, and J.P. Tsai, Software Environment and Architecture for Intelligent Control, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 482-485.

A.A. Rizzi, L.L. Whitcomb, and D.E. Koditschek, Distributed Real-Time Control of a Spatial Robot Juggler, IEEE Computer, May 1992, pp. 12-24.

J.L. Rose, Z.K. Li, C. Thomas, G. Chen, and R. Tupil, Sensor Monitoring and Control Concepts for Intelligent Polymer Processing, Proceedings of the Fifth IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, September 5-7 1990), pp. 1193-1198.

J.K. Rosenblatt, and D.W. Payton, A Fine-Grained Alternative to the Subsumption Architecture for Mobile Robot Control, Proceedings of International Joint Conference on Neural Networks (Washington, D.C., June 1989), pp. 317-324.

E.D. Sacerdoti, Planning in a Hierarchy of Abstraction Spaces, Artificial Intelligence, Vol. 5, 1974, pp. 115-135.

J. Sanborn, and J. Hendler, A Model of Reaction for Planning in Dynamic Environments, International Journal of Artificial Intelligence in Engineering, Vol. 3, No. 2, April 1988, pp. 95-102.

G.N. Saradis, Knowledge Implementation: Structures of Intelligent Control Systems, Proceedings of the 1987 IEEE International Symposium on Intelligent Control, (Philadelphia, Pennsylvania, January 19-20, 1987), IEEE Computer Society Press, pp. 9-17.

C.F. Schmidt, J.L. Goodson, S.C. Marsella, and J.L. Bresina, Reactive Planning Using a "Situation Space", Proceedings of the Annual AI Systems in Government Conference (Washington, D.C., March 27-31, 1989), pp. 50-55.

M.J. Schoppers, Automatic Synthesis of Perception Driven Discrete Event Control Laws, Proceedings of the Fifth IEEE International Symposium on Intelligent Control (Philadelphia, Pennsylvania, January 19-20, 1987), IEEE Computer Society Press.

M.J. Schoppers, In Defense of Reaction Plans as Caches, AI Magazine, Winter 1989, pp. 51-60.

M.J. Schoppers, Universal Plans for Reactive Robots in Unpredictable Environments, Proceedings of the 10th International Joint Conference on AI (August 1987), pp. 1039-1046.

M.J. Schoppers, and T. Linden, The Dimensions of Knowledge-Based Control Systems and the Significance of Meta-Levels, Proceedings of the AAAI Stanford Spring Symposium on Planning in Uncertain, Unpredictable or Changing Environments, J. Hendler, Ed. Tech Rep. 45-90, Systems Research Center, University of Maryland, College Park, 1990.

R.B. Schudy, and C.N. Duarte, Advanced Autonomous Underwater Vehicle Software Architecture, Proceedings of the Symposium on Autonomous Underwater Vehicle Technology (Washington, D.C., June 5-6, 1990), pp. 9-22.

L. Sha, and J.B. Goodenough, Real-Time Scheduling Theory and Ada, IEEE Computer, Vol. 23, 1990, pp. 53-62.

D.D. Sharma, Plan Execution in Dynamic and Unanticipated Environments, 1991 IEEE Conference on AI Applications, pp. 361-367.

K.G. Shin, and X. Cui, Design of a Knowledge-Based Controller for Intelligent Control Systems, IEEE Transactions on Systems, Man and Cybernetics, March/April 1991, pp. 368-375.

R. Shoureshi, and K. Rahmani, Intelligent Control of Building Systems, Intelligent Control Systems, Presented at the Winter Annual Meeting of ASME (December 10-15, 1989, San Francisco, California), pp. 7-15.

T.S. Srivatsan, and S. Chandrashekhar (ed.), Processing and Manufacturing of Composite Materials, ASME, Atlanta, Georgia, Dec. 1-6, 1991.

M. Stefik, Planning with Constraints (MOLGEN: Part 1), Artificial Intelligence, Vol. 16, pp. 111-139.

M. Stefik, Planning and Meta-Planning (MOLGEN: Part 2), Artificial Intelligence, Vol. 16, pp. 141-170.

J. Sticklen, A. Kamel, M.C. Hawley, and V. Adegbite, Fabricating Composite Materials - A Comprehensive Problem-Solving Architecture Based on Generic Tasks, IEEE Expert, April 1992, pp. 43-53.

J. Sticklen, and W.E. Bond, Functional Reasoning and Functional Modeling, IEEE Expert, April, 1991, pp. 20-21.

J. Sticklen, A. Kamel, M.C. Hawley, and J. Delong, An Artificial Intelligence-Based Design Tool for Thin Film Composite Materials, MSU.

J. Sticklen, Problem Solving Architecture at the Knowledge Level, Journal of Experimental and Theoretical Artificial Intelligence, Summer 1989.

J. Sticklen, and B. Chandrasekaran, Integrating Classification-Based Compiled Level Reasoning with Function-Based Deep Level Reasoning, Applied Artificial Intelligence Vol. 3, No. 2, 1989.

A. B. Strong, Fundamentals of Composites Manufacturing: Materials, Methods and Applications, Society of Manufacturing Engineers, 1989.

M. Tambe, and P.S. Rosenbloom, A Framework for Investigating Production System Formulations with Polynomially Bounded Match, Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI -'90), August 1990, pp. 693-700.

J. Tenenberg, Planning with Abstraction, Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI '86), pp. 76-80.

E. Torgerson, and F.W. Paul, A Hierarchical Semi-Autonomous Control Scheme for Apparel Manufacturing, Intelligent Control Systems 1990, Presented at the Winter Annual Meeting of ASME (November 25-30, 1990, Dallas, Texas), pp. 57-63.

E. Tse, and M. Fehling, An Intelligent Control System, Proceedings of the 1989 IEEE International Symposium on Intelligent Control, (Albany, New York, September 25-26, 1989), IEEE Computer Society Press, pp. 194-199.

R.M. Turner, A Schema-Based Model of Adaptive Problem Solving, Doctoral Dissertation, Georgia Institute of Technology, GIT-ICS-89/42, December, 1989, 229 pages.

L. Tychonievich, T.C. Smith, and R. Evans, Influence Networks: A Reactive Planning Architecture, 1991 IEEE Conference on AI Applications, pp. 354-360.

A. Valente, Knowledge-Level Analysis of Planning Systems, SIGART Bulletin, Vol. 6, No. 1, January, 1995, pp. 33-41.

W. Van deVelde, Issues in Knowledge Level Modelling, Second Generation Expert Systems, Jean-Marc David, et. al. (editors), Springer-Verlag, Berlin, 1993.

S.M. Walsh, Artificial Intelligence: Its Application to Composite Processing, 35th International SAMPE Symposium, April 2-5, 1990, pp. 1280-1291.

J. Wei, J.D. DeLong, M. DeMeuse and M.C. Hawley, Comparison of Microwave and Thermal Cure of Epoxy Resins, Polymer Engineering and Science, Vol 33, No. 17, 1993, pp. 1132-1149.

M.P. Wellman, The STRIPS Assumption for Planning Under Uncertainty, Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI '90), pp. 198-203.

J.E. Whitlow, and K.A. Debelak, A Knowledge-Based Structure for Process Control, Proceedings of the 1989 American Control Conference, (Pittsburgh, Pennsylvania, June 21-23, 1989), American Automatic Control Council, pp. 1354-1359.

D.E. Wilkens, Domain-independent Planning: Representation and Plan Generation, Artificial Intelligence, Vol. 22, 1984, pp. 269-301.

D.J. Williams, A.A. West, and C.J. Hinde, A Hierarchical Model of Intelligent Materials Processing, 1991 IEEE International Symposium on Intelligent Control, pp. 330-333.

L. Winkelbauer, K. Fedra, ALEX: Automatic Learning in Expert Systems, Proceedings of the 1991 IEEE Conference on AI Applications, pp. 59-62.

M. Winslett, Validating Generalized Plans in the Presence of Incomplete Information, Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI '87), pp. 261-266.

P.H. Winston, Artificial Intelligence, Addison-Wesley, Reading, Massachusetts, Third edition, 1992.

H.T. Wu, and B. Joseph, Model-Based and Knowledge-Based Control of Pultrusion Processes, SAMPE Journal, Vol. 26, No. 6, November/December 1990, pp. 59-70.

H.T. Wu, and B. Joseph, Knowledge-Based Control of Autoclave Curing of Composites, SAMPE Journal, Vol. 26, No. 6, November/December 1990, pp. 39-54.

Q. Yang, and J.D. Tenenberg, ABTWEAK: Abstracting a Nonlinear Least Commitment Planner, Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI '90), pp. 204-209.

X. Zheng, and S. Srivastava, Events and Actions: An Object-Oriented Approach to Real-Time Control Systems, Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (June 1-2, 1989), pp. 269-272.

S. Zilberstein, The Utility of Planning, SIGART Bulletin, Vol. 6, No. 1, January, 1995, pp. 42-47.